

*Unbounded circumambient patterns in segmental phonology**

Adam G. McCollum

Rutgers University

Eric Baković

Anna Mai

Eric Meinhardt

University of California San Diego

We present an empirical challenge to Jardine’s (2016) assertion that only tonal spreading patterns can be unbounded circumambient, meaning that the determination of a phonological value may depend on information that is an unbounded distance away on both sides. We focus on a demonstration that the ATR harmony pattern found in Tutrugbu is unbounded circumambient, and we also cite several other segmental spreading processes with the same general character. We discuss implications for the complexity of phonology and for the relationship between the explanation of typology and the evaluation of phonological theories.

1 Introduction

Phonological theories (of the ‘free-standing’ variety; Prince 2007) make typological claims, predictions about what are and what are not possible phonological patterns. Much work has been devoted to questions concerning the expressivity of a given phonological theory. Is the theory powerful enough to correctly predict the existence of a given attested pattern?; if not, how can its expressive power be appropriately extended? Is the theory so powerful that it incorrectly predicts the existence of an unattested, potentially ‘pathological’ pattern?; if so, how can its expressive power be appropriately constrained?

Recent work has approached this question from the perspective of formal language theory, aiming to characterise particular attested or

* E-mail: ADAM.MCCOLLUM@RUTGERS.EDU, EBAKOVIC@UCSD.EDU, ACMAI@UCSD.EDU, EMEINHARDT@UCSD.EDU.

We would like to thank Judith Glover for her willingness to share her language, and James Essegbey for all his help. We are also grateful to the reviewers at *Phonology*, and to Sharon Rose, Jeff Heinz, Andrew Lamont, Charlie O’Hara and Caitlin Smith, as well as members of audiences at NELS 48, SCAMP 3 and the University of Chicago. Any errors are our own.

unattested phonological patterns in terms of well-defined, theory-independent computational classes of string sets (formal languages) or mappings between string sets (string functions): any given pattern can be characterised by the minimum level of complexity (expressivity) required to capture it. The larger goal of this work is to formally delimit the boundary between possible and impossible grammatical patterns, as evidenced by attested and unattested language patterns. Heinz (2011a, b, 2018), Heinz & Lai (2013) and Jardine (2016), among others, have advanced the broad hypothesis that phonological patterns are SUBREGULAR, meaning that they can be characterised by computational classes that are strictly less expressive than the class of regular string relations, which characterises the expressivity of ordered *SPE*-style rewrite rules (Johnson 1972, Kaplan & Kay 1994). Figure 1 summarises some subregular function classes and gives some relevant examples, including the one that is the empirical focus of this article, Tutrugbu ATR harmony (McCollum & Essegbey 2018, Essegbey 2019).

Chandlee (2014: §7.1) and Heinz (2018: §6.2.1) note that the vast majority of phonological patterns appear to be contained within the innermost, least expressive classes of the subregular hierarchy. However, there are some patterns that are relatively more complex, and Heinz & Lai (2013) and Jardine (2016) contend that the most complex segmental patterns require no more expressivity than what is afforded by WEAKLY DETERMINISTIC regular functions. These types of patterns include prototypical bidirectional harmony processes. In contrast, Jardine (2016) demonstrates that a number of tonal spreading patterns are more complex than this. These patterns all require at least NON-DETERMINISTIC regular functions to be described, and all exhibit UNBOUNDED CIRCUMAMBIENCE, a dependence on information (e.g. spreading triggers and blocking conditions) that are an unbounded distance from the spreading target in both directions.¹

The unbounded circumambient nature of the Tutrugbu ATR harmony pattern can be seen in the comparison between the pairs of forms in (1). [ATR] spreads leftwards from the root to prefixes, targeting both [+high] and [-high] vowels (1a, b). When the initial prefix vowel is [+high], however, harmony is blocked by [-high] vowels (1d). In other words, the surface realisation of the medial vowels in (1) depends on both the [ATR] value of the root and the presence or absence of an initial-syllable [+high] vowel. We demonstrate below that this pattern is non-deterministic in exactly the same way that the tonal spreading

¹ In what are perhaps more familiar terms, a subset of these unbounded circumambient patterns are non-myopic (Wilson 2003, 2006; see also Finley 2008, Walker 2010, Kimper 2012, Mascaró 2019, McCollum *et al.* 2020, Stanton 2020). In the descriptively most straightforward hypothetical case of a non-myopic pattern, spreading is simply not triggered when there is a blocking segment anywhere in the form; this is dubbed sour grapes by Wilson (2003, 2006), adapting a term from Padgett (1995). The Tutrugbu ATR harmony pattern described below is a variation on the sour grapes pattern.

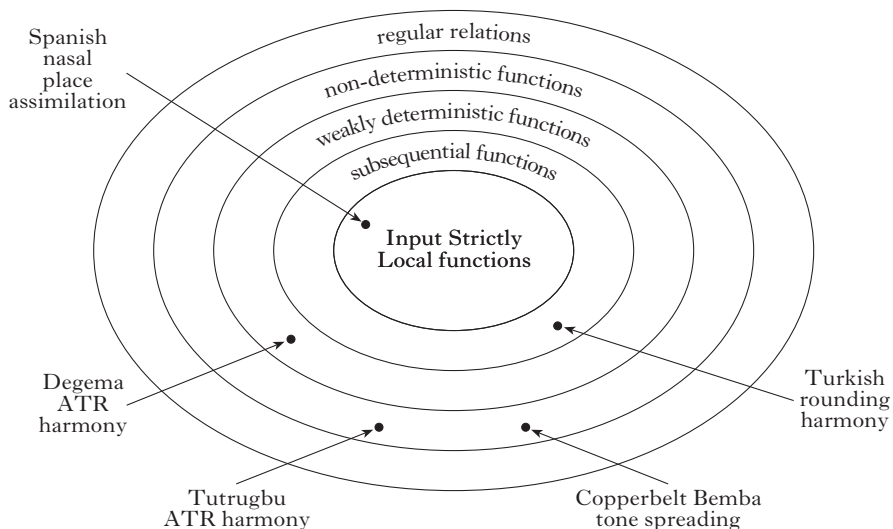


Figure 1

Some regions of the subregular hierarchy of string-to-string mappings, with phonological examples. Examples of Input Strictly Local functions include nasal place assimilation and metathesis (Chandlee 2014); examples of subsequential functions include unidirectional harmony (Heinz & Lai 2013) and dissimilation (Payne 2017); examples of weakly deterministic functions include bidirectional harmony (Heinz & Lai 2013); examples of non-deterministic functions include unbounded circumambient processes (Jardine 2016); examples of regular relations include optional iterativity (Heinz 2018).

patterns discussed by Jardine (2016) are, because this pattern also satisfies Jardine’s criteria for unbounded circumambience.²

(1) *Tutrugbu ATR harmony*

- | | | | |
|------------------|-----------|------------------|-----------|
| a. /a-tí-ba-bá/ | [atíbabá] | c. /ɪ-tí-ba-bá/ | [ítíbabá] |
| 3SG-NEG-FUT-come | | 1SG-NEG-FUT-come | |
| b. /a-tí-ba-ǰē/ | [etibeǰē] | d. /ɪ-tí-ba-ǰē/ | [ítíbaǰē] |
| 3SG-NEG-FUT-grow | | 1SG-NEG-FUT-grow | |

We introduce ATR harmony in Tutrugbu in §2, and in §3 we show that it satisfies Jardine’s definition of an unbounded circumambient process.

² The following abbreviations are used, in addition to those in the Leipzig glossing rules (Comrie *et al.* 2015): ASSOC = associative, AUG = augmentative, CL = class marker, CONT = continuative, DEP = dependent pronoun, EPI = epipatetic, EXC = excessive, FREQ = frequentive, FV = final vowel, GER = gerundial, HAB = habitual, ITV = itive, K = ‘moveable k’ (Dimmendaal 1983), OM = object marker, REV = reversive, SE = ‘stem-enlarging vowel’ (Novelli 1985), SM = subject marker, TMA = tense, mood and aspect, VENT = ventive and VOI = voice.

We introduce the computational machinery necessary to account for unbounded circumambient processes in §4, providing examples for three complexity classes within the subregular hierarchy: subsequential, weakly deterministic and non-deterministic. We provide a finite-state analysis of Tutrugbu ATR harmony in §5, and show that it is non-deterministic. In §6 we discuss other examples of unbounded circumambient segmental patterns, and the implications of our findings for explaining the typology of complexity and evaluating phonological theories, before concluding in §7.

2 Tutrugbu ATR harmony

Tutrugbu is a Ghana-Togo Mountain language (Kwa), spoken in south-eastern Ghana. The data for this article come from formal elicitation, as well as a documentary corpus of natural speech. Data-collection practices are described in Essegbey (2019: 11–13). Almost all of the data presented here were collected by James Essegbey over the course of around 15 years of fieldwork in Ghana, with only recent collaboration with the first author of this paper. The pattern described here was first presented in Essegbey (2009), and more extensively discussed in McCollum & Essegbey (2018) and Essegbey (2019: 36–39). Example data are found throughout Essegbey (2019), and, to the best of our knowledge, are representative of the entire speech community, without exceptions or variation, except as described in §2.2.³

2.1 Data

Tutrugbu has an inventory of nine oral vowels, /a ɔ o u ε e i i/, with contrasts in height, backness, rounding and ATR. As Essegbey (2009) notes, there are only seven surface vowels in the language. The vowels we transcribe as high [–ATR] vowels, /ʊ/ and /ɪ/, always surface as the mid vowels [ɔ] and [ε], but pattern as high vowels. To reflect their phonological status, we will transcribe them as [ʊ] and [ɪ] throughout (see McCollum & Essegbey 2018, 2020 for a different approach). We defend this abstract analysis in §2.3.2 below. Nasal counterparts of these oral vowels are also phonemic in the language, and behave just like oral vowels with respect to ATR harmony. ATR harmony proceeds from right to left in Tutrugbu, from roots to prefixes. Suffixation is rare, and suffixes do not generally undergo harmony.

Observe the ATR pairings demonstrated by regressive harmony on noun class prefixes in (2). In (2a), prefixal [a] alternates with [e], while in (b), [ɔ] alternates with [o]. In (c), [ɪ] alternates with [i], and finally, in (d), [ʊ] alternates with [u]. Note that [ε] does not occur in affixes. We

³ The same pattern is also attested in Tafi, a closely related language (Bobuafor 2013; see also §6.2.1).

analyse [+ATR] as the active (or dominant) feature value in the language, and assign affixes a [-ATR] value underlyingly.

(2) *Harmony in nouns*

| | | | |
|----|----------------------------|---------------------------------|--------|
| | [-ATR] | | [+ATR] |
| a. | /a-ɲɪ/ [aɲɪ] CL1-man | /a-bú/ [ebú] CL1-dog | |
| b. | /ɔ-da/ [ɔda] CL3-copper | /ɔ-pétē/ [opétē] CL3-vulture | |
| c. | /ɪ-da/ [ɪda] CL4-copper | /ɪ-pétē/ [ipétē] CL4-vulture | |
| d. | /bɔ-wɪ/ [bɔwɪ] CL8-axe | /bɔ-ju/ [buju] CL8-war | |

Noun-class prefixes undergo ATR harmony, but concatenating more prefixes to a nominal root is not possible. Verbal morphology, on the other hand, allows for more complexity. In (3) we see that words with only [+high] prefix vowels show full harmony. In (3a–c), [-ATR] roots are preceded by [-ATR] prefixes, while in (3d–f), [+ATR] roots are preceded by [+ATR] prefixes.

(3) *Harmony on [+high] prefixes*

| | | | |
|----|-------------------------------------|----|-------------------------------------|
| | [-ATR] | | [+ATR] |
| a. | /ɪ-tí-bá/ [ítíbá] 1SG-NEG-come | d. | /ɪ-tí-ǰē/ [ítíǰē] 1SG-NEG-grow |
| b. | /bɔ-tí-bá/ [butíbá] 1PL-NEG-come | e. | /bɔ-tí-ǰē/ [butíǰē] 1PL-NEG-grow |
| c. | /kɪ-tí-bá/ [kítíbá] CL5-NEG-come | f. | /kɪ-tí-ǰē/ [kítíǰē] CL5-NEG-grow |

Full harmony also obtains when all prefix vowels are [-high], as in (4). In (4d–f), [+ATR] roots propagate their [+ATR] feature to the left edge of the word. Note also in (4c, f) that [-high] round vowels in the initial syllable trigger progressive rounding harmony in the preverbal domain (McCollum & Essegbey 2020). Rounding harmony is triggered by and targets [-high] vowels; [+high] vowels are transparent.

(4) *Harmony on [-high] prefixes*

| | | | |
|----|-------------------------------------|----|-------------------------------------|
| | [-ATR] | | [+ATR] |
| a. | /a-ba-bá/ [ababá] 3SG-FUT-come | d. | /a-ba-ǰē/ [ebeǰē] 3SG-FUT-grow |
| b. | /ka-ba-bá/ [kababá] CL7-FUT-come | e. | /ka-ba-ǰē/ [kebeǰē] CL7-FUT-grow |
| c. | /ɔ-ba-bá/ [ɔbɔbá] 2SG-FUT-come | f. | /ɔ-ba-ǰē/ [obɔǰē] 2SG-FUT-grow |

In (3) and (4), all prefix vowels agree in [\pm high]. Harmony in forms with prefix vowels of differing values of [high] is illustrated in (5). The initial-syllable vowel is [$-$ high] in these examples, as it is in (4), and harmony obtains throughout the word.

