Computationally, tone is different[*]

Adam Jardine

November 25, 2015

1 Introduction

This paper establishes and characterises a typological difference between segmental and tonal phonology: UNBOUNDED CIRCUMAMBIENT PROCESSES are well-attested in tone but rare in segmental phonology. An unbounded circumambient process is one in which triggers or blockers appear on *both sides* of a target, and there is *no bound*, on *either side*, on the distance between these triggers or blockers and the target. This paper argues that these processes are more complex than those that are commonly attested in segmental phonology, and so the asymmetry can be characterised in a unified way by positing that tone is more computationally complex than segmental phonology.

A common unbounded circumambient process in tonal phonology is UNBOUNDED TONAL PLATEAUING (henceforth UTP; Kisseberth and Odden, 2003; Hyman, 2011), in which any number of tone-bearing units in between two underlying high tones also become high. A simple example from Luganda (Hyman et al., 1987; Hyman and Katamba, 2010; Hyman, 2011) is given below, with high toned vowels (both underlying and surface) marked with an acute accent (ó) and the 'plateau' underlined:

(1) UTP in Luganda (Hyman, 2011, p.231 (52))
/bikópo byaa-walúsiimbi/ $\rightarrow$ bikópó byáá-wálúsiimbi
cups of-Walusimbi
'the cups of Walusimbi'

UTP is an unbounded circumambient process because the triggering H tones can be any distance away from the affected tone-bearing units. Hyman (2011) observes that

[*]Manuscript to be published in *Phonology*. The author would first like to thank Jeffrey Heinz for going over countless versions of this paper, as well as three anonymous reviewers and associate editor who provided thoughtful and detailed critiques and advice. Thanks also go to Larry Hyman, William Idsardi, Kevin Ryan, the computational linguistics group at the University of Delaware, the members of Jeffrey Heinz's 2012 phonology seminar, and an inquisitive audience at NECPHON 7, all of whom provided valuable feedback over the course of the project. Although many helped with this paper, all errors are my own.

UTP is commonly attested in tone, but similar plateauing effects are, with one exception, unattested in segmental phonology.

The first contribution of this paper is to document this asymmetry in detail, and, by comparing UTP to the Sour Grapes vowel harmony pathology (Baković, 2000; Wilson, 2003; McCarthy, 2010; Heinz and Lai, 2013), show that it is part of a larger generalisation: unbounded circumambient processes are well-attested tone, but extremely rare in segmental phonology. The second contribution is to show that UTP and Sour Grapes, by virtue of being unbounded circumambient, are *formally* similar. This is because for these processes, each target must 'look ahead' in either direction to see crucial information in the environment. It is argued that this places unbounded circumambient processes outside of the WEAKLY DETERMINISTIC class of maps, a complexity class in FORMAL LANGUAGE THEORY (FLT) defined in terms of FINITE STATE TRANSDUCERS. The third contribution of this paper is then to understand the unbounded circumambient asymmetry in terms of such classes of maps. Previous work has found segmental processes to be at most weakly deterministic (Chandlee, 2014; Chandlee et al., 2012; Heinz and Lai, 2013; Payne, 2014). The unbounded circumambient asymmetry can thus be captured in terms of a complexity bound on segmental phonology: segmental phonology is restricted to weakly deterministic maps, but tone is not.

The structure of this paper is as follows. §2 establishes the unbounded circumambient asymmetry between tone and segmental phonology. §3 introduces the FLT notions of complexity and reviews previous work applying them to phonology. §4 shows that unbounded circumambient processes are not subsequential, and §5 argues that they are not weakly deterministic. This leads to the computational characterisation proposed by this paper: segmental processes are at most weakly deterministic, but tone is not restricted in this way. It is also discussed how two potential exceptions to this characterisation, Sanskrit *n*-retroflexion and KiYaka vowel harmony, fit into this proposal. §6 discusses how Optimality Theory does not offer a unified way of characterising the typological asymmetry. §7 concludes, and mathematical definitions and a proof that UTP is neither left- nor right-subsequential are given in an appendix (§8).

2 The Unbounded Circumambient Asymmetry

This section defines UNBOUNDED CIRCUMAMBIENT PROCESSES and shows how they are common in tonal phonology but extremely rare in segmental phonology. CIRCUMAMBIENT refers to a process whose application is dependent on the existence of triggers or blockers on both sides of a target; an UNBOUNDED circumambient process is one in which there is *no bound*, on *either side*, on the distance between these triggers/blockers and the target. These terms are discussed in detail below in §2.1.

The bulk of the evidence for the asymmetry comes from Unbounded Tone Plateauing (Kisseberth and Odden, 2003; Hyman, 2011), noted by Hyman (2011) as a tonal process with no common correlate in segmental phonology. §2.2 surveys attestations of UTP in the tonal literature, as a thorough documentation of Hyman (2011)'s claim. §2.3 reviews two known segmental unbounded circumambient processes of mid-vowel harmony in KiYaka (Hyman, 1998, 2011) and *n*-retroflexion in Sanskrit (Whitney, 1889; Macdonell, 1910; Schein and Steriade, 1986; Hansson, 2001; Graf, 2010).

§§2.4 and 2.5 summarize related generalizations in other typological work which support the conclusion that unbounded circumambient processes are rare in segmental phonology, and discuss why Sanskrit and KiYaka are exceptional for these generalizations as well. Thus, the asymmetry is not confined to UTP but really the *class* of unbounded circumambient processes. This is highlighted in §2.5 by way of the Sour Grapes vowel harmony pathology (Wilson, 2003; McCarthy, 2010; Heinz and Lai, 2013), which in Copperbelt Bemba H-spreading (Bickmore and Kula, 2013, 2015) has an attested correlate in tone.

2.1 Definition of 'unbounded circumambient'

A precise definition of UNBOUNDED CIRCUMAMBIENT PROCESS is given in (2). Crucially, this property is atheoretical and thus agnostic to specific theories of representation and processes.

(2) An UNBOUNDED CIRCUMAMBIENT PROCESS is a process for which:

a. its application is dependent on information (i.e., the presence of a *trigger* or *blocker*) on *both* sides of the target, and
b. *on both sides*, there is no bound on how far this information may be from the target

To help illustrate the concept, a rewrite rule representation of this type of process is given in (3). In (3), X and Y are nonempty, they surround the target, and there is no bound on the distance between them.

(3) $A \rightarrow B \;/\; X(U)\_\_(V)Y$ (X and Y are nonempty, U and V may be of any length)

An example is the imaginary rule in (4):

(4) $[+\text{syl}] \rightarrow [+\text{back}] \;/ \begin{bmatrix} +\text{syl} \\ +\text{back} \end{bmatrix} (U) \_\_ \; (V) \begin{bmatrix} +\text{syl} \\ +\text{back} \end{bmatrix}$ (U and V are any string of segments)

For the rule in (4), whether a $[-\text{back}]$ vowel becomes $[+\text{back}]$ depends on the presence of $[+\text{back}]$ vowels on both sides of the target. However, these $[+\text{back}]$ vowels may be separated from the target by strings U and V. Because U and V can be of any length, the process must be able to 'look behind' for a $[+\text{back}]$ segment *over any distance* in the left context (i.e., over U) *and* 'look ahead' over any distance for a [+back] vowel in the right context (i.e., over V).

Importantly, X or Y might contain 'blocking' information which prevents the process from applying. An imaginary such rule is given in (5):

(5) [−sonorant] → [+nasal] / [+nasal] (U) __
(except in this situation: [+nasal] (U) __ (V) [+nasal])

Here, the target will nasalise only when it is preceded by a [+nasal] segment and *not* followed by another [+nasal] segment. While (5) has been written out as a rule, such conditions are more intuitively expressed with Optimality Theory (henceforth OT; Prince and Smolensky, 1993, 2004) constraints.

Furthermore, the definition of unbounded circumambient applies also to autosegmental representations. From an autosegmental standpoint, the 'target' referred to in (2) is any unit affected by the changing of association lines and the distance from the 'trigger' is measured on the timing tier. These choices will become clear in §4.4.

Because 'unbounded' is critical to the definition in (2), we must set the criteria for an unbounded process. Intuitively, an unbounded process is one which operates over multiple units, like segments or tone-bearing units (TBUs), for which the correct generalisation does not refer to a bound on how many units over which it may operate. As linguists may differ as to what constitutes evidence for a process having 'no bound,' this paper considers the criteria in (6):

(6)
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

On its own, criterion (6a) does not constitute strong evidence for a process being unbounded, and thus no processes are included here which only meet (6a).

For some, (6b) is enough evidence that a process is unbounded, especially if there are examples of it operating over three or more units—as Kenstowicz (1994) put it, "phonological rules do not count past two" (p.372). In contrast, some researchers may only consider (6c) sufficient. However, the evidence presented here shows there is a typological asymmetry regardless of which criterion one considers. By either (6b) or (6c), unbounded circumambient processes are far more common in tone than in segmental phonology.

2.2 Unbounded Tone Plateauing

Having established the criteria for classifying a process as unbounded and circumambient, we can now discuss individual unbounded circumambient processes. This section surveys eight languages with some form of Unbounded Tone Plateauing (UTP, Kisseberth and Odden, 2003; Hyman, 2011), in which any number of L–toned or unspecified (∅) TBUs (here, TBU will be assumed to be the mora) surface as H if they are in between two Hs, but surface as L otherwise.

Kisseberth and Odden (2003) motivate UTP as a repair for a constraint against "toneless moras between Hs" (p.67). Hyman and Katamba (2010) formalise UTP in Luganda (which they refer to as 'H tone plateauing') as follows:

(7) 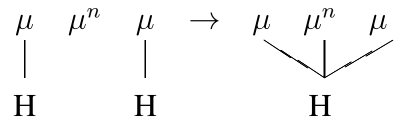

UTP thus fits the definition of an unbounded circumambient process because whether or not the process applies depends on two Hs that a) are on both sides of the affected TBUs and b) can be arbitrarily far away from any one of the affected TBUs.[1]

Using data from a number of sources, the following establishes Hyman (2011)'s claim that UTP is a well-attested tonal process, and that it is unbounded by criterion (6b) in all cases and by criterion (6c) in most. Data from Luganda (which is pointed out by Hyman) and Digo are examined in depth first, and then examples in other languages are briefly reviewed.

2.2.1 Luganda

UTP occurs in Luganda (Hyman et al., 1987; Hyman and Katamba, 2010) both word-internally and in the phrasal phonology. Luganda TBUs can be either H or unspecified (∅) underlyingly. Lexically, a L tone is inserted after a H, causing a falling HL contour when the H is on a final syllable. In some cases, these intermediate L tones contrast with ∅, but not with regards to UTP (technically, they delete in the UTP environment), so they will not factor into the discussion. Tones are marked using an acute accent [á] for underlying and surface H tones, with underlying unspecified and surface low tones unmarked. A circumflex [â] on the vowel marks a falling contour.[2]

(8) Luganda nouns (from Hyman and Katamba, 2010, (2),(3),(26),(52))

a.	/ki-kópo/	kikópo	'cup'
b.	/ki-sikí/	kisikî	'log'
c.	/ki-tabo/	kitabo	'book'
d.	/mu-tund-a/	mutunda	'seller' (from /-tund-/ 'to sell')
e.	/mu-tém-a/	mutéma	'chopper' (from /-tém-/ 'to chop')

[1]This can be contrasted with *bounded* plateauing, in which only one ∅ TBU becomes H in between two Hs. This pattern is attested, for example, in Kihunde (Goldsmith, 1990).

[2]The tones transcribed here are for what Hyman and Katamba call the 'intermediate' forms before the superimposition of phrasal boundary tones. The boundary tones are inserted late in the derivation and play no role in the phonology.

If a toneless noun and a noun with an underlying H tone form a compound, both nouns are pronounced as in isolation (9a). However, when both nouns have an underlying H, a 'plateau' of high tones occurs between them (9b). Autosegmental diagrams accompany the following examples to illustrate the process. In compounds, the plural /bi-/ is used for nouns from the IV /ki-/ noun class. It carries no tone.

(9) Luganda compounds (Hyman and Katamba, 2010, (26) & (52))

```
a.  /mu-tund-a/ + /bi-kópo/   →  mutunda-bikópo   'cup seller'
                       |                    |
                       H                    H
b.  /mu-tém-a/ + /bi-sikí/    →  mutémá-bísíkî    'log-chopper'
         |             |
         H             H             H
```

In (9b) there is a plateau over three unspecified TBUs. This shows target toneless TBUs three TBUs away from their triggers—the mora (represented by) /a/ in /mu-tém-a/ 'chopper' is two TBUs from its right trigger H /í/ in /bi-sikí/ 'log,' and the second /i/ in /bi-sikí/ is three TBUs from its left trigger /é/ in /mu-tém-a/. Thus Luganda satisfies the 'multiple unit' criterion for unboundedness in (6b).