(5) *Harmony on varying height prefixes with a [$-$ high] initial prefix*

| | | | |
|----|--|----|--|
| | [$-$ ATR] | | [+ATR] |
| a. | /a-tí-ba-bá/ [atíbabá] 3SG-NEG-FUT-come | d. | /a-tí-ba-ǰē/ [etíbeǰē] 3SG-NEG-FUT-grow |
| b. | /ka-tí-ba-bá/ [katíbabá] CL7-NEG-FUT-come | e. | /ka-tí-ba-ǰē/ [ketíbeǰē] CL7-NEG-FUT-grow |
| c. | /ɔ-tí-ba-bá/ [ɔtíbɔbá] 2SG-Z-NEG-FUT-come | f. | /ɔ-tí-ba-ǰē/ [otíboǰē] 2SG-NEG-FUT-grow |

In contrast to the previous examples, [$-$ high] vowels block harmony if the vowel in the initial syllable is [$+$ high]. Regardless of the root's [ATR] value, the vowel of the [$-$ high] future prefix and all preceding vowels surface as [$-$ ATR] under this condition, as shown in (6).

(6) *(Dis)harmony on varying height prefixes with a [$+$ high] initial prefix*

| | | | |
|----|-------------------------------------|----|-------------------------------------|
| | [$-$ ATR] | | [+ATR] |
| a. | /ɪ-ba-bá/ [ɪbabá] 1SG-FUT-come | d. | /ɪ-ba-ǰē/ [ɪbaǰē] 1SG-FUT-grow |
| b. | /bu-ba-bá/ [bubabá] 1PL-FUT-come | e. | /bu-ba-ǰē/ [bubaǰē] 1PL-FUT-grow |
| c. | /kɪ-ba-bá/ [kɪbabá] CL5-FUT-come | f. | /kɪ-ba-ǰē/ [kɪbaǰē] CL5-FUT-grow |

These examples show that the [$-$ high] vowel blocks harmony when the initial-syllable vowel is [$+$ high]. The [$-$ high] vowel immediately precedes the root in (6), but Tutrugbu does allow at least one [$+$ high] prefix to intervene between a root and a [$-$ high] vowel. In words with a [+ATR] root, a [$+$ high] initial-syllable vowel and a medial [$-$ high] prefix vowel (satisfying the two conditions necessary to block harmony), a [$+$ high] prefix vowel intervening between the root and the medial [$-$ high] vowel undergoes harmony. In (7a) and (b), the itive prefix alternates based on the [ATR] value of the root, establishing that this particular morpheme regularly undergoes harmony. In (7c) and (d), this prefix is the only one to undergo harmony. In essence, harmony spreads as far as the blocking [$-$ high] vowel, and then stops.

(7) *Partial harmony on [+high] prefixes*

| | | | |
|---------------|---------|-------------------|------------|
| a. /i-dɪ-tó/ | [ɪdɪtó] | c. /ɪ-ba-dɪ-wu/ | [ɪbadiwu] |
| 1SG-ITV-cook | | 1SG-FUT-ITV-climb | |
| b. /i-dɪ-wu/ | [ɪdiwu] | d. /bɔ-ba-dɪ-wu/ | [bɔbadiwu] |
| 1SG-ITV-climb | | 1PL-FUT-ITV-climb | |

The data above show that [–high] vowels are **CONDITIONAL BLOCKERS**: they block harmony only in the presence of an initial-syllable [+high] vowel. Two [+high] vowel prefixes do not block harmony, as in (3), and two [–high] vowel prefixes do not block harmony, as in (4) and (5). It is only the combination of an initial-syllable [+high] vowel and a medial [–high] vowel that blocks harmony. In other words, the realisation of a [–high] prefix vowel depends not only on the [ATR] value of the vowel in the immediately following morpheme (the root, or a prefix closer to the root), but also on the [high] value of the initial-syllable vowel.

In (6) and (7), the initial-syllable [+high] vowel and the medial [–high] vowel are in adjacent syllables. In (8) we see that harmony is blocked by the co-occurrence of these two conditions, even when separated by intervening syllables.

(8) *Unbounded conditional blocking of harmony*

| | | |
|-----------------------|----------------------------------|----------------|
| a. no #[+high] prefix | /a-atí-wu/ | [etíwu] |
| | 3SG-NEG-climb | |
| b. no [–high] prefix | /i-tí-wu/ | [itíwu] |
| | 1SG-NEG-climb | |
| c. adjacent | /ɪ-ba-wu/ | [ɪbawu] |
| | 1SG-FUT-climb | |
| d. 1-syllable gap | /ɪ-tí-ka-wu/ | [ɪtíkawu] |
| | 1SG-NEG-PFV-climb | |
| e. 2-syllable gap | /ɪ-tí-ka-á-wu/ | [ɪtíkaáwū] |
| | 1SG-NEG-PFV-PROG-climb | |
| f. 3-syllable gap | /ɪ-tí-ka-á-ba-wu/ | [ɪtíkaábāwū] |
| | 1SG-NEG-PFV-PROG-VENT-climb | |
| g. 4-syllable gap | /ɪ-tí-ka-á-ba-ba-wu/ | [ɪtíkaábābāwū] |
| | 1SG-NEG-PFV-PROG-VENT-VENT-climb | |
| h. no #[+high] prefix | /a-tí-ka-á-ba-ba-wu/ | [etíkeébēbēwū] |
| | 3SG-NEG-PFV-PROG-VENT-VENT-climb | |

ATR harmony is blocked only when two conditions are met: (i) the initial-syllable vowel is [+high], and (ii) another prefix vowel is [–high], as in (c)–(g). When only one of these conditions is met, as in (a), (b) and (h), harmony obtains. In (c), the [+high] initial-syllable vowel and the [–high] prefix vowel are adjacent, and harmony fails. In (d)–(g), one, two, three and four syllables intervene between these two interacting conditions on harmony. Thus the blocking of regressive ATR harmony

depends on decidedly non-local information – the [high] value of the initial-syllable vowel – and the presence of a [–high] prefix vowel, which may occur a number of syllables from the initial syllable, with no principled upper bound.

2.2 Variation

Before we move on to the analysis, there is an additional aspect of the pattern worth noting. Essegbey (2009: 40) describes variation in the blocking context. When an initial-syllable [+high] prefix is followed by a medial [–high] prefix, as in (9a), the [–high] prefix and all preceding prefixes may surface as [–ATR]. But the initial-syllable [+high] prefix may also surface as [+ATR], even though the following [–high] prefix is [–ATR]. Finally, as in (9b), the [–high] vowel undergoes harmony when the initial-syllable vowel is [–high], just as in our data above, with no variation. (As noted in §2.1, we use /ɪ/ and [ɪ] rather than Essegbey's /ɛ/ and [ɛ].)

- (9) a. /ɪ-ba-tɛ́nú vɪ/ [ɪbatɛ́nú vɪ] ~ [ibatɛ́nú vɪ]
 1_{SG-FUT}-be.able go
 b. /a-ba-tɛ́nú vɪ/ [ɛbetɛ́nú vɪ]
 3_{SG-FUT}-be.able go

In the pattern most widely attested in our data, the medial [–high] vowel conditionally blocks harmony, preventing [+ATR] from spreading to prefixes further from the root. For the alternative form, Essegbey suggests that harmony skips the [–high] vowel to target the initial-syllable [+high] vowel. This in turn suggests that, for these speakers, the medial [–high] vowel is (optionally) CONDITIONALLY TRANSPARENT. We have very few data on conditional transparency when more than one [+high] vowel occurs to the left of medial /a/, but preliminary data suggest that all [+high] vowels are realised as [+ATR] for speakers exhibiting conditional transparency. In any event, note that for both patterns the realisation of medial /a/ depends on both the initial-syllable vowel to its left and the root vowel to its right, even when either or both of these dependencies is long-distance.

2.3 Analysis

2.3.1 *Active value.* Determining which feature value is active, as van der Hulst & van de Weijer (1995: 504) note, can be challenging in patterns of root-controlled harmony. That being said, several facts point to [+ATR] being the active feature value in Tutrugbu (see also Casali 2012 on diagnosing feature-value activity). First, when a vowel is not assimilated to the [ATR] value of the root, it always surfaces as [–ATR]. In (8c–g), for example, all prefix vowels surface as [–ATR], because both parts of the blocking condition are present. Additionally, morphemes to the right of

the root do not generally undergo harmony, surfacing as [-ATR]. In (10a, b), the plural suffix, which may only attach to human roots, surfaces as [-ali], regardless of the root's [ATR] value. In (10c-f), the definite enclitic surfaces as [=í], irrespective of the root's [ATR] value.

- | | | | | | |
|---------|---------------|----------|----|--------------|---------|
| (10) a. | /a-ka-ali/ | [akaali] | d. | /a-ɲí=í/ | [aɲí] |
| | CL1-father-PL | | | CL1-man-DEF | |
| b. | /a-lě-ali/ | [elěali] | e. | /ɔ-sí=í/ | [osí] |
| | CL1-mother-PL | | | CL3-tree-DEF | |
| c. | /a-dzĩ=í/ | [adzĩí] | f. | /ka-dě=í/ | [kedéí] |
| | CL1-woman-DEF | | | CL7-back-DEF | |

In other words, when harmony does not apply, affix vowels surface as [-ATR], which is good evidence that [+ATR] is active. For Akan, Casali (2012) suggests that cross-word harmony and the form of the independent pronouns provide further support for the activity of [+ATR]. As in Akan, [+ATR] may optionally spread across word boundaries in Tutrugbu, as in (11).

- | | | |
|---------|-----------------------------------|---------------------------------|
| (11) a. | /gɪ bɔ-b ^h ɪtɪ kɪ-bwi/ | [gɪ bɔb ^h ɪtɪ kibwi] |
| | REL 1PL-do CL5-issue | |
| b. | /b ^h ɪtɪ mɪ kɪ-bwi/ | [b ^h ɪtɪ mi kibwi] |
| | do 1SG.OBJ CL5-issue | |
| c. | /kɪ-b ^h ɪtɪ kɪ-bwi/ | [kɪb ^h ɪtɪ kibwi] |
| | CL5-do CL5-issue | |
| d. | /pɪ a-sɪ ba-baɖi/ | [pɪ asɪ bebedi] |
| | then 3SG-say 3PL-FUT-look | |
| e. | /ɪ-sɪ i-ɖi ko/ | [isi iɖi ko] |
| | 1SG-say 1SG-look just | |

As far as we are aware, this phrasal ATR harmony may only target [+high] vowels. In (11a) the [+ATR] root /bwi/ triggers harmony on the prefix /kɪ-, and in (b) the same root triggers harmony on both its prefix and the preceding object pronoun, /mɪ/. Observe that the verb /b^hɪtɪ/ is unaffected by phrasal harmony here, and is realised with its underlyingly [-ATR] specification. And yet, in (11c), /bwi/ triggers harmony on its prefix and the preceding verb, /b^hɪtɪ/. In this context, the underlyingly [-ATR] vowels of the verb are assimilated to the [+ATR] value of /bwi/. Similarly, in (d), the verb /sɪ/ is unaffected by the [ATR] value of the following root, surfacing as [-ATR]. In (11e), though, this verb undergoes harmony from the following verb root /ɖi/, and is pronounced as [sɪ]. In both sets of examples, verbs that are underlyingly [-ATR] optionally undergo phrasal [+ATR] harmony. While phrasal [+ATR] spreading is widespread in our data, [-ATR] spreading is unattested.

Casali (2012) also considers the independent form of the personal pronouns to be indicative of the underlying [ATR] value of harmonic affixes more generally. Independent pronouns occur as separate words, while their dependent counterparts occur as prefixes. In Akan, the independent forms are all [-ATR], supporting the claim that [-ATR] is the unmarked underlying form of alternating affixes. This same generalisation holds for Tutrugbu. Independent pronouns are very similar to their dependent (prefixal) counterparts in segmental form, and in all cases surface with [-ATR] vowels, as in (12). According to Casali's diagnostics, the [-ATR] status of the independent pronouns lends further support to [+ATR] spreading.

| (12) <i>independent</i> | <i>dependent</i> | | <i>independent</i> | <i>dependent</i> | |
|-------------------------|------------------|-----|--------------------|------------------|-----|
| [mɪ] | [ɪ] ~ [i] | 1SG | [bʊɔ] | [bʊ] ~ [bu] | 1PL |
| [wɔ] | [ɔ] ~ [o] | 2SG | [wɔnɔ] | [nɔ] ~ [no] | 2PL |
| [jɪ] | [a] ~ [e] | 3SG | [balɪ] | [ba] ~ [be] | 3PL |

2.3.2 *Abstract vowels.* As noted in §2.1, we assume that Tutrugbu has a nine-vowel inventory requiring two abstract [+high, -ATR] vowels, /ɪ ʊ/, which are phonetically realised as mid [ɛ ɔ], but which we consistently represent as [ɪ ʊ]. This abstract analysis is justified for several reasons.

First, these vowels pattern consistently as [+high] with respect to rounding harmony. In (5b, e), even the [-ATR] allomorph of the negation prefix, which is phonetically mid (but transcribed in (5b) as [tɪ]), does not undergo progressive rounding harmony, and we can attribute this to the fact that rounding harmony only targets [-high] vowels. Rounding harmony is also only triggered by [-high] vowels, so the abstract high vowel /ʊ/ of the 1st person plural prefix does not trigger rounding harmony on the future prefix in (6b) or (7d), despite being phonetically mid.

Along with their inertness as triggers and targets of rounding harmony, the phonological status of these vowels is evident in their actual ATR alternations. For example, the 2nd person singular prefix alternates between phonetic [ɔ] and [o], while the 1st person plural prefix alternates between phonetic [bɔ] and [bu] (our phonological surface [bʊ] and [bu]). If, as we argued in §2.3.1, [-ATR] is the underlying feature value, its [+ATR] counterpart is not able to be inferred from its faithful phonetic realisation. Specifically, which phonetic [ɔ] alternates with [o], and which alternates with [u]? This problem appears insurmountable if both surface [ɔ] vowels are treated equivalently.

We have thus chosen to mark the vowels that alternate with [+high, +ATR] vowels as [+high], but this is not the only possible analysis. One alternative is to specify the underlying vowels we have analysed as /ɪ ʊ/ as [+ATR] /i u/. Possible support for this alternative comes from speakers who (optionally) produce the pattern of conditional transparency

described in §2.2. Recall that when a medial [–high] vowel and an initial-syllable [+high] vowel co-occur, these speakers may produce the [+high] vowel(s) as [+ATR], despite the [–high] vowel surfacing as [–ATR] [a]. A possible alternative to the conditional transparency analysis is that in this context, prefix vowels to the left of blocking [a] may surface in accordance with their underlying [ATR] specifications. For these speakers at least, there would be a seven-vowel system, with [+high] vowels underlyingly specified as [+ATR], in contrast with the [–high] vowels. Note that, regardless of whether the conditional transparency analysis or this alternative is correct, the generalisation holds that harmony is blocked when an initial-syllable [+high] vowel co-occurs with a medial [–high] vowel, in unbounded circumambient fashion.

A seven-vowel analysis with underlying /i u/ cannot account for the phrasal [+ATR] harmony described in (11), however. As noted above, in all these examples the verb roots are clearly [–ATR], but may surface as [+ATR] if followed by a [+ATR] word. Since these are roots, not affixes, we cannot simply stipulate that the vowels in a word like /b^hɪtɪ/ are [+ATR]. This verb root is always preceded by [–ATR] prefixes in our data, even in (11c), where it undergoes phrasal [+ATR] harmony. This brings us back to the original point: we cannot infer the [+ATR] counterpart of the vowels in question if they are underlyingly specified as [–high, –ATR]. If the vowels in /b^hɪtɪ/ are simply [–high, –ATR], then why do these vowels surface as [i], and not [e], when affected by phrasal [+ATR] harmony? To account for these facts, we must analyse these vowels as underlyingly [+high]. Since phrasal harmony prevents us from assuming that they are [+high, +ATR], the best analysis requires a nine-vowel inventory, with two abstract [+high, –ATR] vowels.

One final typological point is worth mentioning here. As an anonymous reviewer notes, African languages in which the [+high] vowels exhibit ATR contrasts show almost exclusive [+ATR] dominance, whereas languages in which the [+high] vowels exhibit no [ATR] contrasts show no clear typological tendencies (Casali 2003, 2008, 2012, Rose 2018). Despite historical changes to its vowel inventory, Tutrugbu vowels pattern like those of neighbouring Tafi, which has maintained what we take to be the proto-language's nine-vowel inventory, with an ATR contrast among the high vowels (see also Essegbey 2009: 40). The structure of the [–high] portion of the inventory is somewhat distinct from canonical African nine-vowel systems: the low vowel /a/ pairs with mid /e/, whereas in most nine-vowel systems the low vowel has no harmonic counterpart. If the presence of contrastive [ATR] values for the [+high] vowels is predictive of feature-value activity, our nine-vowel analysis is consistent with the expected dominance of [+ATR] in an African language with nine phonological vowels.

3 Unbounded circumambience

The Tutrugbu ATR harmony pattern described in the previous section instantiates what Jardine (2016) dubs an UNBOUNDED CIRCUMAMBIENT process, defined in (13).

- (13) An unbounded circumambient process is a process:
- a. whose application is dependent on information (e.g. the presence of a trigger or blocker) on both sides of the target [CIRCUMAMBIENT];
 - b. in which, on both sides, there is no bound on how far this information may be from the target [UNBOUNDED].

In Tutrugbu, the surface quality of prefix vowels is dependent on three factors: the [ATR] value of the root, the height of the initial prefix vowel and the presence or absence of a medial non-high prefix vowel. As shown in (8), there appears to be no bound on the potential distance between these three factors; the only constraint on their distance is the productivity of the language's morphology. For these reasons, Tutrugbu ATR harmony satisfies the criteria for unbounded circumambience.

Compared to the patterns previously identified as unbounded circumambient, Tutrugbu ATR harmony stands out as a segmental, rather than a tonal process. In fact, for Jardine (2016), one reason to single out unbounded circumambient processes as a class is to characterise a typological asymmetry between tonal and segmental phenomena, namely that unbounded circumambience is more prevalent (if not unique) among tonal processes than among segmental processes. Jardine attributes this asymmetry to the computational expressivity required to describe unbounded circumambient patterns. He argues that tonal patterns are capable of commanding greater computational expressivity than segmental patterns, and concludes that tone should be viewed as an exception to hypotheses concerning the upper bound on the expressivity of phonology. In the following section, we briefly review these hypotheses, and introduce the computational machinery necessary to capture unbounded circumambient processes.

4 Computational requirements for unbounded circumambience

The expressivity of phonological grammars comprising ordered *SPE*-style rewrite rules (Chomsky & Halle 1968) corresponds to the class of regular string relations (Johnson 1972, Kaplan & Kay 1994). However, many phonological input–output maps do not require the full expressivity of regular relations. Synchronically attested metathesis, partial reduplication, consonant harmony and dissimilation patterns are all SUBSEQUENTIAL (Chandlee *et al.* 2012, Chandlee & Heinz 2012, Luo 2017, Payne 2017). Moreover, Chandlee's (2014) analysis of PBase (Mielke 2007), a database

of phonological patterns, shows that most are describable using a highly restricted subclass of the subsequential functions.⁴ Finally, the vowel-harmony patterns analysed in Nevins (2010) have also been shown to be subsequential (Gainor *et al.* 2012), with the sole exception of bidirectional harmony in Woleaian. For this reason, Gainor *et al.* argue that a more restrictive hypothesis better captures the minimum computational expressivity necessary to describe phonological patterns: the SUBSEQUENTIAL HYPOTHESIS (see also Heinz 2018).

Under this hypothesis, the most complex computational machinery necessary to model phonological patterns is that of subsequential functions (perhaps first proposed by Mohri 1997: 279). However, Heinz & Lai (2013) demonstrate that bidirectional stem-controlled and dominant-recessive vowel-harmony patterns are not subsequential, and propose instead that phonological patterns are at most weakly deterministic, a class they tentatively define in terms of a restriction on compositions of two subsequential functions.

Jardine (2016) discusses a number of tonal patterns that are neither subsequential nor weakly deterministic, and shows that this class of unbounded circumambient tonal patterns requires the strictly greater expressivity of non-deterministic regular functions. Most significantly in the present context, Jardine argues that there is a computational difference between tonal and segmental phonology: while segmental phonology may be restricted to the class of subsequential functions and their weakly deterministic compositions, tonal phonology requires the expressive power of non-deterministic regular functions.

Regular relations and their subclasses can be represented in many ways, including ordered *SPE*-style rewrite rules, suitably restricted Optimality Theory grammars (Riggle 2004) and finite-state transducers (FSTs). In this paper we represent such mappings as FSTs, for two related reasons: much work characterising regular relations and their subclasses uses this formalism, and consequently, so does existing work on the complexity of phonology. FSTs represent string-to-string functions in a way that emphasises incremental calculation, proceeding symbol by symbol from one end of the input string to the other, and restricting the amount of information about the observed prefix of the input that can be ‘remembered’ while deciding what to output at each step. Such differences in the amount and type of information a transducer is able to store can define more *vs.* less expressive classes of transducers, which can in turn be used to describe differences in the amount and type of information required to express different phonological patterns. For more detailed introductions to FSTs and their relation to phonology, we refer the reader to Kaplan & Kay (1994), Mohri (1997) and Jardine (2016). In the

⁴ As Chandlee (2014) notes, the PBase is a set of analyses, not a typologically representative set of cross-linguistic data – but it is the closest thing to such a typologically representative set that is currently available.

next three subsections we review definitions, key properties and attendant phonological examples of subsequential functions, weakly deterministic regular functions and non-deterministic regular functions.

4.1 Subsequentiality

A subsequential FST τ can be defined by the seven parameters in (14).

- (14) a. Q a finite set of states
 b. Σ a finite set of symbols representing the input alphabet
 c. Δ a finite set of symbols representing the output alphabet
 d. $q_0 \in Q$ a distinguished initial state
 e. $\delta: Q \times \Sigma \rightarrow Q$ a state-transition function
 f. $\omega: Q \times \Sigma \rightarrow \Delta^*$ an output function⁵
 g. $\sigma: Q \rightarrow \Delta^*$ a partial function that assigns strings to final states

Intuitively, a subsequential transducer is an FST whose incremental behaviour is always DETERMINISTIC. That is, state transitions and output strings are deterministic functions of the current state and input symbol, and the string-to-string function $\tau: \Sigma^* \rightarrow \Delta^*$ defined by a subsequential FST is also a deterministic mapping: every full input string is associated with at most one full output string.

Subsequential FSTs can be divided into two partially overlapping classes, based on the directionality of their computation (Chandlee 2014: ch. 3). Left-subsequential FSTs read input strings from left to right, while right-subsequential FSTs read input strings from right to left. When using subsequential FSTs to model vowel-harmony patterns, this distinction in directionality of computation maps intuitively onto the directionality of the harmony pattern. Canonical progressive harmony patterns are modelled with left-subsequential FSTs, and canonical regressive harmony patterns are modelled with right-subsequential FSTs.

Some data exhibiting the progressive rounding harmony pattern found in Turkish are shown in (15). Rounding spreads left to right from roots to suffixes; [+high] vowels undergo harmony (15f, g) while [-high] vowels block harmony (15h–j).

⁵ A transducer with an output function that can map each input symbol to a string of output symbols is a WORD (OR ASYNCHRONOUS) transducer; if ω can only map any given input symbol to a single output symbol, it is a LETTER (OR SYNCHRONOUS) transducer. Almost any word transducer can be converted to a letter transducer (Roche & Schabes 1997: ch. 1). We make use of letter transducers everywhere except when analysing Copperbelt Bemba in §4.3.

(15) Turkish progressive rounding harmony

| [-round] root | | [+round] root | |
|-------------------|----------|----------------------|-----------|
| a. /ip-in-i/ | [ipini] | f. /jyz-in-i/ | [jzyzny] |
| rope-2SG.POSS-ACC | | face-2SG.POSS-ACC | |
| b. /el-in-i/ | [elini] | g. /køj-in-i/ | [køjyny] |
| hand-2SG.POSS-ACC | | village-2SG.POSS-ACC | |
| c. /ip-ler-i/ | [ipleri] | h. /jyz-ler-i/ | [jyzleri] |
| rope-PL-ACC | | face-PL-ACC | |
| d. /el-ler-i/ | [elleri] | i. /køj-ler-i/ | [køjleri] |
| hand-PL-ACC | | village-PL-ACC | |
| e. /ip-in-de/ | [ipinde] | j. /jyz-in-de/ | [jzynde] |
| rope-2SG.POSS-LOC | | face-2SG.POSS-LOC | |

Appendix A1 gives an example of a left-subsequential FST which models this pattern,⁶ and in (16) we provide a ‘running tape’ representation of the mapping performed by the FST, using the word [jzynde] to illustrate both rounding harmony and its blocking. Since harmony is stem-controlled in the language, the FST outputs all segments to the left of the root–suffix boundary (symbolised as \vee) without modification (16a). Note that the FST does not as its first operation map the entire substring $\times jyz \vee$ to itself; it processes each input symbol incrementally. To save space, we do not show each of the first five steps separately, and adopt the same strategy in later examples.

(16) Turkish left-subsequential FST tape

a. Stage 1: faithful root map

| | | | | | | | | | | |
|--------|----------|---|---|---|--------|---|---|---|---|----------|
| input | \times | j | y | z | \vee | i | n | d | e | \times |
| output | \times | j | y | z | \vee | | | | | |

b. Stage 2: [+high] suffix map

| | | | | | | | | | | |
|--------|----------|---|---|---|--------|---|---|---|---|----------|
| input | \times | j | y | z | \vee | i | n | d | e | \times |
| output | \times | j | y | z | \vee | y | n | | | |

c. Stage 3: [-high] suffix map; computation ends

| | | | | | | | | | | |
|--------|----------|---|---|---|--------|---|---|---|---|----------|
| input | \times | j | y | z | \vee | i | n | d | e | \times |
| output | \times | j | y | z | \vee | y | n | d | e | \times |

The FST then reads a suffix with a [+high] vowel, and, since the roundness of the root vowel is known, the output function emits a vowel matching the roundness of the root in (16b). (Again, the two symbols of the suffix are technically processed incrementally. Since consonants do not participate in the harmony process, they are always mapped faithfully; we do not show these steps separately.)

⁶ The appendices are available as online supplementary materials at <https://doi.org/10.1017/S095267572000010X>. All running tape representations in the paper are presented as FSTs in Appendix A.

Next, the FST reads an input suffix with a [-high] vowel. This vowel is mapped faithfully, since [-high] vowels block rounding harmony. The end of the word is then reached and the computation ends (16c).

In this way, progressive harmony patterns can be modelled with left-subsequential FSTs – and in similar fashion, regressive harmony patterns can be modelled with right-subsequential FSTs. In the next subsection, we define and discuss weakly deterministic regular functions, characterised by a restricted composition of left- and right-subsequential FSTs.

4.2 Weak determinism

The ordering of rewrite rules $\rho_1 < \rho_2 < \dots < \rho_n$ in an *SPE*-style analysis corresponds to the ordered composition of associated string-to-string functions $\varphi_{\rho_n} \circ \dots \circ \varphi_{\rho_2} \circ \varphi_{\rho_1}$. While the composition of any two subsequential functions going in the same direction can only yield another subsequential function (Mohri 1997), a function defined by the composition of subsequential functions going in opposite directions can capture any regular function, as detailed in the next subsection. Weakly deterministic regular functions (see Fig. 1) are described by Heinz & Lai (2013) as those FSTs that can be defined as the composition of two subsequential functions going in opposite directions, such that the two functions do not use an intermediate alphabet containing symbols not present in the input alphabet; see (17).⁷

- (17) A regular string-to-string function $\tau: X^* \rightarrow Y^*$ is weakly deterministic iff there exist two subsequential functions $I: X^* \rightarrow X^*$ and $O: X^* \rightarrow Y^*$ such that O is either left- or right-subsequential, I is the other and $\tau = O \circ I$.