UTP operates over syntactic phrases as well, satisfying (6c). Under certain morpho-syntactic conditions, noun-verb sequences can also form a domain for UTP. The noun in (10a) below is an adjunct and thus outside of the domain of plateauing. In (10b), however, the noun and verb form a phonological phrase, and plateauing occurs accross all TBUs in between the surviving Hs.

(10) Luganda verb+noun combinations (Hyman and Katamba, 2010, (14))

```
a.  /tw-áa-mú-láb-a walúsimbi/   →  tw-áá-mu-lab-a walúsimbi
        |   |  |        |               |/              |
        H   H  H        H               H               H
    'we saw him, Walusimbi'
b.  /tw-áa-láb-w-a walúsimbi/    →  tw-áá-láb-wá wálúsimbi
        |   |        |
        H   H        H                      H
    'we were seen by Walusimbi'
```

A more extreme example can be seen when proclitics are added to the noun. Proclitics generally do not have tone, as the lack of H in the following data shows:

(11) Luganda toneless proclitics (Hyman and Katamba, 2010, (33))

a. byaa=ba=mulondo '(it's) those of the Mulondos'
b. na=ku=byaa=ba=mulondo 'and on those of the Mulondos'

Luganda UTP 'sees' over these as well, no matter how many are stacked on to the noun:

(12) Luganda verb+noun sequences with proclitics (Hyman and Katamba, 2010, (35))

a. /tw-áa-láb-a byaa=walúsimbi/ → tw-áá-láb-á byáá-wálúsimbi

H H H H

'we saw those of Walusimbi'

b. /tw-áa-génd-a na=byaa=ba=walúsimbi/ → tw-áá-génd-á ná=byáá=bá=wálúsimbi

H H H H

'we went with those of Walusimbi'

These examples show UTP operating over a six-TBU span of toneless TBUs created by prefixation, illustrating triggers separated from their targets by five TBUs on each side. As such, Luganda UTP also satisfies criterion (6c) for unboundedness, and satisfies (6b) with five TBUs.

2.2.2 Digo

Digo verbs (Kisseberth, 1984) show complex interactions between underlying H tones, including UTP. Underlyingly, Digo is a privative H/∅ system. Both verb roots and affixes may carry a H tone, although this is not obligatory. A single H tone in a verbal H shifts to the end of the word, surfacing as a rising/falling pattern on the final two TBUs. This is illustrated with the addition of the third person plural object prefix /á/ to the toneless root /tsukur/ 'take' in (13b) (originating TBU of the H underlined); an autosegmental analysis is given in (14).

(13) Digo (Kisseberth, 1984, (29))

a. ni-na+tsukur-a
'I am taking'

b. ni-na+<u>a</u>-tsukǔr-â
'I am taking them'

(14) /ni-na+á-tsukur-a/ →ni-na+a-tsukǔr-â

H H

'I am taking them (=13b)'

The forms in (13) have the toneless first person prefix /ni/; if this is substituted for the toned third person singular subject prefix /á/, a pleateau occurs from the object prefix to the end of the root:

(15) <u>a</u>-na-<u>á</u>-tsúkúr-â 'He/she is taking them' (Kisseberth, 1984, (29))

Again, we see the presence of two H tones creating a long-distance plateau across the length of the root. The Digo example is complicated by tone shifting; Kisseberth analyses it as a two-step process in which the first H shifts to the initial vowel of a 'verbal complex' (marked with the '+' boundary), the second H shifts to the end of the word (as in (14)), and then the first H then triggering a plateau across the root between them.This analysis is illustrated by the AP derivation in (16) below.

(16) Underlying /á-na+á-tsukur-a/ 'He/she is taking them' (=15)
H H
Right shift a-na+a-tsukur-a
H H
Plateauing a-na+a-tsukur-a
H
Surface a-na+á-tsúkúr-â

The domain for both right shift and plateauing is larger than the verb—they also apply to verb+noun constructions. In the following, (17) shows right shift applying when there is only one H in the verb, and (18) shows plateauing over the phrase resulting from two Hs associated with the verb.

(17) a. ku+afŭnâ 'to chew'
b. nazi 'coconut'
c. ku+afun-a năzî 'to chew a coconut' (Kisseberth, 1984, (63))

(18) a. a-ka+tsúkú ts-â 'he has cleaned'
b. chi-ronda 'wound'
c. a-ka+tsúkú ts-á chí-róndâ 'he has cleaned a wound (Kisseberth, 1984, (65))

The phrase in (18c) results in a plateau over six TBUs, which shows target TBUs three TBUs away from their triggers and thus satisfies criterion (6b) for unboundedness. Also shown in (18c) is a plateau created over a syntactic domain, satisfying criterion (6c).

2.2.3 Other languages

Kisseberth and Odden (2003) cite a plateauing process in Xhosa similar to Digo. In Xhosa, underlying H tones shift to the antepenult. A phrase with two H tones shows a plateau between the first H and the second, shifted H on the antepenult.

(19) Xhosa (Kisseberth and Odden, 2003, pp. 67-8)
a. u-ku-qonóndis-a 'to emphasise'
b. ndi-fún-a 'I want'
c. ndi-fún ú-kú-qónónóndis-a 'I want to emphasise'

The second H shifts to the third /o/ in /ú-ku-qononondis-a/ 'to emphasise,' and becomes the right trigger for UTP; this shows triggers on both sides affecting targets four TBUs away. Kisseberth and Odden (2003) note that such processes are 'common in the Nguni languages' (Digo is not a member of this family). In Zulu (Laughren, 1984; Cassimjee and Kisseberth, 2001; Downing, 2001), another Nguni language, a single H shifts to the antepenult (if it originates on a prefix) or penult (if it originates on a stem). In forms with two Hs, a plateau forms between the Hs. In Zulu, the two H tones do not fuse; instead, the first spreads up to the second, creating a downstep (marked with !):

(20) Zulu (Yip, 2002, p. 158, citing Laughren (1984))

a. i-si-hlaːlo ‘seat’ → i-sí-hlaːlo
H H

b. ámàkhòsánà → ámákhós!ánà (no gloss)
H H H H

Laughren (1984) states that “the rule only applies to a H which is followed by a LH tonal sequence” (p. 221; Yip represents the LH as a single H, which is then downstepped), and while she only gives examples of the plateau operating over two TBUs, analyses the process as operating over an arbitrary number of TBUs.

Other examples of UTP can be found throughout Bantu. In KiYaka (also known as Yaka; Kidima, 1990, 1991), “all toneless syllables flanked by Hs become H by rightward spreading of the H to the left domain” (Kidima, 1991, p.44). The following show plateauing alternations. Kidima (1991) gives KiYaka tonal assignment a complex accentual analysis; the underlying Hs marked in the following data result from a tonal assignment rule.[3] In the following examples, *bakhoko* ‘chickens’ is not assigned a tone, prompting the alternations.

(21) KiYaka (Kidima, 1991, (68), p. 180)

a. bakhoko ba ngwaá̲si
chickens of uncle
‘Uncle’s chickens’

b. bakhoko ba kabeé̲nga
chickens of red
‘The red chickens’

c. bakhoko ba kabeé̲ngá bá ngwáá̲sí Málóó̲ngi
chickens of red of uncle Maloongi
‘Uncle Maloongi’s red chickens’

Example (21c) shows two plateaus, each with their target TBUs separated from their triggers by two TBUS.

UTP also occurs outside of Bantu. In Saramaccan (Roundtree, 1972; Good, 2004; McWhorter and Good, 2012), a creole spoken in Suriname, a phrasal version of UTP occurs accross words in certain syntactic configurations. This analysis follows Good (2004), who posits an underlying H/L/∅ distinction, as in the nouns below.

[3] A surface distinction between regular H and raised H—the latter occuring on accented syllables with an associated H—is ignored in these transcriptions.

(22) Saramaccan nouns (Good, 2004)
a. wómì 'man' /wómi/
b. mùjɛ́ɛ̀'woman' /mujɛ́ɛ/
c. wàjàmákà 'iguana' /wajamáka/
d. sɛ̀mbɛ̀'person' /sɛ̀mbɛ̀/
e. àmɛ̀ɛ̀ká 'American' /amɛɛká/

According to Roundtree (1972), "all changeable low tones [=Good (2004)'s ∅ TBUs] between highs in successive morphs in certain syntactic positions are changed to high..." (p. 314). One such syntactic position is an adjective-noun sequence. In the following, the final /o/ in /hánso/ 'handsome' is realised as low [ó] before /sɛ̀mbɛ̀/ 'person' (23a) but high before /wómi/ 'man' (23b) and /mujɛ́ɛ/ 'woman' (23c). Relevant vowels are emphasised in bold.

(23) Saramaccan phrases
a. /dí hánso sɛ̀mbɛ̀/ → díháns**ò** sɛ̀mbɛ̀
the hansome person "the hansome person" (Roundtree, 1972, p.315)
b. /dí hánso wómi/ → dí háns**ó** wómì
" " man "the handsome man" (Roundtree, 1972, p.324)
c. /dí hánso mujɛ́ɛ/ → dí háns**ó** m**ú**jɛ́ɛ̀
" " woman "the handsome woman" (Roundtree, 1972, p.316)
d. /dí wajamáka=dé á óbo/ → dí w**á**j**á**máká=dé á óbo
the iguana=there have eggs "the iguana there has eggs" (Good, 2004, p. 28)
e. /dí taánga amɛɛká wómi/ → dí tàáng**á** **á**m**ɛ́ɛ́**ká wómì
the strong American man "the strong American man" (McWhorter and Good, 2012, p.48)

Note that in (23b) the initial /u/ of /mujɛ́ɛ/ 'woman' also surfaces as H, illustrating plateauing over two TBUs. In (23e) a plateau occurs over four TBUs, the final /a/ of /taánga/ 'strong' and the first three vowels of /amɛɛká/ 'American,' thus Saramaccan UTP satisfies (6b), with its targets and triggers separated by three TBUs on each side.

Similar to the system in Saramaccan (and noted by Good) is the intonational phonology of the Uto-Aztecan language Papago (Hale and Selkirk, 1987), in which H tones are associated "to each stressed vowel and to all vowels in between"; they list examples of plateaus created in between stressed vowels three TBUs apart. Finally, Do and Kenstowicz (2011) discuss plateauing in between Hs in certain intonational phrases in South Kyungsang Korean, giving a spectrogram showing a plateau between the two H tones in *seccók khaliphonía* 'Western California' (pp. 3 & 11), which shows unspecified TBUs affected by triggers two TBUs away (on either side).

2.3 Unbounded circumambient processes in segmental phonology

Turning to segmental phonology, there are the only two potential attestations of unbounded circumambient processes of which the author is aware. One is Sanskrit *n*-

retroflexion (Whitney, 1889; Macdonell, 1910; Schein and Steriade, 1986; Hansson, 2001; Graf, 2010; Ryan, 2015). The other is ‘plateauing harmony’ in KiYaka (Hyman, 1998, 2011), which also has UTP, as explained in §2.2.3. Both satisfy the unboundedness criterion in (6b), but neither satisfy (6c).

2.3.1 Sanskrit

In Sanskrit *n*-retroflexion, an underlying alveolar /n/ becomes retroflex [ɳ] after retroflex /r,ʂ/, which can appear far to the left of the target /n/. In the following examples, both trigger and target will be highlighted with underlining:

(24) Sanskrit *n*-retroflexion (Hansson, 2001, p. 225, citing Schein and Steriade (1986))

	UR	SR	Gloss
a.	/iʂ̱-ṉaː-/	iʂ̱-ɳ̱aː-	‘seek (pres. stem)’
b.	/cakʂ̱-aːṉa-/	cakʂ̱-aːɳ̱a-	‘see (middle part.)’
c.	/kṟp-a-maːṉa-/	kṟp-a-maːɳ̱a-	‘lament (middle part.)’