This definition of weakly deterministic regular functions is designed to pick out functions with two notable properties. First, this class of functions is able to capture bidirectional patterns that no right- or left-subsequential function can on its own. Second, the increased expressivity of this class is constrained by a restriction: no extra intermediate symbols. This is ensured by the condition that the output alphabet of the first ('inner') subsequential function, I , is the same as the input alphabet of both functions. Without such a restriction, the composition of two subsequential functions is guaranteed to be a function, but not necessarily a subsequential one (Elgot & Mezei 1965). In a weakly deterministic function, the behaviour of the second ('outer') function does not depend on any form of 'mark-up' deposited by the first function into the intermediate representation, and the restrictiveness of this class of patterns is intended to follow from the definition's prohibition of an intermediate alphabet with extra symbols. Without this restriction, the inner function could use additional

⁷ The full definition includes one additional condition that we set aside for clarity of exposition here; our goal here is to communicate the intuition behind Heinz & Lai's definition.

symbols to effectively provide unbounded look-ahead for the outer function, allowing the outer function to behave deterministically, given this marked-up version of the input string.

At this juncture, it is important that we clarify that splitting a phonological process into the composition of multiple finite transductions does not impact the formal status or complexity of the overall transduction. In formalisms that use cyclic rule application or level-ordering, derivational stages of input–output maps are given formal interpretations and theoretical significance. In contrast, no formal status is given to the non-surface outputs of intermediate transductions in current work exploring the subregular hierarchy (Chandlee & Heinz 2018, Chandlee *et al.* 2018). In this way, the formal language theory adopted here is similar to Optimality Theory, where the analytical focus is on properties of holistic input–output maps, regardless of whether they are characterised as multiple independent processes in other formalisms; here, too, only properties of the total input–output mapping are relevant to a pattern’s complexity. However, unlike Optimality Theory, the way in which inputs are mapped to their final output in formal language theory can vary, and therefore requires explication and justification. To this end, breaking a pattern into multiple transductions serves two interrelated purposes which are orthogonal to the minimum expressivity required to describe the overall pattern: (i) to aid human interpretability and reasoning about the overall transduction, and (ii) to validate that the simplest overall transduction has been selected (for example, by showing that a pattern can be generated as the composition of multiple single-direction subsequential transductions, we substantiate the claim that the pattern is subsequential). Although some individual transductions in a composition may resemble what would be considered an independent process in other formalisms, it is critical to keep in mind that this is for convenience and interpretability, and that these individual transductions do not have the same formal status as independent processes in other analytical frameworks.

For example, Heinz & Lai (2013) show that weakly deterministic functions are capable of describing bidirectional harmony patterns, such as the ATR harmony pattern found in Degema (Kari 1997, Archangeli & Pulleyblank 2007). The pattern in Degema differs from that in Turkish in two key respects: the harmonising feature and directionality. First, the feature which participates in harmony is [ATR], resulting in the alternations /i–ɪ/, /u–ʊ/, /e–ɛ/, /o–ɔ/ and /ə–a/. Second, affixes on *both* sides of a root agree in [ATR] with the root, as shown in (18).⁸ (FSTs modelling the Degema pattern can be found in Appendix A2.)

⁸ Note that Turkish lacks prefixes, so harmony can only proceed in one direction. See Baković (2000) for discussion.

(18) *Degema bidirectional ATR harmony*

- | | |
|---|---|
| <p>a. [-ATR] root</p> <p>/ʊ-fʊ-a/ [ʊfóā]</p> <p>GER-be.white-GER</p> <p>/ʊ-bòm-am/ [ʊbómām]</p> <p>GER-beat-GER</p> | <p>b. [+ATR] root</p> <p>/ʊ-bi-a/ [ubiā]</p> <p>GER-be.black-GER</p> <p>/ʊ-der-am/ [udárām]</p> <p>GER-cook-GER</p> |
|---|---|

In (19) and (20) we provide a running tape representation of the mapping, using the word [ubiā] in (18b). We begin with the left-subsequential function.⁹

(19) *Degema inner left-subsequential FST tape*

- | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|--|--|
| <p>a. <i>Stage 1: faithful prefix map</i></p> <p>input <table border="1" style="display: inline-table; border-collapse: collapse; text-align: center;"><tr><td>×</td><td>ʊ</td><td>√</td><td>b</td><td>i</td><td>√</td><td>a</td><td>×</td></tr></table></p> <p>output <table border="1" style="display: inline-table; border-collapse: collapse; text-align: center;"><tr><td>×</td><td>ʊ</td><td></td><td></td><td></td><td></td><td></td><td></td></tr></table></p> | × | ʊ | √ | b | i | √ | a | × | × | ʊ | | | | | | | <p>b. <i>Stage 2: faithful root map</i></p> <p>input <table border="1" style="display: inline-table; border-collapse: collapse; text-align: center;"><tr><td>×</td><td>ʊ</td><td>√</td><td>b</td><td>i</td><td>√</td><td>a</td><td>×</td></tr></table></p> <p>output <table border="1" style="display: inline-table; border-collapse: collapse; text-align: center;"><tr><td>×</td><td>ʊ</td><td>√</td><td>b</td><td>i</td><td>√</td><td></td><td></td></tr></table></p> | × | ʊ | √ | b | i | √ | a | × | × | ʊ | √ | b | i | √ | | |
| × | ʊ | √ | b | i | √ | a | × | | | | | | | | | | | | | | | | | | | | | | | | | | |
| × | ʊ | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| × | ʊ | √ | b | i | √ | a | × | | | | | | | | | | | | | | | | | | | | | | | | | | |
| × | ʊ | √ | b | i | √ | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| <p>c. <i>Stage 3: suffix vowel map; right edge reached</i></p> <p>input <table border="1" style="display: inline-table; border-collapse: collapse; text-align: center;"><tr><td>×</td><td>ʊ</td><td>√</td><td>b</td><td>i</td><td>√</td><td>a</td><td>×</td></tr></table></p> <p>output <table border="1" style="display: inline-table; border-collapse: collapse; text-align: center;"><tr><td>×</td><td>ʊ</td><td>√</td><td>b</td><td>i</td><td>√</td><td>ə</td><td>×</td></tr></table></p> | | × | ʊ | √ | b | i | √ | a | × | × | ʊ | √ | b | i | √ | ə | × | | | | | | | | | | | | | | | | |
| × | ʊ | √ | b | i | √ | a | × | | | | | | | | | | | | | | | | | | | | | | | | | | |
| × | ʊ | √ | b | i | √ | ə | × | | | | | | | | | | | | | | | | | | | | | | | | | | |

Following the prefix vowel, the FST encounters the prefix–root boundary, notated here with the symbol √. The transducer progresses through the input string, and faithfully outputs all symbols between the two root-boundary symbols (19b). In (19c), the FST reaches a suffix vowel. Because the [ATR] value of the root is known, the FST outputs [ə] from input /a/.

Following the application of this inner function, we see that affixes to the right of the root have been appropriately harmonised, but affixes to the left of the root have not. To capture the bidirectionality of harmony in Degema, we compose the left-subsequential function with the right-subsequential function, taking the output of (19c) and applying the right-subsequential function to it. To keep track of the derivation, we add another tape to our running tape representation below. The top tape represents the INITIAL INPUT, prior to application of the left-subsequential function; the middle, INTERMEDIATE, tape represents the output of the left-subsequential function and the input of the right-subsequential function. The new bottom tape represents the FINAL OUTPUT of the computation.

The right-subsequential function in (20a) mirrors the left-subsequential function that was previously applied. The computation begins from the right end of the input string, and upon reaching the first vowel in the string, the vowel is output faithfully.

⁹ The decision to treat the left-subsequential function as the inner function is arbitrary.

(20) Degema outer right-subsequential FST tape

a. Stage 1: faithful suffix vowel map

| | | | | | | | | | |
|---------------------|---|---|---|---|---|--|---|---|---|
| <i>input</i> | × | u | √ | b | i | | v | a | × |
| <i>intermediate</i> | × | u | √ | b | i | | v | ə | × |
| <i>output</i> | | | | | | | | ə | × |

b. Stage 2: faithful root map

| | | | | | | | | | |
|---------------------|---|---|---|---|---|--|---|---|---|
| <i>input</i> | × | u | √ | b | i | | v | a | × |
| <i>intermediate</i> | × | u | √ | b | i | | v | ə | × |
| <i>output</i> | | | √ | b | i | | v | ə | × |

c. Stage 3: prefix vowel map; computation ends

| | | | | | | | | | |
|---------------------|---|---|---|---|---|--|---|---|---|
| <i>input</i> | × | u | √ | b | i | | v | a | × |
| <i>intermediate</i> | × | u | √ | b | i | | v | ə | × |
| <i>output</i> | × | u | √ | b | i | | v | ə | × |

Following the suffix vowel, the FST encounters the root–suffix boundary symbol, and the transducer outputs all characters between the root–boundary characters faithfully (20b). Finally, in (20c) the FST reaches the initial prefix vowel, which is currently disharmonic. The [ATR] value of the root is known at this point, and the FST outputs a vowel matching the root value for the harmonic feature.

The computation is now complete, and the initial input has been successfully harmonised to the final output form observed in Degema. The analysis of this harmony pattern requires the composition of two sub-sequential functions, but maintains the alphabet size of the initial input throughout the application of both functions, satisfying the definition of a weakly deterministic function given by Heinz & Lai (2013). In the next subsection, we describe non-deterministic functions and a set of phonological maps which require non-determinism to be described – unbounded circumambient patterns.

4.3 Non-determinism

As noted in §4.1, a subsequential transducer $\tau: X^* \rightarrow Y^*$ defines a string-to-string function that is deterministic in its input string and has state transition and output functions that are deterministic in the current state and input symbol: every input string $w \in X^*$ is mapped by τ to at most one string in Y^* and any given (state, input symbol) pair is mapped by δ and ω to at most one state and at most one output string respectively. If τ can map at least one input string to more than one output string, then τ is a non-deterministic function on strings; it is a string RELATION. In contrast, a transducer that maps every input string to at most one output string is said to be FUNCTIONAL or SINGLE-VALUED. If there are any (state, input symbol) pairs such that a transducer can ‘choose’ from among a set of

two or more states to transition to, or from among a set of two or more strings to output, then that transducer's transition and/or output functions are non-deterministic.

Following Heinz (2018: §6.2.5), we use the term REGULAR RELATIONS to refer exclusively to the most general class of string-to-string mappings definable using FSTs – the class that includes transducers which are not functional. We also follow Heinz in using NON-DETERMINISTIC REGULAR FUNCTIONS to refer to the class of transducers that are functional, but can have non-deterministic state and/or output functions. Such transducers can have temporarily and incrementally ambiguous input strings, but the point of disambiguation may be an unbounded distance away from the location of the read/write head at the moment of ambiguity. As Elgot & Mezei (1965) show, any regular function can be decomposed into two subsequential functions going in opposite directions, as long as the first function in the composition is allowed to enlarge the input alphabet, as shown in (21).

- (21) A function $\tau: X^* \rightarrow Z^*$ is regular iff there exist two subsequential functions $I: X^* \rightarrow Y^*$ and $O: Y^* \rightarrow Z^*$ such that O is either left- or right-subsequential, I is the other, $X \subseteq Y$ and $\tau = O \circ I$.

As mentioned in the previous subsection, the intuition is that the first, inner function can effectively perform unbounded look-ahead for the second function by 'marking up' intermediate strings with extra information (in the form of extra symbols that are in Y but not in X) that the second, outer function can then use to behave in a manner that is incrementally deterministic.

Jardine (2016) conjectures that unbounded circumambient processes are a class of patterns which require non-determinism, the full expressivity of regular relations.¹⁰ For this reason, to introduce this complexity class, we walk through Jardine's analysis of tonal spreading in Copperbelt Bemba.

In Copperbelt Bemba (Bickmore & Kula 2013, Kula & Bickmore 2015; see also Pater 2018 for discussion), a high tone spreads unboundedly to the right word edge in phrase-final forms (22a). However, if another high tone intervenes between the first high tone and the word edge, then bounded ternary spreading takes place instead (22b).

¹⁰ Smith & O'Hara (2019) and O'Hara & Smith (2019) take advantage of loopholes in Heinz & Lai's (2013) tentative definition of weakly deterministic functions to offer analyses of Copperbelt Bemba that *appear* to be weakly deterministic, obeying the letter but violating the spirit of Heinz & Lai's definition.