One restriction on *n*-retroflexion potentially gives the process an unbounded circumambient quality. Hansson (2001) states that retroflexion fails “when there is also an /ʂ/ or /r/ *later* in the word” (p.230, emphasis original).[4] He cites the following data from (Macdonell, 1910) (second, blocking /ʂ/ or /r/ also underlined; syllabic rhotics are not transcribed for typographic clarity):[5]

(25) Blocking of Sanskrit *n*-retroflexion (Hansson, 2001, p. 225, citing Macdonell (1910))

	Attested	Unattested	Gloss
a.	pṟaː-ṉṟtyat	*pṟaː-ɳ̱ṟtyat	from -nrt- ‘dance’
b.	paṟi-ṉakʂ̱ati	*paṟi-ɳ̱akʂ̱ati	‘encompasses’
c.	-ṉiʂ̱ʈʰaː-	never *-ɳ̱iʂ̱ʈʰaː-	‘eminent’
d.	-ṉiʂ̱ʂ̱idʰ-	never *-ɳ̱iʂ̱ʂ̱idʰ-	‘gift’
e.	-ṉiṟɳiɟa-	never *-ɳ̱iṟɳiɟa-	‘adornment’
f.	-ṉṟmɳa-	never *-ɳ̱ṟmɳaː-	‘manhood’
g.	pṟa-ṉaŋkʂ̱yati	*pṟa-ɳ̱aŋkʂ̱yati	‘causes to dissappear’ (Monier-Williams, 1899)

Given the data in (24), one would expect, for example, that the underlying /n/ in (25b) [pari-nakʂati] ‘encompasses’ would surface as [ɳ], because it follows a trigger /r/ for *n*-retroflexion. However, it instead surfaces as [n], which is attributed to the ‘blocking’ /ʂ/ three segments to the right. Thus, Sanskrit *n*-retroflexion fits the definition for circumambient in a slightly different way than UTP: the crucial context includes the presence of a trigger on the left side and a blocker on the right. The evidence in the examples here also fit criterion (6b) for unboundedness; (25g) shows a blocker 3 segments away to the right

[4]There are additional restritions; see Ryan (2015) to a detailed discussion of Sanskrit retroflexion.

[5]Many thanks to Kevin Ryan for pointing (25g) out and for enlightening me about Sanskrit in general.

of the target, and (24c) shows a trigger 5 segments to the left. Thus, by criterion (6b), Sanskrit *n*-retroflexion is an unbounded circumambient process. However, Ryan (2015)'s study of several Sanskrit corpora finds no examples of retroflexion being blocked when a long vowel or multiple syllables intervene between the target and blocker. Thus, whether or not it is truly unbounded is the subject of some doubt. Furthermore, there is no evidence satisfying (6c), i.e., that the distance this second blocker can be from the target may be extended by a morphological or syntactic process.

2.3.2 KiYaka vowel harmony

The other attested unbounded circumambient segmental process is vowel height harmony in KiYaka. Hyman (2011) cites KiYaka (Hyman, 1998, where it is referred to as Yaka) as a unique example of vowel 'plateauing.' In KiYaka, the initial vowel of the perfective suffix /ile/ lowers to a mid [e] when the vowel in the stem is also mid (26c and d below). Otherwise, a progressive harmony converts the final /e/ to [i] ([l] and [d] alternate depending on the following vowel, and with [n] following a root nasal; thus the allomorph [ene] in (27b)). The lowering of this middle /i/ to [e] does not happen in the applicative suffix /ila/, which does not end in a mid vowel. The following examples show this alternation, with the /i/ to [e] change in question underlined:

(26) KiYaka mid-vowel 'plateauing' (Hyman, 2011, p. 501)

	Gloss	Root	+Applicative /ila/	+Perfect /ile/
a.	'obstruct'	/kik/	kik-ila	kik-idi
b.	'bind'	/kas/	kas-ila	kas-idi
c.	'pay attn.'	/keb/	keb-ila	keb-ele
d.	'clear brush'	/sol/	sol-ila	sol-ele

That this is part of a more general process lowering high vowels to mid if and only if they are between two mid vowels can be seen in (27). This process can take place at least over three vowels:

(27) KiYaka mid-vowel plateauing (Hyman, 1998, p. 19(6a,e,&f))

	Gloss	Stem	+Final Vowel /a/	+Perfect /ile/
a.	'to send'	/hit-ik/	hit-ik-a	hit-ik-idi
b.	'lower'	/bet-ilik/	bet-idik-a	bet-elek-ele
c.	'to do an about-face'	/kel-umuk/	kel-umuk-a	kel-omok-ene

As in UTP, the plateau of mid-vowels shows triggers on both sides of the target vowels. Examples (b) and (c) in (27) thus show triggers two vowels from the right and two vowels from the left from their targets. As such, KiYaka vowel harmony satisfies criterion (6b) for an unbounded process. However, according to Hyman (1998), plateauing alternations can only be seen with verb roots and the perfective /-ile/; it thus does not satisfy (6c).

2.4 Empirical summary

The preceding sections presented eight attestations of UTP, an unbounded circumambient process in tonal phonology, and two examples of separate unbounded circumambient processes in segmental phonology. All satisfied criterion (6b), with the greatest distance between target and trigger being five TBUs. Seven examples of UTP were shown to operate over domains extended by morphology or syntax, thus satisfying criterion (6c). Neither segmental process satisfied (6c), and in KiYaka, the distribution of the process was quite limited. Both satisfied (6b), although in KiYaka the greatest attested distance between trigger and target was three vowels, and in Sanskrit there was no evidence for blocking beyond the syllable following the target.

Most importantly, to the best of the author's knowledge, these are the only examples of such processes in segmental phonology. This is notable given the wide attestation of long-distance segmental processes in general, as documented in the comprehensive surveys on feature-spreading harmony (Rose and Walker, 2004), vowel harmony (Baković, 2000; Nevins, 2010), consonant harmony (Rose and Walker, 2004; Hansson, 2001, 2010), and consonantal disharmony (Suzuki, 1998; Bennett, 2013). In fact, Wilson (2003, 2006b), reviewing the typologies of nasal, emphasis, and vowel harmonies, characterises segmental spreading processes as 'myopic'. This means that even segmental spreading processes which affect multiple segments proceed in a local fashion, never 'looking ahead' beyond immediately adjacent segments. As unbounded circumambient processes by definition depend on information unboundedly far away from the target, any myopic process is not an unbounded circumambient process, and any unbounded circumambient spreading process is necessarily not myopic. KiYaka vowel harmony and Sanskrit *n*-retroflexion are thus exceptions to the myopic spreading generalisation; this highlights how atypical these two processes are.[6]

There thus is a typological asymmetry between tonal and segmental processes: unbounded circumambient patterns are extremely rare in segmental processes, but well-attested in tone, at the very least in the variants of UTP discussed. A comparison between proportions of circumambient unbounded processes found in typological surveys (such as those just mentioned) of segmental and tonal processes would be ideal, but comparable surveys of tonal processes, or even particular kinds of processes, do not exist (to the best of the author's knowledge). Regardless, the evidence reviewed in this paper clearly shows an asymmetry despite the absence of such surveys for tone.

It should be noted that bidirectional spreading processes, in which a feature spreads outwards from a single trigger in two directions, are common in segmental harmony. An example is Arabic emphasis spreading (Al Khatib, 2008), in which an emphatic gesture spreads in an unbounded fashion in both directions from an underlying emphatic segment. In the following example from Southern Palestinian Arabic, emphasis spreads to the left (28a), to the right (28b), and to both the left and right (28c) (the spread is blocked by high

[6] See Gafos (1999) and Hansson (2001) for arguments that *n*-retroflexion is best analysed as spreading.

front segments, such as /j/).[7]

(28) Palestinian Arabic (Al Khatib, 2008, (1))

a. /balːaːsˤ/ → [bˤaˤlˤːaˤːsˤ] 'theif'
b. /sˤajːaːd/ → [sˤaˤjːaːd] 'hunter'
c. /ʔatˤfaːl/ → [ʔaˤtˤfˤaˤːlˤ] 'children'

Other bidirectional processes are nasal spread in Capanahua and Southern Castillian (Safir, 1982) and stem-control analyses of vowel harmony (Baković, 2000) and consonant harmony (Hansson, 2001, 2010). While these processes apply in an unbounded fashion, and they operate in two directions, they only have one trigger, and thus are not circumambient.

2.5 Sour Grapes in tone

One final piece of evidence for the unbounded circumambient asymmetry regards the unattested 'sour grapes' vowel harmony pattern (henceforth 'Sour Grapes'; Baković, 2000; Wilson, 2003; McCarthy, 2010; Heinz and Lai, 2013).[8] Sour Grapes is a non-myopic harmony pattern predicted to exist by ranking permutations of classic OT with AGREE constraints, but not attested in segmental harmony. Sour Grapes is also unbounded circumambient, and, as will be discussed momentarily, a Sour Grapes-like process appears in tone.

Sour Grapes works as follows. Given a spreading [+F] feature (which targets underlying [–F] segments) and a blocking feature [!F], [–F] segments become [+F] after another [+F], provided that there is no blocking segment, [!F], following in the word:

(29)

a. [–F]^n → [–F]^n (no trigger, no blocker → no harmony)
b. ...[+F][–F]^n → ...[+F]**[+F]**^{n} (trigger, no blocker → **harmony**)
c. ...[+F][–F]^n[!F]... → ...[+F][–F]^n[!F]... (trigger, blocker → no harmony)

In other words, [+F] spreads if and only if it can spread all the way.[9] Note that Sour Grapes is not myopic: the spread of the [+F] has to 'look ahead' to see if there is a blocker before it can apply. As such, it is also an unbounded circumambient pattern, because the presence or absence of both [+F] and [!F] segments, which can be any distance apart, bears on the realisation of [–F] segments in between (this blocking aspect makes it circumambient in a similar way to Sanskrit *n*-retroflexion).

[7] Thanks to an anonymous reviewer for providing this example.

[8] The term 'sour grapes', originally due to Padgett (1995), refers to behaviours under certain formulations of OT in which either all features in a particular domain assimilate or none do. See also McCarthy (2010).

[9] This characterisation is thanks to an anonymous reviewer.

However, a Sour Grapes-like pattern does exist in tone. In Copperbelt Bemba (Bickmore and Kula, 2013, 2015), underlying H tones undergo one of two spreading processes, bounded ternary spreading or unbounded spreading. The latter is blocked by the presence of another H tone. In phrase-final forms, unbounded spreading applies to the rightmost H.

(30) Bemba unbounded spreading (Bickmore and Kula, 2013, (1)&(17))

	UR	SR	gloss
a.	/u-ku-tul-a/	ù-kù-tùl-à	'to pierce'
b.	/bá-ka-fik-a/	bá-ká-fíká	'they will arrive'
c.	/bá-ka-mu-londolol-a/	bá-ká-mú-lóóndólól-á	'they will introduce him/her'
d.	/tu-ka-páapaatik-a/	tù-kà-páápáátík-á	'we flatten'

Bounded spreading occurs when another H appears to the right. Bounded spreading obeys the Obligatory Contour Principle; it will spread up to two additional TBUs, maintaining at least one L TBU before the second H. All other intervening TBUs surface with a L tone.

(31) Bemba bounded spreading (Bickmore and Kula, 2013, (18), 2015, (8))

	UR	SR	gloss
a.	/bá-ka-pat-a kó/	bá-ká-pát-à kó	'they will hate'
b.	/bá-ka-londolol-a kó/	bá-ká-lóòndòlòl-à kó	'they will introduce them'
c.	/tu-ka-béleeng-el-an-a kó/	tù-kà-bélééng-él-àn-à kó	'we will read for e.o'
d.	/tu-ka-lás-a Kapembuá/	tù-kà-lás-á Kápèèmbwá	'we will hit Kapembwa'

The formalisations in (32) summarise the facts. When no Hs are present, as in (32a), all TBUs surface as L (c.f. (30a) ù-kù-tùl-à 'to pierce'). When one H is present, it spreads to all remaining TBUs in the domain (and the rest surface as low, as in (30d)). When two Hs are present (32c), the first only spreads to the next two TBUs (c.f. (31c) tù-kà-béléénggél-àn-à kó 'we will read for e.o').

(32) a. $\mu^n \rightarrow \begin{array}{c}\mu^n\\ |\\ \mathrm{L}\end{array}$ (c.f (30a)) b. $\begin{array}{ccc}\mu^m & \mu & \mu^n\\ & | & \\ & \mathrm{H} & \end{array} \rightarrow \begin{array}{ccc}\mu^m & \mu & \mu^n\\ | & \multicolumn{2}{c}{/\!/}\\ \mathrm{L} & \multicolumn{2}{c}{\mathrm{H}}\end{array}$ (c.f (30d)) c. $\begin{array}{cccc}\mu^m & \mu & \mu^n & \mu\\ & | & & |\\ & \mathrm{H} & & \mathrm{H}\end{array} \rightarrow \begin{array}{cccccc}\mu^m & \mu & \mu & \mu & \mu^{n-2} & \mu\\ | & \multicolumn{3}{c}{/\!/\!/} & | & |\\ \mathrm{L} & \multicolumn{3}{c}{\mathrm{H}} & \mathrm{L} & \mathrm{H}\end{array}$ (c.f (31c))

Note that, modulo the bounded spreading, the formalisations in (32) are almost identical to the Sour Grapes generalisations in (29). In other words, in Copperbelt Bemba the second H can be seen as a blocker for unbounded spread. This makes it an unbounded circumambient process, because the realisation of unspecified TBUs depends on the presence or absence of Hs on *both* sides which can be arbitrarily far away. As a tonal process, it does not seem particularly aberrant, in contrast to Sanskrit *n*-retroflexion and KiYaka vowel harmony. Thus, it provides more evidence for the unbounded circumambient asymmetry.