(22) Copperbelt Bemba tone spreading

- | | |
|---|----------------------|
| a. /u-ku-tul-a/ AUG-CL15-pierce-FV | [ùkùtùlà] |
| b. /bá-ka-fik-a/ 2SM-FUT-arrive-FV | [bákáfiká] |
| c. /bá-ka-mu-londolol-a/ 2SM-FUT-1OM-introduce-FV | [bákámúlóóndólólá] |
| d. /tu-ka-páapaatik-a/ 1PL-FUT-flatten-FV | [tùkàpáápáátiká] |
| e. /bá-ka-pat-a-kó/ 2SM-FUT-hate-FV=CL17.LOC | [bákápátàkó] |
| f. /bá-ka-londolol-a=kó/ 2SM-FUT-introduce-FV=CL17.LOC | [bákálóóndólólákó] |
| g. /tu-ka-béleeng-el-an-a=kó/ 1PL-FUT-read-APPL-RECP-FV=CL17.LOC | [tùkàbélééngélànakó] |

We schematise this pattern in (23).¹¹

- | | |
|---|---|
| (23) a. $L^n \rightarrow L^n$ | no H present |
| b. ... $HL^m \rightarrow \dots HH^m$ | unbounded H spreading in the absence of a following H |
| c. ... $HL^n H \dots \rightarrow \dots HHH L^{n-2} H \dots$ | ternary H spreading in the presence of a following H ($n \geq 2$) |
| d. ... $HLH \dots \rightarrow \dots HHH \dots$ | binary spreading up to a following H |

In (24) and (25) we show ternary spreading with a running tape representation of the input string /HLLLH/. The input–output mapping here is composed of a left-subsequential FST, as well as a right-subsequential FST that reads the output of the left-subsequential FST as its input, and outputs the actual attested form (FSTs modelling the mapping can be found in Appendix A3). The order between the two FSTs is not arbitrary: the left-subsequential FST adds mark-up that the right-subsequential FST then takes advantage of.¹²

Generally speaking, the left-subsequential FST outputs all L tones without modification until encountering a H tone, at which point the H ‘spreads’ to two following L tones. For all following L tones, the FST

¹¹ Ternary spreading is analysed as two separate processes in Bickmore & Kula (2013) and Kula & Bickmore (2015). The motivation for their analysis and the constraints on ternary spreading are not relevant to the present paper.

¹² As Jardine (2016: 265ff) describes in more detail, it is possible to allow the FST to ‘wait’ a finite number of states; this allows the machine to look ahead and access information on the far side of the target segment before deciding what to output for that target segment. To preserve subsequentiality while introducing ‘waiting’, there must be an *a priori*-known bound on the number of wait states (here equivalent to the look-ahead context having an *a priori* bounded length). That is not possible here, as a result of unbounded circumambience.

outputs a distinct symbol (notated as Ψ here) not contained in the input alphabet, because it cannot determine whether the output should be an L (or an H) tone until it does (or does not) encounter another H tone later in the word. Since a second H tone may in principle occur an unbounded distance from the first, the FST cannot ‘wait’, and so instead outputs the placeholder Ψ symbol, indicating that the preceding context matches the lefthand side of the structural description for unbounded spreading. The FST iteratively outputs all subsequent input L tones as Ψ until it encounters either an input H, which is output without modification, or the right word edge.

With the input /HLLLH/, the inner left-subsequential FST first reads the word-initial H tone, and initiates ternary spreading to the immediately following two L tones (24a).

(24) *Copperbelt Bemba inner left-subsequential FST tape*

a. *Stage 1: ternary spreading*

| | | | | | | | |
|--------|---|---|---|---|---|---|---|
| input | × | H | L | L | L | H | × |
| output | × | H | H | H | | | |

b. *Stage 2: mark-up*

| | | | | | | | |
|--------|---|---|---|---|--------|---|---|
| input | × | H | L | L | L | H | × |
| output | × | H | H | H | Ψ | | |

c. *Stage 3: right edge reached*

| | | | | | | | |
|--------|---|---|---|---|--------|---|---|
| input | × | H | L | L | L | H | × |
| output | × | H | H | H | Ψ | H | × |

Next, the FST reads the third L tone from the input, but cannot determine at this point whether to output an H or an L tone, because the presence or absence of a following H tone is unknown. The FST thus outputs a new symbol not contained in the input alphabet, Ψ (24b). This new symbol will ultimately provide the outer right-subsequential FST with the information necessary to determine all tone values for the word. Progressing through the input string, the inner FST continues as described above. In (24c), the FST reads a second input H tone and maps it faithfully, and then reaches the right edge.

The outer right-subsequential FST now reads the output just produced by the inner left-subsequential FST as its input, and completes the input–output mapping. Generally speaking again, the outer FST outputs all input L and H tones without modification, but maps intermediate Ψ according to the previously read context. Proceeding from right to left, if the FST reads an H tone, then it outputs all Ψ as L; if it does not encounter H, then it outputs all Ψ as H. The outer FST is able to discern whether a second H tone is present in the word, and uses the mark-up passed from the inner FST to determine whether ternary or unbounded spreading occurs.

This is shown in (25) for the more specific form under discussion. First, the outer right-subsequential FST reads the word-final H tone, and outputs it faithfully.

(25) Copperbelt Bemba outer right-subsequential FST tape

a. Stage 1: faithful H map

| | | | | | | | |
|---------------------|---|---|---|---|---|---|---|
| <i>input</i> | × | H | L | L | L | H | × |
| <i>intermediate</i> | × | H | H | H | ψ | H | × |
| <i>output</i> | | | | | | H | × |

b. Stage 2: mark-up resolution

| | | | | | | | |
|---------------------|---|---|---|---|---|---|---|
| <i>input</i> | × | H | L | L | L | H | × |
| <i>intermediate</i> | × | H | H | H | ψ | H | × |
| <i>output</i> | | | | L | H | × | |

c. Stage 3: left edge reached; computation ends

| | | | | | | | |
|---------------------|---|---|---|---|---|---|---|
| <i>input</i> | × | H | L | L | L | H | × |
| <i>intermediate</i> | × | H | H | H | ψ | H | × |
| <i>output</i> | × | H | H | H | L | H | × |

Since an input H has been encountered, the FST outputs all ψ as L in (25b) and all H faithfully in (25c). The result in this case is thus ternary, rather than unbounded, spreading from the initial H.

Unbounded circumambient processes like tone spreading in Copperbelt Bemba can thus be analysed as regular non-deterministic maps, either in the form of a single non-deterministic FST, as in Jardine (2016), or as the composition of two subsequential functions that may use an enlarged alphabet containing some symbols not present in the initial input alphabet.

4.4 Summary

The three levels of expressivity defined and exemplified in this section can be summarised as follows. Subsequential regular functions can describe unidirectional processes with bounded look-ahead (and unbounded ‘look-behind’), as in Turkish harmony. Weakly deterministic functions are intended to be able to describe bidirectional processes where the first ‘pass’ is *not* allowed to behave as look-ahead for, or otherwise affect the behaviour of, the second pass, as in Degema. Non-deterministic regular functions can describe compositions of unidirectional processes going in opposite directions, where the first pass may serve as unbounded look-ahead for, or otherwise affect the behaviour of, the second pass, as in Copperbelt Bemba.

Jardine (2016: §5.4) acknowledges two attested cases of apparent unbounded circumambience in segmental phonology, Sanskrit *n*-retroflexion (Ryan 2017) and Yaka height harmony (Hyman 1998). He does not consider them equivalent to tonal patterns, however, claiming that such segmental patterns are ‘extremely rare’. He suggests that the harmony patterns in Sanskrit and Yaka may not actually be unbounded. In §5 we analyse ATR harmony in Tutrugbu, which provides a further challenge to the claim that segmental phonology is at most weakly deterministic. Looking ahead, we identify a number of other cases in §6.2 as evidence that unbounded circumambient vowel-harmony patterns are more widely attested than previously thought.

should surface, [ba] or [be]. This indeterminacy precludes a subsequential analysis of Tutrugbu. If, however, we mimic the structure of the analysis provided for Copperbelt Bemba in §4.3, where the first FST is able to introduce intermediate mark-up into the derivation, the Tutrugbu pattern is analysable.

Generally speaking, the first, right-subsequential FST proceeds leftward, outputting all high vowels in accordance with the root's [ATR] value. If, after reading a [+ATR] root, the FST reads an input [−high] vowel, the transducer first outputs Ψ and then all subsequent vowels as either Ψ (for further [−high] vowels) or \mathcal{H} (for [+high] vowels). By doing so, this first FST passes information about the [ATR] value of the root and the presence of a potential blocker to the second, left-subsequential FST.

This analysis is illustrated in (28) with the example [ɪbawu]. First, the right-subsequential FST faithfully maps the root, starting from the right edge of the input.

(28) *Tutrugbu inner right-subsequential FST tape*

a. *Stage 1: faithful root map*

| | | | | | | | | |
|---------------|---|---|---|---|---|---|---|---|
| <i>input</i> | × | ɪ | b | a | √ | w | u | × |
| <i>output</i> | | | | | √ | w | u | × |

b. *Stage 2: [−high] prefix vowel mark-up*

| | | | | | | | | |
|---------------|---|---|---|--------|---|---|---|---|
| <i>input</i> | × | ɪ | b | a | √ | w | u | × |
| <i>output</i> | | | b | Ψ | √ | w | u | × |

c. *Stage 3: [+high] prefix vowel mark-up; left edge reached*

| | | | | | | | | |
|---------------|---|---------------|---|--------|---|---|---|---|
| <i>input</i> | × | ɪ | b | a | √ | w | u | × |
| <i>output</i> | × | \mathcal{H} | b | Ψ | √ | w | u | × |

Next, this inner FST encounters a prefix with a [−high] vowel (28b). Having not yet determined the height of the initial prefix vowel, the [−high] prefix vowel is output with the new symbol Ψ , which encodes its height and leaves its [ATR] value to be determined by the outer, left-subsequential FST. The inner FST then encounters a prefix with a [+high] vowel (28c). It happens to be the initial prefix in this case, but as with all [+high] prefix vowels, it is output with \mathcal{H} , which encodes its height and leaves its [ATR] value to be determined. The left edge is then reached, and this pass ends.

The outer left-subsequential FST then proceeds from left to right, using the output of the inner FST as its input, as in (29).

(29) *Tutrugbu outer left-subsequential FST tape: mark-up resolved; right edge reached; computation ends*

| | | | | | | | | |
|---------------------|---|---|---|---|---|---|---|---|
| <i>input</i> | × | ɪ | b | a | √ | w | u | × |
| <i>intermediate</i> | × | ʌ | b | ψ | √ | w | u | × |
| <i>output</i> | × | ɪ | b | a | √ | w | u | × |

Upon encountering a ʌ in the initial syllable, the outer FST ‘knows’ that the conditions for blocking have been met, because initial ʌ indicates both that the root is [+ATR] and that there is a medial [–high] prefix vowel. Because ʌ is in the initial syllable, this FST outputs all ʌ and ψ as [ɪ] and [a] respectively. With the introduction of additional symbols into the derivation, this analysis of Tutrugbu parallels the non-deterministic analysis of Copperbelt Bemba in §4.3.

As noted earlier in §4.3, a subsequential FST can utilise ‘wait’ states to achieve (bounded) look-ahead without increasing expressivity, but the maximum number of wait states (here equivalent to the amount of look-ahead) must be known and fixed *a priori*. Given that the context to the left of a target [–high] vowel in Tutrugbu may be unboundedly long, then there is no single number of ‘waiting’ transition states that can model all of the possible data and capture the principled generalisation describing it. In (8) we demonstrated that [–high] vowels could be separated from the initial syllable by a number of syllables with no principled upper bound. The generalisation is that the realisation of [ATR] on [–high] prefix vowels depends on both the [ATR] value of the root and the height value of the initial syllable. Furthermore, a given [–high] vowel may occur a potentially unbounded distance from both the root and the initial syllable. The conditional blocking pattern in Tutrugbu is thus unbounded circumambient, like the tonal processes analysed in Jardine (2016).

6 Discussion

Formal language theory offers a valuable framework-independent meta-language for comparing linguistic formalisms and classifying phonological patterns in terms of formally well-characterised and independently motivated categories of computational complexity and expressivity (Heinz 2018). While the subregular programme has excelled at identifying salient categories for phonological typology, we think the programme is at its weakest when it comes to predicting and explaining typology. It has primarily done so by positing, on a post hoc basis, innate categorical constraints on the hypothesis space that human learners apply to phonology via a phonology-specific learning mechanism. In the subsections below we present evidence that segmental phonology is substantially more expressive than previously predicted, outlining what a more productive relationship between formal language theoretic phonology and

typological explanation could look like, and how phonological theories should be evaluated.

6.1 How complex is phonology?

The first counterexamples to the subregular hypothesis articulated in Heinz (2011a, b, 2018) came from the tonal patterns in Jardine (2016), which require the expressivity of non-deterministic regular functions. To uphold the subregular hypothesis, Jardine (2016) adds a qualification: *segmental* phonology is subregular, but *tonal* phonology is not constrained in the same way. Heinz (2018) writes:

Jardine's result [i.e. the existence of unbounded circumambient phonological patterns] is perhaps the most serious challenge to the Subsequential Hypothesis (or a revised Weakly Deterministic hypothesis) because the best characterization of Yaka vowel harmony seems to be that it is circumambient unbounded (Hyman 1998). However, this is the only known example of this type, and it is probably premature to reject the hypothesis on these grounds alone.

In addition to the Tutrugbu pattern detailed above, we consider in §6.2 a comparable number of segmental phenomena requiring the same level of expressivity as the tonal patterns discussed by Jardine (2016), as well as those in Yaka, adding more counterexamples to any proposed subregular upper bound on the complexity of phonology (see Avcu 2018 for experimental evidence for the learnability of more complex patterns). This evidence shows that segmental phonology, like tonal phonology, is more expressive than previously thought, requiring the expressivity of non-deterministic regular functions.

In fact, there is some evidence that segmental patterns may require strictly *more* expressivity than the non-deterministic regular functions. As Heinz (2018: §6.2.5) notes, optional processes require non-deterministic mappings that are, unlike the Copperbelt Bemba and Tutrugbu patterns, non-functional (recall the discussion at the beginning of §4.3). Optional processes thus require the more expressive class of regular relations. In the face of optionality, Heinz (2018: 175) suggests that optionality may be 'handled at a higher level of control than the individual transformation'. Although offloading optionality to some other part of the grammar effectively downplays the significance of these potential counterexamples to the subregular hypothesis, patterns of optionality like those listed in Vaux (2008) and others like iterative optionality in Icelandic umlaut (Anderson 1974) present challenges to the strong claim that segmental phonology is categorically subregular. In particular, in cases where an iterative process may be optional, there is no *a priori* bound on the number of potential alternations (Mohri 1997).

Moreover, some recent work suggests that segmental phonology may require more expressivity than the regular relations: Bowler (2013) and Bowler & Zymet (2019) demonstrate from a corpus of Warlpiri nouns

that disharmonic nouns exhibit optional ‘majority rules’ harmony (Lombardi 1999, Baković 2000), a pattern that numerous linguists have claimed to be unattested, impossible and pathological, given that it requires even more expressivity than the regular relations (Riggle 2004, Heinz & Lai 2013, Hulden 2017, Heinz 2018, Lamont 2019). Whatever conclusion future work comes to about the significance of the last two pieces of evidence, these patterns suggest that there is presently no clear subregular cap on the empirically observed expressivity of phonology – segmental or tonal.

6.2 Additional unbounded circumambient segmental patterns

In §6.2.1–§6.2.4 we briefly describe four additional cases of unbounded circumambient segmental patterns, all involving ATR harmony: Tafi, Turkana, Karimojong and Liko. A few other segmental patterns whose unboundedness is less clear are discussed in §6.2.5. The ATR harmony patterns of Tutrugbu, Tafi, Turkana, Karimojong and Liko all quite clearly meet Jardine’s (2016) definition of unbounded circumambience in (13). Jardine’s claim that unbounded circumambient segmental patterns are significantly rarer than their tonal counterparts, and that they are therefore not within the computational bounds of segmental phonology, thus appears to have been premature.

6.2.1 *Tafi*. We have claimed that Tafi (Kwa), a closely related language, exhibits the same pattern of harmony as Tutrugbu. Bobuafor (2013) does not give an explicit description of conditional blocking in Tafi, but its existence is clear from the data presented, as illustrated in (30) (from Bobuafor 2013 and Mercy Bobuafor (personal communication)). Harmony obtains when a [+ATR] verb is preceded by only [+high] prefixes (30a) or only [–high] prefixes (30b). It is also found when a [+ATR] root is preceded by vowels of both heights, and the initial-syllable vowel is [–high] (30c). However, when the initial-syllable vowel is [+high] and is followed by a [–high] vowel, harmony is blocked by the [–high] vowel, just as in Tutrugbu (30d, e). The additional examples in (30f, g) demonstrate that blocking occurs at longer distances, since in these examples the initial-syllable high vowel and medial non-high vowel are not syllable-adjacent.¹⁴

¹⁴ We assume that input tonal specifications are identical to those in the output in these Tafi examples.

- (30) a. /bu-ó-bī/ [buúbī] e. /kɪ-ba-dzí/ [kɪbadzí]
 SM-PROG-be.cooked 3SG.DEP-FUT-be
 b. /a-ba-dzi/ [ebedzi] f. /lɪ-tí-ba-dzi/ [lɪtɪbadzí]
 3SG.DEP-FUT-become 3SG.DEP-NEG-FUT-be
 c. /á-tí-ba-bí/ [étíbebí] g. /ɪ-tí-ba-dzu/ [ɪtɪbadzu]
 3SG.DEP-NEG-be.cooked 1SG-NEG-FUT-build
 d. /lǒ-ba-bí/ [lǒbabí]
 3SG.DEP-FUT-be.cooked

6.2.2 *Turkana*. Besides Tutrugbu and Tafi, Turkana (Nilotic) also features a vowel-harmony pattern that is unbounded circumambient (Dimmendaal 1983, Baković 2000).¹⁵ In Turkana, [+ATR] spreads bidirectionally from a dominant root or suffix, as shown in (31). Prefix and other recessive vowels alternate, but the behaviour of /a/ depends on whether it is to the left or the right of a dominant vowel: /a/ to the left is opaque to harmony and does not alternate (as shown by the gerundial prefix in these examples), while /a/ to the right alternates between the [+low, -ATR] vowel [a] and the [-low, +ATR] vowel [o], with [o] occurring after [+ATR] dominant vowels (as shown by the epipatetic vowel in these examples).¹⁶

(31) *Turkana ATR harmony*

- a. /a-k-ɪnɔk-a-kɪn-ɪ/ [akɪnɔkakɪnɪ]
 GER-K-light.fire-EPI-DAT-VOI
 b. /a-k-ibus-a-kɪn-ɪ/ [akibusokini]
 GER-K-drop-EPI-DAT-VOI

A small subset of dominant suffixes in Turkana are [-ATR], and when a dominant [-ATR] suffix co-occurs with a [+ATR] root, the realisation of the epipatetic vowel between them is neither [+low, -ATR] [a] nor [-low, +ATR] [o], but rather [-low, -ATR] [ɔ], as shown in (32). As above, in (32a) the epipatetic vowel surfaces as [o] after the [+ATR] root /ibus/ when no dominant [-ATR] suffix vowel follows. However, in (32b), a dominant [-ATR] suffix occurs to the right of the epipatetic vowel, which surfaces as [ɔ].

- (32) a. /ɛ-ibus-a-kɪn-ɪ/ [eibusokini]
 3SG-drop-EPI-DAT-VOI
 b. /ɛ-ibus-a-kɪn-a/ [ɛibusɔkɪnɔ]
 GER-drop-EPI-DAT-VOI

¹⁵ See Noske (1996, 2000) for more on Turkana vowel harmony.

¹⁶ Word-final vowels in the examples below are devoiced; see Dimmendaal (1983: §2.3.5) for discussion.

The surface quality of a suffixal low vowel is thus not determinable based solely on whether there is an [+ATR] vowel to its left; it also depends on the presence or absence of a dominant [-ATR] suffix to its right. In other words, the realisation of a suffixal low vowel depends on information a potentially unbounded distance away in both directions (see Appendix A5), and the realisation of a suffix low vowel is therefore non-deterministic. ATR harmony in Turkana is thus unbounded circumambient.

6.2.3 *Karimojong*. The pattern described above for Turkana is also attested in Karimojong (Nilotic; Novelli 1985, Lesley-Neuman 2012). As in Turkana, the low vowel suffix alternates between [a] and [o] for [ATR]. Note specifically the alternation of the frequentive suffix, /-an/, in (33a, b). There are also dominant [-ATR] suffixes, like the itive suffix /ɔr/, shown in (33c, d). In (33c), the dominant [-ATR] itive suffix attaches to the right of the root, spreading its [-ATR] value to the root. Moreover, the itive suffix causes both the frequentive suffix and the root to surface as [-ATR] in (d). Similarly to Turkana, the surface quality of the frequentive suffix depends on both the [ATR] value of the root and the presence or absence of a dominant [-ATR] suffix, surfacing as [an], [ɔn] or [on].

(33) *Karimojong ATR harmony*

- | | |
|--|------------------|
| a. /akɪ-tɔ-dɔŋ-an-akin/ INF-CAUS-handle.firmly-FREQ-DAT | [akɪtɔdɔŋanakin] |
| b. /akɪ-zi-dɔŋ-an-akin/ INF-CAUS-castrate-FREQ-DAT | [akizidɔŋonokin] |
| c. /akɪ-dɔŋ-ɔr/ INF-castrate-ITV | [akɪdɔŋɔr] |
| d. /akɪ-dɔŋ-ɔn-ɔr/ INF-castrate-FREQ-ITV | [akɪdɔŋɔnɔr] |

In (33d), the frequentive suffix exhibits a dependency on the [ATR] value of two syllable-adjacent vowels. Longer-distance dependencies are also reported in Karimojong. In reduplicated stems, an epenthetic vowel is inserted immediately before suffixes. In (34a, b), the epenthetic vowel (which Novelli 1985: 224 calls the ‘stem-enlarging vowel’) alternates based on the quality of the root: it surfaces as [ɔ] after the reduplicated stem /dɔŋɔdɔŋ/, but as [o] after /dɔŋidɔŋ/. Thus the stem-enlarging vowel is like the frequentive in (33). Like the frequentive suffix, the epenthetic vowel takes on the [-ATR] value of a following itive suffix (34c). Yet, in (34d, e), both the stem-enlarging vowel and the itive suffix take on the value of a following dominant [+ATR] suffix. In other words, the surface quality of the stem-enlarging vowel is conditioned by the potentially non-local [ATR] values of dominant morphemes on both its left and its right.

- (34) a. /akɪ-tɔ-dɔŋɔdɔŋ-ɔ-un/ [akɪtɔdɔŋɔdɔŋɔun]
 INF-CAUS-pinch-SE-FREQ
 b. /akɪ-dɔŋɔdɔŋ-ɔ-un/ [akɪdɔŋɔdɔŋɔun]
 INF-castrate-SE-FREQ
 c. /akɪ-dɔŋɔdɔŋ-o-ɔr/ [akɪdɔŋɔdɔŋɔɔr]
 INF-castrate-SE-ITV
 d. ɛ-dɔŋɔdɔŋ-ɔ-ɔr-i/ [ɛdɔŋɔdɔŋɔoori]
 1SG-pinch-SE-ITV-1PL.IND.T1
 e. /ɛ-dɔŋɔdɔŋ-ɔ-ɔr-ete/ [ɛdɔŋɔdɔŋɔoorete]
 3PL-pinch-SE-ITV-TMA

6.2.4 *Liko*. *Liko* (Bantu) displays a slightly different pattern of unbounded circumambience (de Wit 2015). ATR harmony is typically controlled by roots in *Liko*. Rightward harmony affects all vowels, while leftward harmony affects only the first prefix. In (35a), all affixes surface with their underlying, i.e. [-ATR] feature values, while in (35b) the [+ATR] root causes input /a/ to surface as [o]. The unboundedness of rightward harmony is seen in the three suffix alternations in (35c) (de Wit 2015: 77, 82).

(35) *Liko* ATR harmony

- a. /ta-pók-a/ [tapóká] c. /ká-lut-án-ág-á/ [kólutónógó]
 1PL-leave-FV CL9B-pull-ASSOC-PL-FV
 b. /ta-ḃín-a/ [toḃíno]
 1PL-dance-FV

Like the other languages discussed in this section, *Liko* possesses dominant [+ATR] and dominant [-ATR] morphemes. The negative enclitic, /=gɔ/, is invariantly [-ATR] in (36). In (36a), the enclitic does not undergo harmony from a [+ATR] root, although the preceding [+high] suffix does. When preceded by an underlying /a/, though, the enclitic not only resists [+ATR] harmony, but also prevents preceding low vowels from undergoing harmony (36b, c). In (b), the enclitic blocks harmony on a single preceding /a/, and in (c), the negative enclitic blocks harmony on both preceding low vowels (de Wit 2015: 94–96).

- (36) a. /ná-ká-ḃín-I=gɔ/ [nákóḃínigɔ]
 1SG.PST-NEG-dance-FV=NEG
 b. /ná-ká-ḃín-á=gɔ/ [nákóḃínágɔ]
 1SG.PST-NEG-dance-FV.PST=NEG
 c. /ná-ká-ḃín-ág-á=gɔ/ [nákóḃínágágɔ]
 1SG.PST-NEG-dance-PL-FV.PST=NEG

De Wit (2015) argues that [+ATR] is the active feature value in the language, further arguing that alternating vowels are specified as [-ATR]

underlyingly. In examples like (36b, c), the root value for [ATR] is not sufficient to determine whether suffixal /a/ undergoes harmony to [o]. The presence or absence of a dominant enclitic must be known in order to map low suffix vowels to their attested output forms.

One might object to treating Turkana, Karimojong and Liko in the same way as Tutrugbu and Tafi, since ATR harmony with conditional blocking in the first two is, at least intuitively, a single pattern. In the other three languages, on the other hand, spreading of [+ATR] and the more restricted pattern of [-ATR] spreading are – again, at least intuitively – two distinct patterns, as in existing analyses (Dimmendaal 1983, Baković 2000, Lesley-Neuman 2012). Recall, however, that the question of interest is the minimum level of complexity required to describe complete input–output mappings (Heinz 2018: §2). Thus the object of study is not the complexity of a particular pattern isolated from the larger phonology in a language, but rather the function that maps inputs to outputs comprising *all* relevant patterns. As a consequence of this framing, the multiple [ATR] spreading patterns in Turkana, Karimojong and Liko are together comparable in complexity to the single pattern found in Tutrugbu or Tafi.

6.2.5 *Other cases.* There are other apparent cases of circumambient harmony described in the literature, but these are less clear, and we cannot say with certainty that they are both unbounded and circumambient. Hyman (1998) argues that in Yaka (Bantu), high suffix vowels lower if they are preceded and followed by mid vowels, thus exhibiting some similarity to unbounded tone plateauing. Although the harmony extends over several syllables, Jardine (2016) doubts the unboundedness of this pattern. But as far as we can tell, harmony in Yaka satisfies all three of Jardine's (2016: 250) stated criteria in (37).

- (37) a. The source authors characterise the process as unbounded, and there is no evidence to the contrary.
 b. Examples exist of the process operating over multiple units.
 c. Examples exist of the process applying even when productive word- or phrase-formation processes extend its domain.

In addition to Yaka, Jardine (2016: §2.3.1) questions the status of retroflex harmony in Sanskrit (Indo-Aryan), based on Ryan (2017). In Sanskrit, a retroflex consonant triggers retroflexion of a following /n/ as long as another retroflex consonant is not present later in the word. Ryan (2017: §4) finds no examples of blocking across more than one syllable, a point which Jardine uses to question the boundedness of harmony. But Ryan shows evidence of triggering across intervening segments and blocking across one, two or three segments, which, in our estimation, suggests that the pattern may be unbounded.

Teso (Nilotic; Rottland & Otaala 1983), Toposa (Nilotic; Schröder & Schröder 1987) and Bondu-So (Dogon; Hantgan & Davis 2012) may

also exhibit unbounded circumambience. These languages all exhibit patterns of conflicting [ATR] dominance, like Turkana and Karimojong, but triggering morphemes in extant descriptions always immediately precede or follow the relevant suffix, so we cannot determine the unboundedness of these harmony patterns. If these are in fact bounded, then they parallel other cases of bounded circumambient harmony, like those described in Kalinowski (2009) and Lionnet (2016).