2.6 Empirical conclusion

This section has made clear an asymmetry between tonal and segmental processes. Unbounded circumambient processes—i.e., processes in which crucial information about the environment lie arbitrarily far away on both sides of the target—are well-attested in tone, but extremely rare in segmental phonology. The remainder of the paper is concerned with understanding this asymmetry in terms of the relative computational complexity of unbounded circumambient processes.

3 The Computational Complexity of Phonological Maps

A relevant measure of computational complexity can be found in Formal Language Theory (FLT). This section introduces FLT, shows how it relates to phonology, and introduces properties of FINITE-STATE TRANSDUCERS that allow us to measure the relative complexity of phonological processes.

3.1 Formal language complexity and cognitive complexity

A FORMAL LANGUAGE is a set of strings, and FLT studies the relationships between formal languages and the expressive power of grammars that describe them. FLT characterisations of natural language patterns have been argued to reflect domain-specific cognitive biases for and against patterns of a certain level of complexity. For example, the REGULAR class of formal languages is insufficient to describe English syntax, and English syntax is at least CONTEXT-FREE (Chomsky, 1956). Phonology, on the other hand, appears to be at most regular (Johnson, 1972; Kaplan and Kay, 1994).[10] This notion of complexity has been explicitly linked to cognitive complexity (Rogers and Pullum, 2011; Rogers and Hauser, 2010; Folia et al., 2010), and results from artificial language learning experiments provide evidence in support of the psychological reality of the context-free/regular division between syntax and phonology (Lai, 2012, 2015).

The goal of this paper is to use this notion of complexity to characterise the unbounded circumambient asymmetry discussed in §2. Instead of formal languages, however, the following sections study the relationships between regular MAPS, or relations between strings in which an input string is paired with at most one output string.[11] Analagous to the regular/context-free division for formal languages, maps can be classified according to their complexity. The distinctions important for this paper center around the property of SUBSEQUENTIALITY (Mohri, 1997), which can be defined in terms of FINITE-STATE

[10] For more on the regular/non-regular split between phonology and syntax, see Heinz (2011) and Heinz and Idsardi (2011, 2013).

[11] Cases of free variation are thus not maps, as one UR can be paired with multiple SRs. The focus of the present study is on subclasses of *maps*, and so such cases shall not be considered, but (finite) free variation can be studies in a similar way with the P-SUBSEQUENTIAL TRANSDUCERS of Mohri (1997) or the SEMI-DETERMINISTIC TRANSDUCERS of Beros and de la Higuera (2014).

TRANSDUCERS (FSTs). FSTs are idealised machines that match pairs of strings; they are described in more detail below. While FSTs in general can describe any regular map, subsequential FSTs describe a more restricted set of maps.

3.2 Overview: Formal language complexity and phonology

The classes of interest to this paper are the LEFT- and RIGHT-SUBSEQUENTIAL maps, the WEAKLY DETERMINISTIC maps, and the REGULAR MAPS. Figure 1 below depicts the relationships between these classes as a nested hierarchy in which more complex classes properly include lesser ones. The left- and right-subsequential maps are provably less complex than the weakly deterministic maps (Heinz and Lai, 2013) and the regular maps (Mohri, 1997), and the weakly deterministic maps have been conjectured to be less complex than the regular maps (Heinz and Lai, 2013). Assuming this conjecture to be true (more on this below), this means that all weakly deterministic maps are also regular maps, but not all regular maps are weakly deterministic. Regular maps conjectured not to be weakly deterministic can be referred to as FULLY REGULAR.

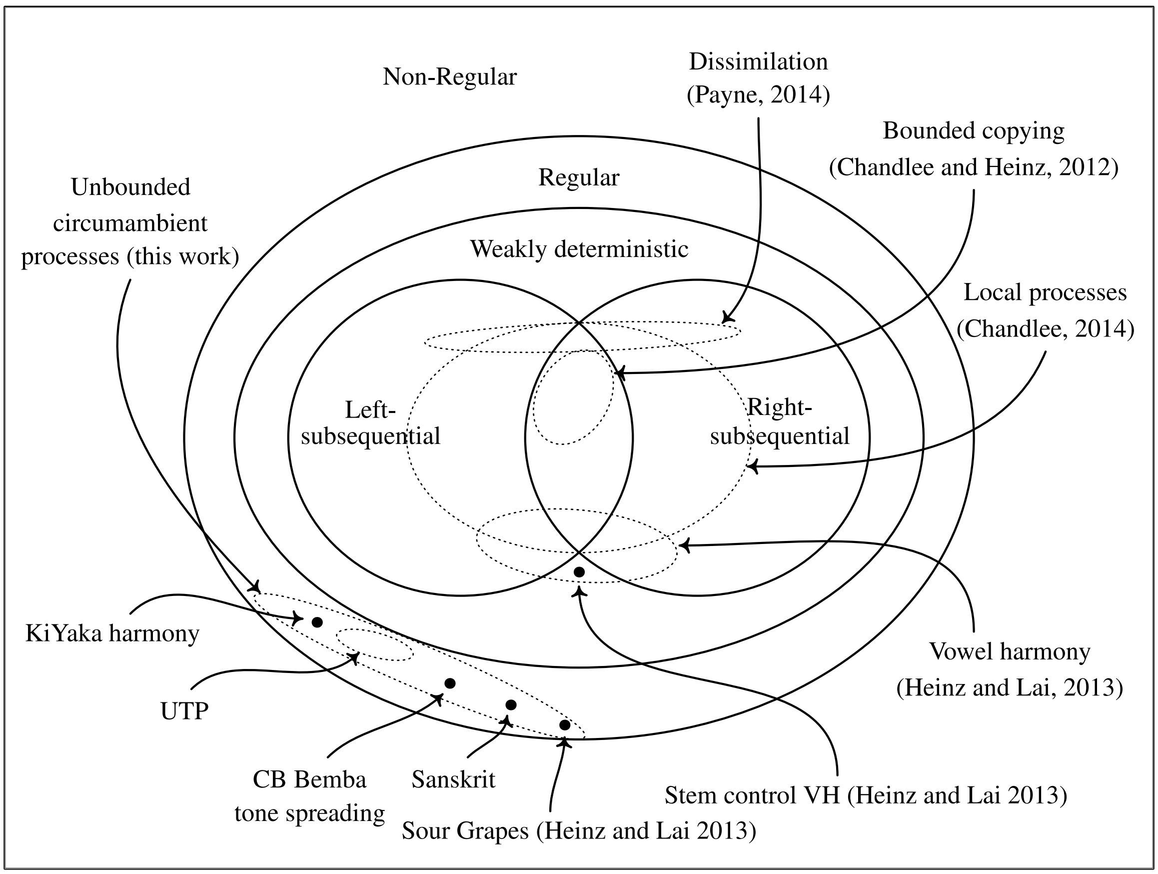

Figure 1: Phonology and the subregular classes

Figure 1 also depicts where typological research examining segmental processes as maps place these processes in the complexity hierarchy. This work has found all of them to be within the weakly deterministic class, with most in the less complex left- and right-subsequential classes. This has led to the hypothesis that that the weakly deterministic class forms a bound on the complexity of phonology. As Heinz and Lai (2013) discuss, this hypothesis is supported by the absence of Sour Grapes, which they prove to be neither left- nor right-subsequential and conjecture to not be weakly deterministic.

Figure 1 also shows the other unbounded circumambient processes discussed in this paper belonging to the fully regular region. As will be explained in detail in §4, this is because, like Sour Grapes, UTP and other unbounded circumambient processes are neither left- nor right-subsequential, and there is no known weakly deterministic characterisation of them. By revising the above hypothesis, then, we have a characterisation for the typological asymmetry established earlier: *segmental* phonology, but not *tonal* phonology, is at most weakly deterministic. The existence of Sanskrit *n*-retroflexion and KiYaka are admitted exceptions, but as already discussed in §2, they are not only rare but also exceptions to Wilson's myopia generalisation. These exceptions shall be discussed in §5.3.

3.3 Finite-state transducers, subsequentiality, and determinism

This subsection introduces FSTs informally and discusses SUBSEQUENTIAL TRANSDUCERS, a type of FST which is strictly less expressive than the fully regular NONDETERMINISTIC TRANSDUCERS (Mohri, 1997). To do this, let us use the regressive nasal place assimilation generalisation stated in (33a) below as an example. A rule for this generalisation is given in (33b).

(33) a. Nasals become labial before a labial consonant.
b. $[+\text{nasal}] \rightarrow [\text{labial}]$ / __ $[\text{labial}]$

The abstraction away from place features other than [labial] is to simplify the following notation. Let 'C' and 'V' represent non-labial, non-nasal consonants and vowels, respectively, and let 'p' represent labial consonants. An FST corresponding to the generalisation in (33a) is given in Figure 2.

A FST comprises a set of STATES (pictorially, the circles labeled with numbered qs) and TRANSITIONS (the labeled arrows) between the states. The FST can be interpreted as reading an input and writing an output as follows. Beginning from the START STATE (marked with the unlabelled arrow; in Figure 2, q_0), it traverses the transition for each symbol (indicated to the left of the colon in a transition label) in the input string, each time writing out the corresponding output (indicated to the right of each label).[12] At the end of the input, it appends to the output the output for the current state (indicated to the right of the colon on the state label). The EMPTY STRING λ indicates that there is

[12]For a more detailed introduction to FSTs the reader is referred to Beesley and Karttunen (2003).

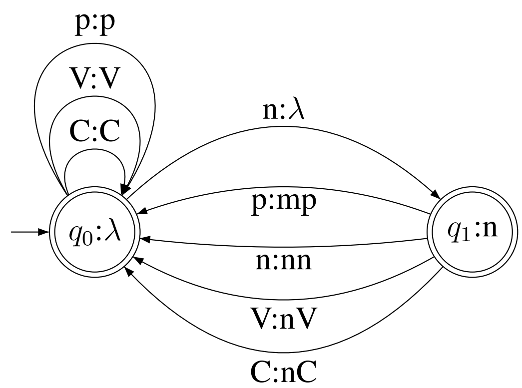

Figure 2: A deterministic nasal place assimilation FST

no output. In this way, the transitions and states define what input/output string pairs are accepted by a machine. As an example, a derivation is given below in (34) showing how an input ‘CVnC’ corresponds to an output ‘CVnC’ (i.e., no change occurs). The derivation highlights each state, input symbol, and output as the machine reads ‘CVnC’:

Input:		C		V		n		C	
State:	q_0	$\rightarrow$	q_0	$\rightarrow$	q_0	$\rightarrow$	q_1	$\rightarrow$	q_0
Output:		C		V				nC	

(34) A derivation for CVnC $\mapsto$ CVnC in the FST in Figure 2

Note that, while on the initial ‘C’ and ‘V’ the FST simply loops on state q_0, outputting ‘C’ and ‘V’, respectively, on the following ‘n’ it takes a transition to state q_1, outputting nothing. This is because the machine cannot yet ‘decide’ the output for this ‘n’, as it could be output as ‘n’ or ‘m’ depending on whether or not the following input symbol is a ‘p’. Let us thus call state q_1 a ‘wait’ state. As the next symbol is a ‘C’, the machine takes the ‘C:nC’ transition from q_1 to q_0, outputting ‘nC’. If the input is instead ‘CVnp’, the derivation is similar, but in the last step the machine takes the ‘p:mp’ transition from state q_1:.

Input:		C		V		n		p	
State:	q_0	$\rightarrow$	q_0	$\rightarrow$	q_0	$\rightarrow$	q_1	$\rightarrow$	q_0
Output:		C		V				mp	

(35) A derivation for CVnp $\mapsto$ CVmp in the FST in Figure 2

In this way, state q_1 allows the machine to ‘look ahead’ one symbol in the input. The reader can verify that this machine will take any permutation of ‘C’s, ‘V’s, ‘n’s and ‘p’s as an input, outputting 'n's as 'm's only before a 'p.' Thus, the FST in Figure (2) describes exactly the map represented in (33).[13]

For all states in Figure (2), there is only one possible transition given a particular input symbol. This means the FST in Figure (2) is DETERMINISTIC. Not all FSTs are deterministic; it is possible to write a FST such that a state may have two separate transitions on a particular output. Figure 3 is such a machine. In Figure 3, there are no outputs specified for the states, and state q_2 is a 'non-accepting' state (indicated by the single rather than double, circle), meaning that the machine rejects an input/output pair for which the input ends on that state.