6.2.6 *Summary.* We contend that claims about the frequency of complex segmental phenomena relative to tonal phenomena are at best premature. As further support for a computational distinction between segmental and tonal processes, Heinz (2018) and Jardine (2016) claim that unbounded circumambient tonal processes are relatively common and equivalent segmental mappings are at most ‘extremely rare’. In support of this, Jardine adduces nine potential examples of such tonal processes, but only two unbounded circumambient segmental processes, Sanskrit *n*-retroflexion and Yaka height harmony. In our view, the difference between two and nine examples is not persuasive evidence of a difference in typological frequency. As Piantadosi & Gibson (2014) contend (and as we elaborate in §6.3), outside particular cases of carefully framed questions and statistical reasoning, we do not currently – nor will we soon – have descriptions of enough languages to draw strong conclusions about what properties are categorically impossible (e.g. non-deterministic segmental phonological patterns) for natural languages to have on the basis of simple comparisons of frequency counts. Furthermore, we simply do not have enough data to even know how many languages have tone, or how many have vowel harmony. At best, we have very rough ballpark estimates. For instance, Yip (2002: 1) notes that ‘by some estimates as much as 60–70 per cent of the world’s languages are tonal’. Such rough estimates demonstrate how much we have to learn before we can evaluate the frequency of unbounded circumambient tonal patterns and their segmental equivalents.

Even if the foregoing were not the case, of the nine tonal patterns marshalled as evidence in Jardine (2016), at least two (Digo and Xhosa) are not actually unbounded circumambient – see Appendix B for a left-subsequential analysis. Moreover, we have identified in this subsection at least four examples of unbounded circumambient segmental patterns in addition to our main example of Tutrugbu ATR harmony. We discuss these patterns, not to make typological claims about the frequency of these patterns relative to similarly complex tonal processes, but instead to demonstrate that they must be dealt with seriously, and not swept aside as outliers.

6.3 How should we establish and explain typological generalisations?

The empirical data from Tutrugbu (§2) and the additional languages described in §6.2 provide strong reasons to revise the typological claims

248 Adam G. McCollum, Eric Baković, Anna Mai and Eric Meinhardt in Jardine (2016), outlined in §3. Additionally, we would like to point out problems with the approach to typology advanced by existing work on the subregular hypothesis (e.g. Heinz & Idsardi 2011, Chandlee & Heinz 2012, 2018, Chandlee *et al.* 2012, Gainor *et al.* 2012, Heinz & Lai 2013, Jardine 2016) and summarised in (38).

- (38) a. The principal explanation for apparent typology is constraints on possible grammars in the hypothesis space humans use for learning phonological patterns.
- b. Naive estimates of frequency counts or literature reviews are a reasonable evidential basis for concluding that a particular type of phenomenon is categorically impossible.

In the particular case of segmental and tonal pattern, even granting both points in (38), our discussion of the vowel-harmony patterns in Tutrugbu and in §6.2 shows that a non-trivial number of unbounded circumambient patterns exist in segmental phonology, just as they do in tonal phonology.

More generally, contemporary literature on typology and linguistic theory offers four reasons for rejecting (38b). First, Piantadosi & Gibson (2014) show that, in general, the number of independent languages that linguists need to observe to achieve reasonable statistical confidence that a particular type of feature or phenomenon is categorically *impossible* is far higher than they are likely to observe and document for many generations. Second, even if we had both detailed descriptions and computational analyses for the phonology of every language currently spoken, this would only constitute a small and likely unrepresentative fraction of all human languages spoken in the last 100,000–200,000 years (Plank 2007, Bowers 2011). Third, even if we had reasonably complete analyses on both current and past human languages, for these data to be strongly informative about what is likely or even possible in potential future human languages requires either good reasons to believe such data are representative of all possible human languages or a dynamic model of how languages change over long timescales that would allow us to make such predictions (see e.g. Maslova 2000, Daniel 2007, Bickel 2011, Croft *et al.* 2011, Cysouw 2011).

The final reason for rejecting (38b) is that it doesn't have a good track record. Work on the subregular hypothesis has not *predicted* the typology of complexity; instead, the particular choice of exactly which class is hypothesised to be an upper bound has been post hoc, and driven by the most complex phenomena known to researchers at the time. The strengths of formal language theoretic methods lie in picking out well-defined classes of patterns and functions, and identifying formal conditions under which such patterns and functions are efficiently learnable. In order to make typological predictions, these strengths must be coupled with other linking hypotheses, which in previous literature have been nativist (38b). Nevertheless, since an empirically adequate theory must predict that

actually attested languages *do* exist, these methods are metatheoretically useful for identifying what a well-specified theory should predict to be *possible*.

While it may traditionally have been considered a virtue in generative linguistics to attempt to predict all and *only* observed typological data, in the context of model selection the phenomenon of ‘memorising’ observed data and failing to make correct generalisations that predict future data is called overfitting – regardless of whether the inductive learner in question is that of a child reasoning about an unknown but partially observed ‘set of possible strings’ or that of a scientist reasoning about a ‘set of possible human languages’ (see Perfors *et al.* 2011, Piantadosi & Gibson 2014 and Rasin *et al.* ms for review, as well as discussion of the relevance of model selection to the role of typology and/or nativism in linguistic theory development and comparison). In either case, such hypotheses should, *ceteris paribus*, be normatively dispreferred relative to simpler hypotheses that predict as yet unobserved data – or at least don’t rule it out.

We offer the following suggestions in place of (38b). First, rather than assuming that the absence of some type of phenomenon is best explained by arbitrary innate constraints, linguists should preferentially consider explanations that offer an independently motivated reason for that absence. Second, per the conclusions of Piantadosi & Gibson (2014) and Tily & Jaeger (2011), multiple methodologies should be used to identify and test putative linguistic universals, and any linguistic universal proposed should be presented alongside an explicit measure of evidential strength.

Turning to (38a), contemporary literature on typology and its relation to linguistic theory has shown the value of considering phylogenetic and areal effects (see e.g. discussion of word-order typology in Bickel 2007: 241), language-specific historical contingencies (Nichols 1992, Blevins 2004, Hansson 2008, Harris 2008) and differential learnability (e.g. Hayes & Wilson 2008, Culbertson 2012, Moreton & Pater 2012, Stanton 2016)¹⁷ as productive alternatives to innate constraints. We independently need explanations and theories of each of these domains, and each is a strictly simpler type of explanation of typological commonality and rarity than the non-specific, catch-all assumption of innate knowledge or constraints.

It may traditionally have been considered a virtue for generative linguistic theories to explain as much about observed typological data as possible. However, in light of all of the mechanisms we now know can and do affect observed linguistic typology and how little we know about the origins of long-distance phonological processes in particular, it would be surprising if a recently explored dimension of typological variation were entirely or largely explained by just one factor. This is especially the case when that factor is an innate constraint whose existence and specificity are not independently motivated or well-evidenced. In the language of model

¹⁷ Note that, in contrast to e.g. Heinz & Idsardi (2011, 2013), these explore relationships between relative learnability and relative *rarity*, not impossibility.

selection, (i) each previously mentioned non-nativist cause is a hypothesis with higher prior probability than a phonology-specific innate constraint of arbitrary complexity and weak independent motivation, and (ii) each of these non-nativist causes is constrained in terms of what kind of data it predicts and can explain, whereas an innate constraint could be invoked to explain almost any typological observation (see e.g. discussion of the ‘size principle’ in Tenenbaum & Griffiths 2001).¹⁸ We conclude that phonologists interested in explaining the typology of complex long-distance phonological processes would profit from exploring hypotheses about why they arise in the first place in the languages that have them and why they have the relative distributions they do (Bickel 2007).

6.4 Evaluating phonological theories and computational phonology

Independent of explaining typology *per se*, exactly fitting observed typology has been commonly used for evaluating generative linguistic theories. We argue that there are better ways of evaluating theories than the all-and-only criterion. We also emphasise that computational phonology has a key role to play: while the goal of identifying maximally restrictive theories of grammar has commonly been used to frame the significance of contemporary work on the subregular hierarchies of languages and relations, this sells short what contemporary work in computational phonology has to offer theory development and evaluation. For this reason, the question of how to evaluate theories of grammar without the all-and-only criterion is particularly relevant to this article.

In terms of model selection (Perfors *et al.* 2011, Rasin *et al.* ms), the all-and-only criterion neglects consideration of both a theory’s PRIOR PROBABILITY as a scientific hypothesis, plus the non-typological predictions it makes. Relevant sources of evidence for evaluating grammatical theories include results and debates concerning psychologically plausible representations and architectures (e.g. Marslen-Wilson 1973, Marslen-Wilson & Tyler 1987, Tanenhaus *et al.* 1995), computational learnability and predictions of learning errors (Tesar 1995, Heinz *et al.* 2009, Perfors *et al.* 2011, Rogers & Pullum 2011, Chandlee *et al.* 2014, Piantadosi *et al.* 2016, among many others), as well as linguistically significant distinctions (e.g. ‘Do phonological grammars count beyond two?’). Significantly, these are all sources of evidence *already* in use: phonologists bereft of the all-and-only criterion are hardly without other means of theory evaluation.

¹⁸ This second problem could be mitigated by a theory of what kinds of innate, phonology-specific constraints exist and do not exist, as well as of what kinds of patterns such a theory cannot predict. This theory should be principally motivated by evidence and methods outside of the kinds of linguistic observations the theory is intended to explain in the first place; for example, it might be motivated by what is known about domain-specificity and development.

7 Conclusion

Jardine (2016) claims that the minimum expressivity required to describe tonal patterns exhibiting unbounded circumambience is categorically greater than that needed to represent any patterns of segmental phonology, which are claimed to require at most weakly deterministic regular functions. Existing evidence for potentially unbounded circumambient segmental processes in Yaka and Sanskrit has been minimised, and these patterns have been construed as aberrant (Jardine 2016, Heinz 2018). We have shown that ATR harmony in Tutrugbu clearly satisfies the definition of unbounded circumambience according to the criteria laid out in Jardine (2016). Moreover, we have noted a variety of other languages that exhibit segmental patterns requiring similar expressivity. Altogether, there is strong evidence that segmental patterns require no less expressivity than tonal patterns.

The claim advanced in this article, that segmental phonology requires the same expressivity as tonal phonology, has significant implications for the continued evaluation of the subregular hypothesis and the expressivity of phonology more generally. Formal language theory provides a framework within which to frame these questions and address issues like the computational properties of human sound patterns. Moreover, we recognise that work like Heinz (2011a, b), Heinz & Lai (2013) and Jardine (2016) has stimulated a growing body of work examining phonology from a different point of view, and we hope further discussion of empirical patterns like ATR harmony in Tutrugbu and others noted here will facilitate a deeper understanding of human sound patterns, their expressivity and their distribution among the world's languages.

REFERENCES

- Anderson, Stephen R. (1974). *The organization of phonology*. New York: Academic Press.
- Archangeli, Diana & Douglas Pulleyblank (2007). Harmony. In de Lacy (2007). 353–378.
- Avcu, Enes (2018). Experimental investigation of the subregular hypothesis. *WCCFL* 35. 77–86.
- Baković, Eric (2000). *Harmony, dominance and control*. PhD dissertation, Rutgers University.
- Bickel, Balthasar (2007). Typology in the 21st century: major current developments. *Linguistic Typology* 11. 239–251.
- Bickel, Balthasar (2011). Statistical modeling of language universals. *Linguistic Typology* 15. 401–413.
- Bickmore, Lee S. & Nancy C. Kula (2013). Ternary spreading and the OCP in Copperbelt Bemba. *Studies in African Linguistics* 42. 101–132.
- Blevins, Juliette (2004). *Evolutionary Phonology: the emergence of sound patterns*. Cambridge: Cambridge University Press.
- Bobuafor, Mercy (2013). *A grammar of Tafi*. PhD thesis, University of Leiden.