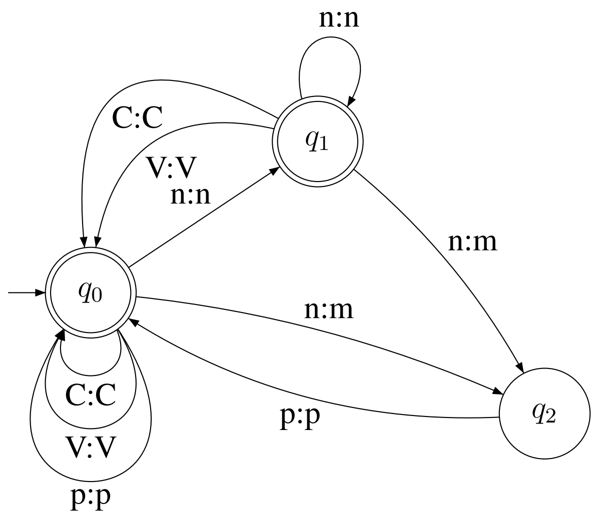

Figure 3: A non-deterministic nasal place assimilation FST

Non-deterministic FSTs are *more powerful* than deterministic FSTs: any map describable by a deterministic FST can be described by a nondeterministic FST, but there are maps describable by non-deterministic FSTs which *cannot* be described with a deterministic FST (Mohri, 1997). The ability to be described by a deterministic FST is central to the definition of the classes of left- and right-subsequential maps (Mohri, 1997). The regular class of maps are those maps which can be described by any *non*-deterministic FST. Thus, the left- and right-subsequential maps form a strict subset of the regular class.

UTP is a concrete example of a map which cannot be described with a deterministic FST, and it will be described in detail in §4. For now, we can understand the restriction determinism places on the subsequential maps in terms of 'wait' states, as in Figure 2, which allow deterministic FSTs to look ahead in the input. Because the number of states is finite, *there can only ever be a finite number of wait states*, and so a deterministic FST has *bounded* lookahead. Nondeterminism, in contrast, allows an FST to 'postpone' a decision about a particular input symbol indefinitely.

[13]This abstracts away from other constraints and phonological processes (for example, those preventing an output 'CCCC'). This is no different from modelling a specific phonological generalization with a single SPE-style rewrite rule or partial OT constraint ranking.

3.4 Subsequentiality and segmental phonology

Computational analyses of typologies of segmental processes have shown that they are describable with deterministic FSTs, with some variation regarding directionality. Mohri (1997) described two kinds of subsequential maps: RIGHT-SUBSEQUENTIAL maps and LEFT-SUBSEQUENTIAL maps, which are describable by deterministic FST reading an input string left-to-right and right-to-left, respectively. Collectively, they can be referred to as the SUBSEQUENTIAL maps.

The discussion from the previous section gives an intuition for how local processes are subsequential. For a thorough survey of the subsequentiality of local processes, including epenthesis, deletion, metathesis, substitution, and partical reduplication, the reader is referred to Chandlee (2014). Work studying long-distance segmental processes such as vowel harmony (Gainor et al., 2012; Heinz and Lai, 2013) and dissimilation (Payne, 2014) has also found them to be largely left- or right-subsequential. To see how, imagine a long-distance progressive consonantal harmony process in which a feature [–F] becomes [+F] after some other consonant specified [+F], no matter how early in the word this consonant appeared (an example is /l/ in KiYaka becoming [n] in suffixes attaching to a root containing a nasal). Such a map, given in (36) (with the changed feature highlighted in bold), is describable with the deterministic FST in Figure 4.

(36) ...[+F]...[–F]... $\mapsto$ ...[+F]...**[+F]**...

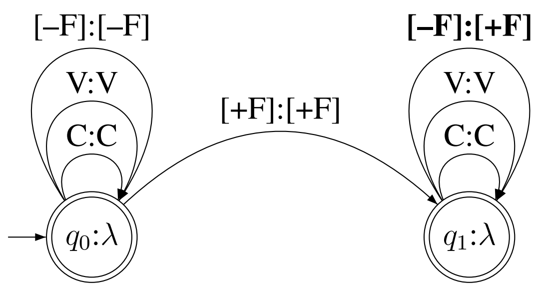

Figure 4: A deterministic FST for progressive harmony (C=consonant, V=vowel segments not participating in harmony)

This is impossible for a *regressive* harmony process, in which [–F] becomes [+F] *before* some [+F] segment an unspecified distance *later* in the word:

(37) ...[–F]...[+F]... $\mapsto$ ...**[+F]**...[+F]...

This requires unbounded lookahead in the left-to-right direction; in order to determine the output for a target [–F] in the input, a FST reading left-to-right would have to wait indefinitely to see if a trigger [+F] appears later in the string. However, this map is *right*-subsequential because reading the input *right-to-left* requires no lookahead. Reading right-to-left can be thought of as reversing the input, feeding it into the FST, and then reversing the output (Heinz and Lai, 2013). If we reverse the strings in the map in (37), we get exactly the same map as in (36). We can thus describe the reverse of (37) with a deterministic FST, and so it is right-subsequential.

Thus, local processes and unidirectional long-distance processes are either left- or right-subsequential. However, there are two classes of phonological processes which are neither left- nor right-subsequential: unbounded circumambient processes and bidirectional spreading processes. These processes are the focus of the following two sections.

4 Unbounded circumambient processes are not subsequential

This section proves that UTP is neither left- nor right-subsequential. This is compared to a similar result for Sour Grapes by Heinz and Lai (2013), and is then generalised to the class of unbounded circumambient processes. Additionally, §4.4 defends the use of a linear representation of UTP.

4.1 UTP as a map

To analyze UTP using the computational framework for studying string maps outlined in the previous section, the following uses a string representation marking associations to H tones on each TBU. The use of this representation shall be defended in §4.4. For now, the reader can understand the result as follows: *if* UTP is viewed with this kind of string representation, *then* it is neither left- nor right-subsequential.

The Unbounded Tone Plateauing (UTP) generalisation, originally formalised in (7) in §2, is repeated below in (38). Its string-based counterpart is given in (39). In (39), H represents a TBU associated to a H tone, and $\varnothing$ represents an unspecified TBU. The superscripts m, n, and p represent any natural number.

(38) a.
$$\begin{array}{ccccccc} \mu & \mu^n & \mu & \to & \mu & \mu^n & \mu \\ | & & | & & \diagdown & | & \diagup \\ \text{H} & & \text{H} & & & \text{H} & \end{array}$$

(39)
a. $\varnothing^n \mapsto \varnothing^n$
b. $\varnothing^m\text{H}\varnothing^n \mapsto \varnothing^m\text{H}\varnothing^n$
c. $\varnothing^m\text{H}\{\varnothing,\text{H}\}^n\text{H}\varnothing^p \mapsto \varnothing^m\text{HH}^n\text{H}\varnothing^p$

The linear map in (39) makes explicit every possible situation in UTP. In (39a) and (39b), in which there are fewer than two H-toned TBUs in the underlying form, no plateauing occurs. Plateauing occurs instead when there are two or more Hs in the input (see, e.g., (10b) from Luganda). This is seen in (39c): all TBUs in between the first and last H-toned TBUs surface as H. The notation $\{\varnothing,\mathrm{H}\}^n$ denotes a string of n TBUs, either H or ∅.

4.2 UTP is not subsequential

The UTP map in (39) is not left- nor right-subsequential. As shall be shown, this is because it requires unbounded lookahead in both directions. As §5.3 will discuss in detail, this holds for any unbounded circumambient process. This section presents an informal illustration of the proof; for the full proof, see the Appendix (§8).

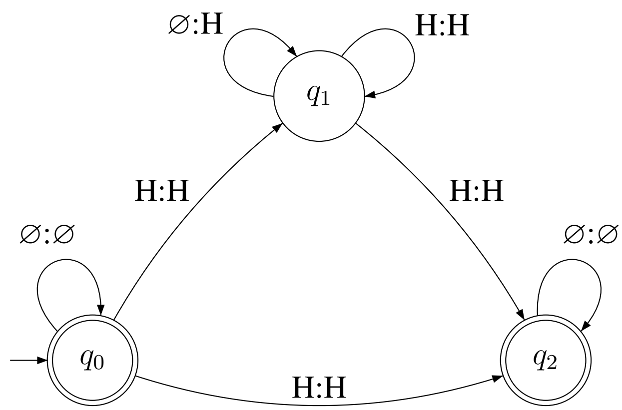

Figure 5: A non-deterministic FST for UTP

The UTP map in (39) is *regular*, as it can be modeled with a nondeterministic FST. This FST is given in in Figure 5. In this FST, underlying ∅ TBUs are output as H in state q_1, from which a final state can only be reached if there is another input H following; one can think of q_1 as the 'plateau' state.

This nondeterminism is *necessary*, as there is no way to capture this map with a deterministic FST. Recall that a deterministic FST must have at most one transition per input symbol at every state. We cannot determinise the FST in Figure 5. To see why not, let us attempt the 'waiting' strategy employed in the Figure 2 FST in §3.3. While there are many ways to try this, this discussion follows one. The proof in the Appendix ensures that all will fail, but the discussion here is intended to give an intuition as to why.

In the deterministic FST in Figure 6 below, q_2 is a wait state representing the knowledge that a sequence H∅—which may be a plateauing environment—has been seen in the input. Here, it 'waits' one symbol to see if the next symbol in the input is a H or a ∅ . It will thus correctly transform ∅H∅H, with only one intervening ∅ TBU, to ∅HHH. However it incorrectly maps inputs for which the second H is farther away; for example,

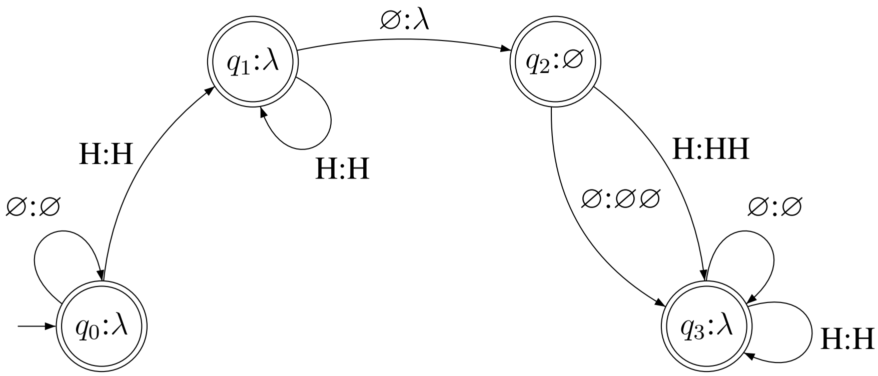

Figure 6: First attempt at a deterministic FST for UTP

the input ∅H∅∅H would be mapped to ∅H∅∅H (itself). Thus, because there is only one wait state, the machine in Figure 6 can only describe plateaus of at most three Hs. The following machine in Figure 7 thus adds an additional wait state to try and remedy this situation.

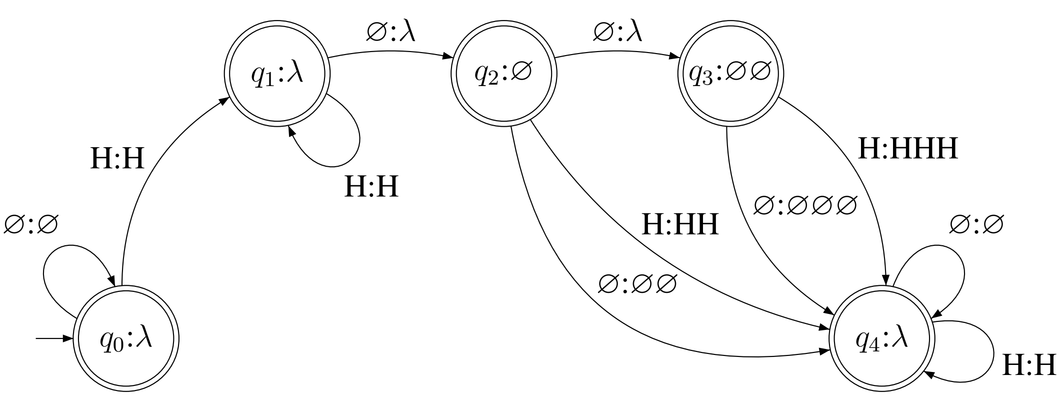

Figure 7: Second attempt at a deterministic FST for UTP

The FST in Figure 7 is much like the one in Figure 6, except that it has two wait states, q_2 and q_3, which represent the lookahead necessary to capture the behavior of one and two ∅s, respectively, following an input H. Thus, Figure 7 correctly maps ∅H∅H $\mapsto$ ∅HHH and ∅H∅∅H $\mapsto$ ∅HHHH. However, it incorrectly maps inputs like ∅H∅∅∅H, where two Hs are separated by three or more ∅s, to themselves. By now it is perhaps obvious that we are on a wild goose chase; any 'wait n symbol' strategy will fail for any mapping in the UTP relation whose input string includes the sequence H$\varnothing^{n+1}$H. However, given the restriction of determinism, 'wait n symbols' is the best we can do. Simply reversing the string, in an attempt to create a right-subsequential transducer, will not help us; the position of the first triggering H is just as arbitrarily far to the right as the second is to the left. Thus, a deterministic FST representation of the UTP map is impossible, and so it lies outside both the left- and right-subsequential classes.

4.3 Interim conclusion: Unbounded circumambient processes are not subsequential

That UTP is neither left- nor right-subsequential follows from its unbounded circumambient nature: as triggers may lie any distance away on either side of a given target, a FST describing the map requires unbounded lookahead in both directions.

Heinz and Lai (2013) prove this is also true for Sour Grapes, for the same reasons. For Sour Grapes, the fate of an input segment that can potentially assimilate rests on whether or not a trigger appears to one side *and* whether or not a blocker appears to the other. They show that this means it cannot be described by a deterministic FST reading in either direction. We can then see why no unbounded circumambient process can be subsequential. For unbounded circumambient processes, it is *definitional* that crucial information may appear on either side of the target, unboundedly far away. This means any unbounded circumambient process will require unbounded lookahead in both directions, and will not be subsequential. This result is then key to characterising the unbounded circumambient asymmetry in terms of computational complexity, as discussed further in §5.

As stated at the outset, this result depends on a particular kind of string-based representation. Thus, before moving on to compare unbounded circumambient processes with bidirectional spreading, as will be taken up in §5, it is necessary to address the potential objection that these results are invalid because they hinge on this string representation for UTP.

4.4 Autosegmental representations and linear representations

Tone is widely analyzed with autosegmental representations (though for recent representational alternatives see Cassimjee and Kisseberth (2001) and Shih and Inkelas (2014)), while the notion of subsequentiality is defined in terms of strings. Thus, the preceding analysis used a string representation of UTP, although it is a tonal process. This section explains and defends two interrelated facts about such string representations:

(40) a. Each symbol in a string represents associations to a timing tier unit in any corresponding autosegmental representation.
b. Therefore, 'unbounded lookahead' is determined by what translates to autosegmental terms as distance on the *timing* tier, not the *melody* tier.

As discussed below, (40a) is a common assumption about the relationship between string representations and autosegmental representations, and is the one implicit in the string representations used in the computational literature cited in §3.4. As for (40b), measuring distance as on the timing tier as opposed to the melody tier is a representational *choice*. Indeed, Kornai (1995) compares different ways of encoding string representations of autosegmental representations. The following sections explain and justify string representations based on (40) in detail.

4.4.1 Translating between string and autosegmental representations

Let us compare string and autosegmental representations of the Sour Grapes map, repeated below in (41) from (29).

(41) a. $[–\text{F}]^n \mapsto [–\text{F}]^n$ (no trigger, no blocker $\rightarrow$ no harmony)

b. $...[+\text{F}][–\text{F}]^n \mapsto ...[+\text{F}]\mathbf{[+F]}^n$ (trigger, no blocker $\rightarrow$ **harmony**)

c. $...[+\text{F}][–\text{F}]^n[!\text{F}]... \mapsto ...[+\text{F}][–\text{F}]^n[!\text{F}]...$ (trigger, blocker $\rightarrow$ no harmony)

We have a number of options for representing the featural contrasts made by the symbols in (41) autosegmentally. Some possible autosegmental representations for the underlying form $[+\text{F}][–\text{F}]^n[!\text{F}]$ from (41c) are given below in (42).

(42) a.
$$\begin{array}{cccccc} \text{V} & \text{V}_1 & \cdots & \text{V}_n & \text{V} \\ | & & \diagdown | \diagup & & | \\ +\text{F} & & –\text{F} & & !\text{F} \end{array}$$

b.
$$\begin{array}{cccccc} \text{V} & \text{V}_1 & \cdots & \text{V}_n & \text{V} \\ | & | & & | & | \\ +\text{F} & –\text{F} & \cdots & –\text{F} & !\text{F} \end{array}$$

c.
$$\begin{array}{cccccc} \text{V} & \text{V}_1 & \cdots & \text{V}_n & \text{V} \\ | & & & & | \\ +\text{F} & & & & –\text{F} \end{array}$$

In (42a) and (b), the target [–F] vowels in the string are represented as underlyingly associated to [–F] features in the autosegmental diagram. In (42a), a single [–F] feature is associated to multiple vowels, whereas in (42b), each [–F] vowel is associated to its own [–F] feature (in violation of the Obligatory Contour Principle (Leben, 1973; McCarthy, 1986)). Here, a [!F] autosegment is used as shorthand for some vowel which is also associated to some other feature [+G] which prevents [+F] from spreading to it (e.g., as in [+low] vowels in Akan (Clements, 1976), which block the spreading of a [+ATR] feature). In (42c), the targets are analyzed as underspecified on the [$\pm$F] tier, and the blocker is analyzed as underlyingly specified as [–F] (as in Clements (1976)'s analysis of [+low] vowels in Akan as underlyingly [–ATR]).

One property that all possible autosegmental analyses share is that each symbol in a string in (41a) corresponds to the featural associations of a particular timing tier unit in (42). To give an explicit example of this, the following is a translation between symbols in UR strings in (41) and the autosegmental information at each timing tier unit in (42a). The translation for each string symbol [+F], [–F], and [+F] is given in (43a), and an example correspondence between the string [+F][–F][–F][!F] and an autosegmental representation is given in (43b).

(43) a. $\begin{array}{c} \mathrm{V} \\ | \\ +\mathrm{F} \end{array} \Leftrightarrow [+\mathrm{F}] \quad \begin{array}{c} \mathrm{V} \\ | \\ -\mathrm{F} \end{array} \Leftrightarrow [-\mathrm{F}] \quad \begin{array}{c} \mathrm{V} \\ | \\ !\mathrm{F} \end{array} \Leftrightarrow [!\mathrm{F}]$ b. $\begin{array}{cccc} [+\mathrm{F}] & [-\mathrm{F}] & [-\mathrm{F}] & [!\mathrm{F}] \\ \Updownarrow & \Updownarrow & \Updownarrow & \Updownarrow \\ \mathrm{V} & \mathrm{V} & \mathrm{V} & \mathrm{V} \\ | & \diagdown & \diagup & | \\ +\mathrm{F} & \multicolumn{2}{c}{-\mathrm{F}} & !\mathrm{F} \end{array}$

While simplified in the sense that it focuses on one feature, translations like in (43) are what phonologists commonly, if implicitly, use when moving back and forth between linear strings of feature bundles and autosegmental representations. Again, what is important is that each symbol in the string corresponds to a timing tier unit in the autosegmental representation—the property originally highlighted in (40a).

As (43) makes clear, information about the *melody* (or featural) tier units in the autosegmental representation is obscured in the string representation. In (43b), the string does not encode the multiple associations of the [–F] autosegment, and it also does not show that [+F] and [!F] are only separated by a single [–F] autosegment. This ambiguity is apparent in (42), in which different autosegmental interpretations of the same string representation have different information on their featural tiers.

4.4.2 Representational assumptions and lookahead

As a result, when a FST reads in a symbol in a string representation like in (41), it can thus be thought of in autosegmental terms as reading the featural information of a particular timing tier unit. Thus, because subsequentiality depends on lookahead in terms of a FST reading such a string representation, the results regarding subsequentiality in the work of Heinz and Lai (2013) and others are based on what translates in autosegmental terms to lookahead on the *timing* tier, and *not* the melody tier.

There are a few arguments for this assumption. One is that locality on the melody tier depends on certain representational assumptions such as underspecification and the Obligatory Contour Principle, both of which have been argued against as phonological universals (see Inkelas (1995) for the former, Odden (1986) for the latter). The string representation is agnostic about such assumptions.

More importantly, as detailed in this paper, it is by measuring unboundedness on the timing tier which distinguishes UTP and Sour Grapes from common segmental processes via its non-subsequentiality. Thus, while this choice is an assumption on which the results in this and the literature cited in §3.4 is based, it appears to be correct in that it helps us to characterise the unbounded circumambient asymmetry. How this relates to different representational assumptions will be discussed momentarily.

Returning to UTP, it is then important to establish that the linearisation used in the previous section is indeed comparable to that of previous studies of subsequentiality, such as for Sour Grapes in (41), in that it also measures lookahead in terms of the timing tier. The linearisation in (39) can be made explicit with a map along the lines of the one in (43) as given below in (44a).

(44) a. $\begin{array}{c} \mu \\ | \\ \mathrm{H} \end{array} \Leftrightarrow \mathrm{H} \qquad \mu \Leftrightarrow \varnothing$ b. $\begin{array}{ccc} \mu & \mu & \mu \\ & | & \\ & \mathrm{H} & \end{array} \Leftrightarrow \varnothing\mathrm{H}\varnothing$

As can be seen in (44b), the string symbols H and ∅ can encode both contrasts in association in the underlying representation and the changes in these associations in the surface, as originally claimed in (40a). Additionally, the string representation in (44b) preserves the linear order of the TBU tier of the autosegmental representation, just as in (43). Thus, lookahead is measured in terms of the timing tier, as originally claimed in (40b). Note again that, as they simply encode the associations to each timing tier unit, linearisations like in (43) and (44) are very general, and can be applied to any set of autosegmental representations for a particular process.

4.4.3 Representation and subsequentiality

This concludes the arguments for why this particular kind of representation is used in this paper. However, an important question remains: what if we *don't* use this particular representation? The answer is that this is a valid direction for research, but any results along these lines will not change the result argued for in this paper. It is known that, formally, representation is related to expressive power (Medvedev, 1964; Rogers et al., 2013). Work studying the computational properties of autosegmental representations (Kay, 1987; Wiebe, 1992; Bird and Ellison, 1994; Yli-Jyrä, 2013; Jardine and Heinz, 2015), while significant, does not yet offer a hierarchy of complexity like the one presented in §3.4. Future research may use this work in computational autosegmental phonology as a starting point to study the relationship between representation and computational complexity of phonological processes.

It may even be possible to derive non-subsequentiality of unbounded circumambient processes from some aspect of representation (such as lookahead on the melody tier). This requires, however, that if such representations were translated into the string representations according to (40), non-subsequentiality in the string representations would somehow emerge. Thus, any such explanation based on representation would *duplicate* the computational results outlined here—it would not *refute* them. That is, it would have to uphold the fact that, when viewed as string maps for for which (40) is true, much of segmental phonology is subsequential while UTP and Sour Grapes are not.

5 Unbounded circumambient processes and weak determinism

Having shown that unbounded circumambient processes are not left- or right-subsequential, there is one final distinction to be made. This section argues that unbounded circumambient processes are computationally distinct from bidirectional spreading processes, introduced in §2.5. This leads to the characterisation of the unbounded circumambient asymmetry in terms of a weakly deterministic complexity bound on segmental phonology which is absent in tone.

5.1 Bidirectional spreading and the weakly deterministic class

Unbounded circumambient processes are not the only class of process which is not left- or right-subsequential. In patterns of stem-control harmony, or cases of bidirectional spreading, as in the Arabic emphasis spreading discussed in §2.5, a feature spreads outward both to the right and the left:

(45) ...[–F]...[+F]...[–F]... $\mapsto$ ...**[+F]**...[+F]...**[+F]**...

Such a map requires unbounded lookahead in either direction: a target may follow or precede the trigger, in any direction. Thus, as Heinz and Lai (2013) also show, such cases (which shall henceforth be referred to under the umbrella term BIDIRECTIONAL SPREADING) are also not subsequential. However, there is a crucial difference between bidirectional spreading and unbounded circumambient processes. Bidirectional spreading hinges on a single trigger whose influence spreads outward. In contrast, unbounded circumambient processes hinge on *two* triggers/blockers whose targets lie between.

Heinz and Lai (2013) observe that bidirectional spreading processes are essentially the same unidirectional map applied left-to-right and then right-to-left. They propose a superclass of the subsequential maps, called the WEAKLY DETERMINISTIC maps, which includes this kind of process. Weakly deterministic maps are those which can be *decomposed* into a left- and right-subsequential map such that the left-subsequential map is not allowed to change the alphabet or increase the length of the string.