- Bowern, Claire (2011). Out of Africa? The logic of phoneme inventories and founder effects. *Linguistic Typology* **15**. 207–216.
- Bowler, Margit (2013). Majority rules effects in Warlpiri vowel harmony. Ms. <https://www.academia.edu/38700903>.
- Bowler, Margit & Jesse Zymet (2019). A count effect in Warlpiri vowel harmony. Poster presented at the 27th Manchester Phonology Meeting. Available (April 2020) at http://linguistics.berkeley.edu/~jzymet/MFM2019_2.
- Casali, Roderic F. (2003). [ATR] value asymmetries and underlying vowel inventory structure in Niger-Congo and Nilo-Saharan. *Linguistic Typology* **7**. 307–382.
- Casali, Roderic F. (2008). ATR harmony in African languages. *Language and Linguistics Compass* **2**. 496–549.
- Casali, Roderic F. (2012). [+ATR] dominance in Akan. *Journal of West African Languages* **39:1**. 33–59.
- Chandlee, Jane (2014). *Strictly local phonological processes*. PhD dissertation, University of Delaware.
- Chandlee, Jane, Angeliki Athanasopoulou & Jeffrey Heinz (2012). Evidence for classifying metathesis patterns as subsequential. *WCCFL* **29**. 303–309.
- Chandlee, Jane, Rémi Eyraud & Jeffrey Heinz (2014). Learning Strictly Local subsequential functions. *Transactions of the Association for Computational Linguistics* **2**. 491–503.
- Chandlee, Jane & Jeffrey Heinz (2012). Bounded copying is subsequential: implications for metathesis and reduplication. In *Proceedings of the 12th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology*. Stroudsburg, PA: Association for Computational Linguistics. 42–51.
- Chandlee, Jane & Jeffrey Heinz (2018). Strict locality and phonological maps. *LI* **49**. 23–60.
- Chandlee, Jane, Jeffrey Heinz & Adam Jardine (2018). Input Strictly Local opaque maps. *Phonology* **35**. 171–205.
- Chomsky, Noam & Morris Halle (1968). *The sound pattern of English*. New York: Harper & Row.
- Comrie, Bernard, Martin Haspelmath & Balthasar Bickel (2015). The Leipzig glossing rules. <https://www.eva.mpg.de/lingua/resources/glossing-rules.php>.
- Croft, William, Tanmoy Bhattacharya, Dave Kleinschmidt, D. Eric Smith & T. Florian Jaeger (2011). Greenbergian universals, diachrony, and statistical analyses. *Linguistic Typology* **15**. 433–453.
- Culbertson, Jennifer (2012). Typological universals as reflections of biased learning: evidence from artificial language learning. *Linguistics and Language Compass* **6**. 310–329.
- Cysouw, Michael (2011). Understanding transition probabilities. *Linguistic Typology* **15**. 415–431.
- Daniel, Michael (2007). Representative sampling and typological explanation: a phenomenological lament. *Linguistic Typology* **11**. 69–78.
- de Lacy, Paul (ed.) (2007). *The Cambridge handbook of phonology*. Cambridge: Cambridge University Press.
- Dimmendaal, Gerrit J. (1983). *The Turkana language*. Dordrecht: Foris.
- Elgot, C. C. & J. E. Mezei (1965). On relations defined by generalized finite automata. *IBM Journal of Research and Development* **9**. 47–68.
- Essegbey, James (2009). Noun classes in Tutsugbu. *Journal of West African Languages* **36**. 37–56.
- Essegbey, James (2019). *Tutsugbu (Nyangbo) language and culture*. Leiden & Boston: Brill.
- Finley, Sara (2008). *Formal and cognitive restrictions on vowel harmony*. PhD dissertation, Johns Hopkins University.

- Gainor, Brian, Regine Lai & Jeffrey Heinz (2012). Computational characterizations of vowel harmony patterns and pathologies. *WCCFL* 29. 63–71.
- Hansson, Gunnar Ólafur (2008). Diachronic explanations of sound patterns. *Language and Linguistics Compass* 2. 859–893.
- Hantgan, Abbie & Stuart Davis (2012). Bondu-so vowel harmony: a descriptive analysis with theoretical implications. *Studies in African Linguistics* 41. 187–212.
- Harris, Alice C. (2008). On the explanation of typologically unusual structures. In Jeff Good (ed.) *Linguistic universals and language change*. Oxford: Oxford University Press. 54–76.
- Hayes, Bruce & Colin Wilson (2008). A maximum entropy model of phonotactics and phonotactic learning. *LI* 39. 379–440.
- Heinz, Jeffrey (2011a). Computational phonology. Part I: Foundations. *Language and Linguistics Compass* 5. 140–152.
- Heinz, Jeffrey (2011b). Computational phonology. Part II: Grammars, learning, and the future. *Language and Linguistics Compass* 5. 153–168.
- Heinz, Jeffrey (2018). The computational nature of phonological generalizations. In Larry M. Hyman & Frans Plank (eds.) *Phonological typology*. Berlin & Boston: De Gruyter Mouton. 126–195.
- Heinz, Jeffrey & William J. Idsardi (2011). Sentence and word complexity. *Science* 333. 295–297.
- Heinz, Jeffrey & William J. Idsardi (2013). What complexity differences reveal about domains in language. *Topics in Cognitive Science* 5. 111–131.
- Heinz, Jeffrey, Gregory M. Kobele & Jason Riggle (2009). Evaluating the complexity of Optimality Theory. *LI* 40. 277–288.
- Heinz, Jeffrey & Regine Lai (2013). Vowel harmony and subsequentity. In András Kornai & Marco Kuhlmann (eds.) *Proceedings of the 13th Meeting on the Mathematics of Language*. Sofia: Association for Computational Linguistics. 52–63.
- Hulden, Mans (2017). Formal and computational verification of phonological analyses. *Phonology* 34. 407–435.
- Hulst, Harry van der & Jeroen van de Weijer (1995). Vowel harmony. In John A. Goldsmith (ed.) *The handbook of phonological theory*. Cambridge, Mass.: Blackwell. 495–534.
- Hyman, Larry M. (1998). Positional prominence and the ‘prosodic trough’ in Yaka. *Phonology* 15. 41–75.
- Jardine, Adam (2016). Computationally, tone is different. *Phonology* 33. 247–283.
- Johnson, C. Douglas (1972). *Formal aspects of phonological description*. The Hague & Paris: Mouton.
- Kalinowski, Cristin (2009). Multidirectional vowel harmony in Esimbi. *BLS* 35. 147–155.
- Kaplan, Ronald M. & Martin Kay (1994). Regular models of phonological rule systems. *Computational Linguistics* 20. 331–378.
- Kari, Ethelbert E. (1997). *Degema*. Munich: Lincom Europa.
- Kimper, Wendell (2012). Harmony is myopic: reply to Walker 2010. *LI* 43. 301–309.
- Kula, Nancy C. & Lee S. Bickmore (2015). Phrasal phonology in Copperbelt Bemba. *Phonology* 32. 147–176.
- Lamont, Andrew (2019). Majority Rule in Harmonic Serialism. In Katherine Hout, Anna Mai, Adam McCollum, Sharon Rose & Matt Zaslansky (eds.) *Supplemental proceedings of the 2018 Annual Meeting on Phonology*. <https://doi.org/10.3765/amp.v7i0.4546>.
- Lesley-Neuman, Diane (2012). Morpho-phonological levels and grammaticalization in Karimojong: a review of the evidence. *Studies in African Linguistics* 41. 99–169.
- Lionnet, Florian (2016). *Subphonemic teamwork: a typology and theory of cumulative coarticulatory effects in phonology*. PhD thesis, University of California, Berkeley.

- 254 Adam G. McCollum, Eric Baković, Anna Mai and Eric Meinhardt
- Lombardi, Linda (1999). Positional faithfulness and voicing assimilation in Optimality Theory. *NLLT* 17. 267–302.
- Luo, Huan (2017). Long-distance consonant agreement and subsequenceality. *Glossa* 2(1):52. <http://doi.org/10.5334/gigl.42>.
- McCollum, Adam G., Eric Baković, Anna Mai & Eric Meinhardt (2020). On the existence of non-myopic harmony. Ms, Rutgers University & University of California San Diego.
- McCollum, Adam G. & James Essegbey (2018). Unbounded harmony is not always myopic: evidence from Tutrugbu. *WCCFL* 35. 251–258.
- McCollum, Adam G. & James Essegbey (2020). Initial prominence and progressive vowel harmony in Tutrugbu. *Phonological Data and Analysis* 2:3. 1–37. <https://doi.org/10.3765/pda.v2art3.14>.
- Marslen-Wilson, William (1973). Linguistic structure and speech shadowing at very short latencies. *Nature* 244. 522–523.
- Marslen-Wilson, William & Lorraine Komisarjevsky Tyler (1987). Against modularity. In Jay L. Garfield (ed.) *Modularity in knowledge representation and natural-language understanding*. Cambridge, Mass.: MIT Press. 37–62.
- Mascaró, Joan (2019). On the lack of evidence for nonmyopic harmony. *LI* 50. 862–872.
- Maslova, Elena (2000). A dynamic approach to the verification of distributional universals. *Linguistic Typology* 4. 307–333.
- Mielke, Jeff (2007). PBase: a database of phonological patterns. Available (April 2020) at <https://pbase.phon.chass.ncsu.edu>.
- Mohri, Mehryar (1997). Finite-state transducers in language and speech processing. *Computational Linguistics* 23. 269–311.
- Moreton, Elliott & Joe Pater (2012). Structure and substance in artificial-phonology learning. Part 1: Structure. *Language and Linguistics Compass* 6. 686–701.
- Nevins, Andrew (2010). *Locality in vowel harmony*. Cambridge, Mass.: MIT Press.
- Nichols, Johanna (1992). *Linguistic diversity in space and time*. Chicago: University of Chicago Press.
- Noske, Manuela (1996). [ATR] harmony in Turkana. *Studies in African Linguistics* 25. 61–99.
- Noske, Manuela (2000). [ATR] harmony in Turkana: a case of Faith Suffix \gg Faith Root. *NLLT* 18. 771–812.
- Novelli, Bruno (1985). *A grammar of the Karimojong language*. Berlin: Reimer.
- O'Hara, Charlie & Caitlin Smith (2019). Computational complexity and sour-grapes-like patterns. In Katherine Hout, Anna Mai, Adam McCollum, Sharon Rose & Matt Zaslansky (eds.) *Supplemental proceedings of the 2018 Annual Meeting on Phonology*. <https://doi.org/10.3765/amp.v7i0.4502>.
- Padgett, Jaye (1995). Partial class behavior and nasal place assimilation. In Keiichiro Suzuki & Dirk Elzinga (eds.) *Proceedings of the 1995 Southwestern Workshop on Optimality Theory (SWOT)*. Tucson: Department of Linguistics, University of Arizona. 145–183.
- Pater, Joe (2018). Substance matters: a reply to Jardine (2016). *Phonology* 35. 151–156.
- Payne, Amanda (2017). All dissimilation is computationally subsequential. *Lg* 93. e353–e371. <http://dx.doi.org/10.1353/lan.2017.0076>.
- Perfors, Amy, Joshua B. Tenenbaum & Terry Regier (2011). The learnability of abstract syntactic principles. *Cognition* 118. 306–338.
- Piantadosi, Steven T. & Edward Gibson (2014). Quantitative standards for absolute linguistic universals. *Cognitive Science* 38. 736–756.
- Piantadosi, Steven T., Joshua B. Tenenbaum & Noah D. Goodman (2016). The logical primitives of thought: empirical foundations for compositional cognitive models. *Psychological Review* 123. 392–424.

- Plank, Frans (2007). Extent and limits of linguistic diversity as the remit of typology – but through constraints on WHAT is diversity limited? *Linguistic Typology* **11**. 43–68.
- Prince, Alan (2007). The pursuit of theory. In de Lacy (2007). 33–60.
- Rasin, Ezer, Iddo Berger, Nur Lan & Roni Katzir (ms). Learning rule-based morpho-phonology. Available (April 2020) at <http://compling.tau.ac.il/papers/RasinBergerLanKatzir%202018%20Learning%20rule-based%20morpho-phonology.pdf>.
- Riggle, Jason (2004). *Generation, recognition, and learning in finite-state Optimality Theory*. PhD dissertation, University of California, Los Angeles.
- Roche, Emmanuel & Yves Schabes (1997). *Finite-state language processing*. Cambridge, Mass.: MIT Press.
- Rogers, James & Geoffrey K. Pullum (2011). Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language and Information* **20**. 329–342.
- Rose, Sharon (2018). ATR vowel harmony: new patterns and diagnostics. In Gillian Gallagher, Maria Gouskova & Sora Heng Yin (eds.) *Proceedings of the 2017 Annual Meeting on Phonology*. <http://dx.doi.org/10.3765/amp.v5i0.4254>.
- Rottland, Franz & Laura A. Otaala, (1983). Mid-vowel assimilation in Teso-Turkana. In Rainer Voßen & Marianne Bechhaus-Gerst (eds.) *Nilotic studies: proceedings of the international symposium on languages and history of the Nilotic peoples, Cologne, January 4–6, 1982*. Berlin: Reimer. 169–181.
- Ryan, Kevin M. (2017). Attenuated spreading in Sanskrit retroflex harmony. *LI* **48**. 299–340.
- Schröder, Helga & Martin Schröder (1987). Vowel harmony in Toposa. *Afrikanistische Arbeitspapiere* **12**. 27–35.
- Smith, Caitlin & Charlie O’Hara (2019). Formal characterizations of true and false sour grapes. In Gaja Jarosz, Max Nelson, Brendan O’Connor & Joe Pater (eds.) *Proceedings of the Society for Computation in Linguistic (SCiL) 2019*. Vol. 2. 338–341. <https://scholarworks.umass.edu/scil/vol2/iss1/41>.
- Stanton, Juliet (2016). Learnability shapes typology: the case of the midpoint pathology. *Lg* **92**. 753–791.
- Stanton, Juliet (2020). Gurindji nasal cluster dissimilation as trigger deletion. *JL* **56**. 157–195.
- Tanenhaus, Michael K., Michael J. Spivey-Knowlton, Kathleen M. Eberhard & Julie C. Sedivy (1995). Integration of visual and linguistic information in spoken language comprehension. *Science* **268**. 1632–1634.
- Tenenbaum, Joshua B. & Thomas L. Griffiths (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences* **24**. 629–640.
- Tesar, Bruce (1995). *Computational Optimality Theory*. PhD dissertation, University of Colorado, Boulder.
- Tily, Harry & T. Florian Jaeger (2011). Complementing quantitative typology with behavioral approaches: evidence for typological universals. *Linguistic Typology* **15**. 497–508.
- Vaux, Bert (2008). Why the phonological component must be serial and rule-based. In Bert Vaux & Andrew Nevins (eds.) *Rules, constraints, and phonological phenomena*. Oxford: Oxford University Press. 20–60.
- Walker, Rachel (2010). Nonmyopic harmony and the nature of derivations. *LI* **41**. 169–179.
- Wilson, Colin (2003). *Analyzing unbounded spreading with constraints: marks, targets, and derivations*. Ms, University of California, Los Angeles.
- Wilson, Colin (2006). Unbounded spreading is myopic. Paper presented at the Phonology Fest Workshop on Current Perspectives on Phonology, Indiana University.
- Wit, Gerrit de (2015). *Liko phonology and grammar: a Bantu language of the Democratic Republic of the Congo*. PhD dissertation, University of Leiden.
- Yip, Moira (2002). *Tone*. Cambridge: Cambridge University Press.