The process in (45) can be decomposed into two left- and right-subsequential maps describable by the consonant harmony FST in Figure 4 in the following way. First the input string is read by the FST left-to-right (applying the left-subsequential map), then the resulting output is fed back into the FST right-to-left (applying the right-subsequential map to the output). This decomposition is schematised below in (46):

(46)	a. (left-subsequential)	...[–F]...[+F]...[–F]...	$\mapsto$	...[–F]...[+F]...**[+F]**...
	b. (right-subsequential)	...[–F]...[+F]...[+F]...	$\mapsto$	...**[+F]**...[+F]...[+F]...

This composition of the two maps is special because it does not change the alphabet or increase the length of the string. Intuitively, this is because these bidirectional processes can be thought of as one unidirectional process operating in two directions. This is highlighted by the fact that both sub-maps use the *same* FST. Thus, bidirectional spreading is weakly deterministic, as first seen in Figure 1.

The weakly deterministic class is defined as such by Heinz and Lai (2013) as a restriction on Elgot and Mezei (1965)'s result that any regular map can be decomposed into a left-subsequential and right-subsequential map, as long as the left-subsequential map is allowed to enlarge the alphabet. This result holds because the left-subsequential map can 'mark-up' the string with extra symbols in the first map, and then erase them in the second.[14] However, no attested segmental maps studied in the literature cited above require such a markup.

5.2 Unbounded circumambient processes and the weakly deterministic class

In contrast, Heinz and Lai (2013) argue that Sour Grapes requires such a markup, and thus is not weakly deterministic.[15] This is because Sour Grapes is not simply the application of the same process in two directions. Because Sour Grapes is unbounded circumambient, for any decomposition into two sub-maps, the left-subsequential process must somehow encode whether or not a [+F] has been seen to the right of the [–F] targets. It is difficult to see how this can be done without intermediate markup, and thus Heinz and Lai (2013) conjecture that Sour Grapes is not weakly deterministic. The exact same arguments apply to UTP. Let us look at why more concretely.

UTP can be decomposed into two subsequential maps with an *augumented* alphabet, as in (47). First, the left-subsequential process marks all ∅ following a H as ?, and then the right-subsequential process changes all ? preceding a H to H (and all ? not preceding a H to ∅).

(47)

Input	∅∅∅H	H∅∅∅	H∅∅H
	↧	↧	↧
a. (left-subsequential)	∅∅∅H	H???	H??H
	↧	↧	↧
b. (right-subsequential)	∅∅∅H	H∅∅∅	HHHH
Output	∅∅∅H	H∅∅∅	HHHH

Crucially, this decomposition relies on the intermediate ? symbols to 'carry forward' the information that a H appears to the left in the string. This allows the right-subsequential map to correctly apply without any unbounded lookahead. However, like Sour Grapes, it is hard to see how there could be a similar decomposition which uses only H and ∅, maintains string length in the left-subsequential sub-map, and obtains the exact same map. This is because using the same alphabet to create an encoding will always lead to distinctions between input strings being lost, and so such an encoding is bound to fail.

[14] This is not unlike the use of abstract intermediate forms in early derivational phonology, e.g. (Clements, 1977).

[15] *Proving* that a pattern is not in a complexity class requires an abstract characterisation of that class. For example, abstract characterizations of the regular and context-free languages allow for 'pumping lemmas' which can be used to prove that a language is not a member of these classes (Hopcroft et al., 2006). No such tools exist yet for the weakly deterministic class. A similar case exists for P, or the class of problems that can be solved in polynomial time. Many problems in the NP class conjectured to be outside of P are considered computationally intractable (see Idsardi, 2006; Heinz et al., 2009, for such discussion on OT), although no proof exists that P≠NP (Fortnow, 2009).

Again, this is based on the unbounded circumambient nature of the process: because there is crucial information on either side of the targets, it is necessary to mark targets to the right of the left trigger in order for the right-subsequential function to correctly process them without unbounded lookahead. Thus, Heinz and Lai (2013)'s conjecture also applies not just to UTP, but to any unbounded circumambient process.

An important implication of Heinz and Lai (2013)'s conjecture is that there exist FULLY REGULAR maps outside of the weakly deterministic class, and thus that the weakly deterministic class of maps is, as depicted in Figure 1, a proper subclass of the regular class of maps. This section has argued that unbounded circumambient processes fall into this fully regular class (also depicted in Figure 1), and that this distinguishes them in complexity from bidirectional spreading.

5.3 The weakly deterministic hypothesis

Thus, this conjecture provides a way to capture the unbounded circumambient asymmetry. Heinz and Lai (2013) posit a WEAKLY DETERMINISTIC HYPOTHESIS for phonology: phonology is at most weakly deterministic. This hypothesis correctly predicts the absence of unbounded circumambient processes like Sour Grapes in the typology of vowel harmony. However, this paper has shown that unbounded circumambient processes are well-attested in tone. This paper thus proposes that the weakly deterministic bound *only* applies to segmental phonology, and not *tonal* phonology. This accurately predicts fully regular maps such as UTP and Copperbelt Bemba H spread to exist in tone. Furthermore, to propose that tone is more computationally complex than segmental phonology, as has been done here, is in line with Hyman (2011) and others' assertions that tone can 'do more' than phonology.

An explanation for how the weakly deterministic bound manifests in the phonological system shall be left for future work, although it is possible to speculate on a few points. One, formal language complexity correlates with an increase in computational resources necessary for parsing and generation. It could be that tone has access to such resources because prosodic information more commonly interacts with syntax (see, e.g., Hyman and Katamba, 2010) and thus requires more powerful computation. This issue of computational power can also be directly related to learning, as empirical work suggests that formal complexity constrains phonological learning (Heinz, 2010; Lai, 2012, 2015; McMullin and Hansson, 2015; Moreton and Pater, 2012; Rogers et al., 2013).

5.4 Explaining the exceptions to the weakly deterministic hypothesis

The weak determinism hypothesis in its strongest form cannot account for the rare cases of unbounded circumambient segmental processes discussed in this paper, Sanskrit *n*-retroflexion and KiYaka vowel harmony. There are a few ways to reconcile these exceptions with the weakly deterministic hypothesis.

One possibility is that the accounts of these processes in the literature were incorrect in classifying them as unbounded. As brought up in §2.3.1, Ryan (2015) has observed that blocking of Sanskrit *n*-retroflexion may be bound to the adjacent syllable. In KiYaka, the greatest attested distance was only three vowels, and was restricted to a particular morphological context.

Another explanation comes from potentially interfering factors. For example, the presence of an unbounded circumambient process in the tonal phonology may license an unbounded circumambient process in segmental phonology. This may be the case for KiYaka, which as noted in §2.2.3, also has UTP. It could be that KiYaka speakers, having first internalized plateauing process in tonal phonology, may then be able to generalise it to their segmental phonology.

Finally, it may simply be that the constraint against fully regular maps in segmental phonology is not categorical but somehow gradient, and thus admits exceptions. One way this may manifest is through a learning bias in which individuals are more receptive to some patterns than others (Moreton, 2008; Wilson, 2006a). Children may show a strong preference for weakly deterministic maps when learning segmental phonology, but may change to a more general learner in the face of sufficient data. As Staubs (2014) shows, gradient typological generalisations may also result from the transmission of patterns among multiple learning agents. It is likely that non-weakly deterministic patterns require more kinds of data to learn, and perhaps this data is only available in tone, which is known to operate over much larger domains that segmental phonology.

Hence, there are a number of reasons why the cases of Sanskrit *n*-retroflexion and KiYaka vowel harmony do not immediately invalidate the proposal offered here. Even if they cannot ultimately be explained away, they are exceptional for other characterisations of segmental phonology, such Wilson (2003, 2006b)'s myopia generalisation. Finally, Sanskrit *n*-retroflexion and KiYaka vowel harmony appear to be the only such cases—Hyman (2011) states, for example, that KiYaka is the "only one example" of such a process of which he is aware (p. 218). As such, the development of the potential explanations above shall be left to future research.

In sum, this paper has characterised the unbounded circumambient asymmetry in terms of computational complexity and identified the weakly deterministic boundary as the relevant difference between tone and segmental phonology which captures this asymmetry.

6 The Unbounded Circumambient Asymmetry in Optimality Theory

The previous sections have laid out a computational explanation for the unbounded circumambient asymmetry. This section briefly outlines why this explanation is superior to any possible explanation using current theories based in Optimality Theory (Prince and

Smolensky, 1993) or its variants. This is because current theories of OT do not provide a unified characterisation of the unbounded circumambient processes. As shall be argued, this means that making OT empirically adequate with regards to the asymmetry runs into a ‘duplication of effort’ problem. Banning the particular non-local effects that generate segmental unbounded circumambient processes requires changes both to how OT manipulates segments autosegmentally but also to how OT compares candidates. Thus, OT does not provide a unified characterisation of the asymmetry comparable to the one based on computational complexity put forth in the preceding sections.

Sour Grapes vowel harmony is an appropriate place to begin. Sour Grapes is unattested in segmental phonology but predicted by parallel OT and local AGREE constraints checking the agreement in a feature of adjacent segments. Given such constraints, candidates in which spreading is blocked still violate AGREE (as the last segment to which the feature has spread still disagrees with the blocker) and are thus harmonically bounded by candidates in which no spreading occurs, as they also violate AGREE but do not register any FAITHFULNESS violations incurred by the spreading feature (see McCarthy (2010) for a thorough explanation of how Sour Grapes is generated in OT). Thus, the nature of optimisation, which allows comparison of candidates with non-local changes to those with local ones, allows for unbounded behaviour with local constraints. In fact, this is exactly this type of behaviour that Hyman (1998) harnesses to describe vowel harmony in KiYaka.

One proposal for dealing with the Sour Grapes problem is McCarthy (2010)’s analysis of spreading in Harmonic Serialism (HS), which restricts GEN to only produce candidates with a single change (with winning candidates fed back into the grammar until no changes are more optimal). As only candidates with single changes are compared, global comparison cannot take place, and thus patterns such as Sour Grapes can no longer be generated with local AGREE constraints. Of course, assuming a HS framework provides no explanation for the presence of a Sour Grapes-like tone pattern in Copperbelt Bemba.

Furthermore, this solution is specific to Sour Grapes vowel harmony resulting from AGREE, and does not extend to autosegmental analyses of plateauing. Plateauing, tonal or otherwise, can be seen as two identical features adjacent on the tier merging to satisfy the Obligatory Contour Principle (OCP), the intervening anchor units associating to the resulting fused feature in order to avoid a gapped structure. This can be schematized derivationally as follows:

(48)

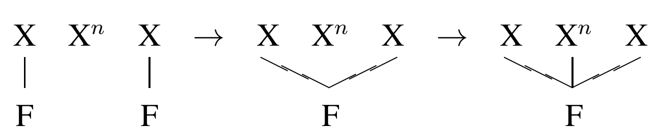

While (48) is a derivational sketch of the process, the intuition can be implemented in either serial or parallel versions of OT given basic constraints governing the behaviour of autosegments proposed in the OT literature. Fusion of identical, adjacent autosegments can be motivated by an OCP markedness constraint (Leben, 1973; McCarthy, 1986; Yip,

2002) ranking over a UNIFORMITY faithfulness constraint militating against the fusion of autosegments (Pater, 2004; Meyers, 1997). The ‘filling in’ of the associations to the intervening anchor units can then be motivated by ranking NOGAP (Ito et al., 1995; Ringen and Vago, 1998; Walker, 1998; Yip, 2002), a constraint against structures in which associations skip over potential anchors, above faithfulness constraints militating against the addition of association lines. In classical OT, given such a ranking only a candidate featuring both fusion and ‘filling’ in will satisfy both the OCP and NOGAP markedness constraints. In HS, a candidate featuring fusion will represent an increase in harmony (as it satisfies the OCP), and then successive candidates representing a gradual filling in of the associations in the intervening anchor units will represent an increase in harmony with respect to NOGAP. Thus, both the fusion and ‘filling in’ can be achieved straightforwardly either in parallel OT or in HS.

Thus, we have a duplication of effort problem: in order to categorically remove segmental unbounded circumambient processes from the typology predicted by OT theories of phonology, it is necessary to at least adopt one of the changes proposed for Sour Grapes *and* remove the segmental versions of the above constraints that achieve autosegmental plateauing. Thus, while it appears technically possible to have a theory of OT that is empirically adequate with regards to the unbounded circumambient asymmetry, it does not provide a unified explanation. In contrast, the FLT-based account presented here characterises the difference: segmental phonology is weakly deterministic, but tonal phonology is not.

7 Conclusion

This paper has made three contributions. One, it has documented the asymmetry in the attestation of unbounded circumambient processes in tonal phonology and segmental phonology. Two, it has shown that UTP is similar to Sour Grapes, in that they are both unbounded circumambient processes, and that this similarity has formal consequences. The third contribution then characterised the asymmetry between segmental phonology and tone by arguing that unbounded circumambient processes are fully regular and are thus more computationally complex than processes which do not require unbounded lookahead in two directions. This has been shown to be a superior characterisation than that offered in Optimality Theory, which cannot account for the asymmetry in a unified way.

The conclusions in this paper raise a number of interesting questions for future research. For one, what computational constraints are there on tone? In other words, is there a sub-regular class of maps which includes (or corresponds to) unbounded circumambient processes? How does this relate to Elgot and Mezei (1965)’s result using intermediate markup to generate any regular map from two subsequential maps? As Heinz and Lai (2013) point out, this idea may be brought to bear on the question of how abstract intermediate representations are in phonology. The questions of representation raised at the end of §4.4 can be approached in a similar way—how do changes in representation correlate with changes in generative capacity? Finally, it remains to be seen how the FLT insight presented here can be incorporated into traditional phonological theory. One potential approach is to appeal to learning, as raised in §5.3.

Unfortunately, going into these concerns in detail is beyond the scope of this paper. Instead, this paper's goal is similar to that of Kisseberth (1970)'s prophetic work on conspiracies in Yawelmani. Kisseberth writes that he is not "principally interested in proposing detailed formalism; instead I would like to encourage phonologists to look at the phonological component of a grammar in a particular way" (p.293 Kisseberth, 1970). This work, too, aims to encourage phonologists to look at phonology in a new way. Regardless of how it is incorporated into our previous understanding of phonology, the unbounded circumambient asymmetry between tonal and segmental phonology is robust, and the best available characterisation of this generalisation comes from computational complexity.

8 Appendix: Mathematical Definitions and Proof

8.1 Notation

Basic knowledge of set theory is assumed. An alphabet is a finite set of symbols; if Σ is an alphabet, let Σ^* denote the set of all finite strings, including the empty string λ, over Σ. Let $|w|$ denote the length of string w. If w and u are strings let wu denote their concatenation. If w is a string and X is a set of strings then let wX denote the set of strings resulting from concatenating w to each string in X. The PREFIXES of a string $w \in \Sigma^*$ are $Pr(w) = \{u \in \Sigma^* | \exists v \in \Sigma^* \text{ such that } w = uv\}$. The prefixes of a set of strings $L \subseteq \Sigma^*$ are $Pr_{set}(L) = w \in \cap_{x \in L} Pr(x)$.

The LONGEST COMMON PREFIX (lcp) of a set of strings is the longest prefix shared by all strings in the set: $lcp(L) = w$ such that $w \in Pr_{set}(L)$ and $\forall w' \in Pr_{set}(L), |w'| \leq |w|$. For example, $lcp(\{aaa, aab\}) = aa$, because aa is the longest prefix shared by both aaa and aab.

If Σ and Δ are alphabets a RELATION is some subset of $\Sigma^* \times \Delta^*$. A relation R is a MAP (or FUNCTION) iff for all $w \in \Sigma^*$, $(w, v), (w, v') \in R$ implies $v = v'$.

The TAILS of x in given a relation R, denoted $T_R(x)$ are $T_R(x) = \{(y, v) | t(xy) = uv, u = lcp(t(x\Sigma^*))\}$. If R is a map, it is a SUBSEQUENTIAL map iff its sets of tails are finite; that is, the set $\bigcup_{w \in \Sigma^*} \{T_R(w)\}$ is of finite cardinality.

8.2 Subsequential Finite State Transducers

A finite-state transducer (FST) is a six-tuple $(q_i, F, Q, \Sigma, \Delta, \delta)$ where Q is the finite set of states, $q_i \in Q$ is the initial state, $F \subseteq Q$ is the set of final states, and $\delta \subseteq Q \times \Sigma^* \times \Delta^* \times Q$ is the transition function. The recursive extension of the transition function δ^* is defined as:

- $\delta \subseteq \delta^*$
- $(q, \lambda, \lambda, q) \in \delta^*$ for all $q \in Q$
- $(q, x, y, r) \in \delta^*$ and $(r, a, b, s) \in \delta$ implies $(q, xa, yb, s) \in \delta^*$

The relation that a FST describes is defined as $R(t) = \big\{(x, y) \in \Sigma^* \times \Delta^* \big| \exists q_f \in F$ such that $(q_i, x, y, q_f) \in \delta^*\big\}$.

A FST is DETERMINSITIC iff $\forall q \in Q$ and for all $\sigma \in \Sigma$, $(q, \sigma, v, r), (q, \sigma, v', r') \in \delta$ implies $v = v'$ and $r = r'$. SUBSEQUENTIAL FSTs (SFSTs) are deterministic FSTs with an added output function $\omega : Q \rightarrow \Delta^*$ which specifies for each state an output string to be written when the machine ends on that state. Thus, a SFST is a 7-tuple $(q_i, F, Q, \Sigma, \Delta, \delta, \omega)$. The relation that a SFST describes is defined as $R(t) = \{(x, yz) \in \Sigma^* \times \Delta^* : \exists q_f \in F$ such that $(q_i, x, y, q_f) \in \delta^*$ and $\omega(q_f) = z\}$.

Theorem 1 (Mohri 1997) *A relation R is a subsequential map iff it is recognized by a SFST for which each state in the machine corresponds to a set of tails in R.*

8.3 Proof UTP is not subsequential

The proof is exactly Heinz and Lai (2013)'s proof for the non-subsequentiality of Sour Grapes.

Proof. Let UTP be the map discussed in the main text. The following shows that for all distinct $n, m \in \mathbb{N}$, $T_{UTP}(\text{H}\varnothing^m) \neq T_{UTP}(\text{H}\varnothing^n)$. As $\mathbb{N}$ is infinite, this means there must be infinitely many states in the canonical SFST for UTP, which by Theorem 1 means there is no SFST describing it.

If $x =$H∅, $lcp(UTP(x\Sigma^*)) =$H, because $UTP(x\Sigma^*)$ includes both HLL (which $= UTP$(H∅∅)) and HHH (which $= UTP$(H∅H)) and thus there no shared prefix of $UTP(x\Sigma^*)$ longer than H. Thus for all $n \neq 2$, (∅,L^n) $\notin T_{UTP}$(H∅); i.e., (∅,LL) is the only possible tail with ∅ as the first member of the tuple.

If $x =$H∅∅, $lcp(UTP(x\Sigma^*)) =$H, because $UTP(x\Sigma^*)$ includes both HLLL and HHHH. Thus for all $n \neq 3$, (∅,L^n) $\notin T_{UTP}$(H∅); i.e., only (∅,LLL) is possible is the only possible tail with ∅ as the first member of the tuple.

We can see then that for any distinct $n \in \mathbb{N}$, (∅,L^k) $\in T_{UTP}$(H∅^{n}) only if $k = n+1$. For $m \in \mathbb{N}, m \neq n$, (∅,L^j) $\in T_{UTP}$(H∅^{m}) only if $j = m+1$. Thus for all distinct n and m, $k \neq j$, and so T_{UTP}(H∅^{n}) $\neq T_{UTP}$(H∅^{m}). □

References

Al Khatib, Sam (2008). On the directionality of emphasis spread. In *Proceedings of the 2008 annual conference of the Canadian Linguistic Association*.

Baković, Eric (2000). *Harmony, dominance and control*. Ph.D. thesis, Rutgers University.

Beesley, Kenneth R. and Karttunen, Lauri (2003). *Finite State Morphology*. CSLI Publications.

Bennett, William (2013). *Dissimilation, Consonant Harmony, and Surface Correspondence*. Ph.D. thesis, Rutgers, the State University of New Jersey.

Beros, Achilles and de la Higuera, Colin (2014). A canonical semi-deterministic transducer. In *Proceedings of the 12th International Conference on Grammatical Inference (ICGI 2014)*, JMLR Workshop Proceedings, 33–48.

Bickmore, Lee S. and Kula, Nancy C. (2013). Ternary spreading and the OCP in Copperbelt Bemba. *Studies in African Linguistics*, 42.

Bickmore, Lee S. and Kula, Nancy C. (2015). Phrasal phonology in Copperbelt Bemba. *Phonology*, 32. 147–176.

Bird, Steven and Ellison, T. Mark (1994). One-level phonology: Autosegmental representations and rules as finite automata. *Computational Linguistics*, 20.

Cassimjee, F. and Kisseberth, C.W. (2001). Zulu tonology and its relationship to other Nguni languages. In Shigeki Kaji (ed.), *Cross-linguistic studies of tonal phenomena: tonogenesis, Japanese accentology, and other topics.*, 327–359. Institute for the Study of Languages and Cultures of Asia and Africa (ILCAA).

Chandlee, Jane (2014). *Strictly Local Phonological Processes*. Ph.D. thesis, University of Delaware.

Chandlee, Jane, Athanasopoulou, Angeliki, and Heinz, Jeffrey (2012). Evidence for classifying metathesis patterns as subsequential. In Jaehoon Choi, E. Alan Hogue, Jeffrey Punske, Deniz Tat, Jessamyn Schertz, and Alex Trueman (eds.), *WCCFL 29*. Somerville, MA: Cascadilla Press.

Chandlee, Jane and Heinz, Jeffrey (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*, 42–51. Association for Computational Linguistics, Montreal, Canada.

Chomsky, Noam (1956). Three models for the description of language. *IRE Transactions on Information Theory*, 113–124.

Clements, G. N. (1976). *Vowel Harmony in Nonlinear Generative Phonology: An Autosegmental Model*. Bloomington: Indiana University Linguistics Club Publications.

Clements, G. N. (1977). Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. In *NELS 7*, 49–64.

Do, Young Ah and Kenstowicz, Michael (2011). A note on phonological phrasing in South Kyungsang. Ms., MIT.

Downing, Laura (2001). How ambiguity of analysis motivates stem change in Durban Zulu. *UBC Working Papers in Linguistics*.

Elgot, C. C. and Mezei, J. E. (1965). On relations defined by generalized finite automata. *IBM Journal of Research and Development*, 9. 47—68.

Folia, Vasiliki, Uddn, Julia, De Vries, Meinou, Forkstam, Christian, and Petersson, Karl Magnus (2010). Artificial language learning in adults and children. *Language Learning*, 60. 188–220. ISSN 1467-9922. doi:10.1111/j.1467-9922.2010.00606.x.

Fortnow, Lance (2009). The status of the p versus np problem. *Communications of the ACM*, 52:9. ISSN 0001-0782.

Gafos, Adamantios (1999). *The articulatory basis of locality in phonology.* Garland.

Gainor, Brian, Lai, Regine, and Heinz, Jeffrey (2012). Computational characterizations of vowel harmony patterns and pathologies. In *WCCFL 29*, 63–71.

Goldsmith, John A. (1990). *Autosegmental & Metrical Theory*. Basil Blackwell, Inc.

Good, Jeff (2004). Tone and accent in Saramaccan: charting a deep split in the phonology of a language. *Lingua*, 114. 575–619.

Graf, Thomas (2010). *Logics of Phonological Reasoning*. Master's thesis, University of California, Los Angeles.

Hale, Ken and Selkirk, Elisabeth (1987). Government and tonal phrasing in Papago. *Phonology Yearbook*, 4. 151–183.

Hansson, Gunnar Ólafur (2001). *Theoretical and Typological Issues in Consonant Harmony*. Ph.D. thesis, University of California, Berkeley.

Hansson, Gunnar Ólafur (2010). *Consonant harmony: Long-distance interaction in phonology*. Berkeley: University of California Press.

Heinz, Jeffrey (2010). Learning long-distance phonotactics. *LI*, 41. 623–661.

Heinz, Jeffrey (2011). Computational phonology part I: Foundations. *Language and Linguistics Compass*, 5:4. 140–152.

Heinz, Jeffrey and Idsardi, William (2011). Sentence and word complexity. *Science*, 333:6040. 295–297.

Heinz, Jeffrey and Idsardi, William (2013). What complexity differences reveal about domains in language. *Topics in Cognitive Science*, 5:1. 111–131.

Heinz, Jeffrey, Kobele, Gregory, and Riggle, Jason (2009). Evaluating the complexity of Optimality Theory. *LI*, 40:2. 277–288.

Heinz, Jeffrey and Lai, Regine (2013). Vowel harmony and subsequentiality. In Andras Kornai and Marco Kuhlmann (eds.), *Proceedings of the 13th Meeting on Mathematics of Language*. Sofia, Bulgaria.

Hopcroft, John, Motwani, Rajeev, and Ullman, Jeffrey (2006). *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, third edition.

Hyman, Larry (1998). Positional prominence and the ‘prosodic trough’ in Yaka. *Phonology*, 15. 14–75.

Hyman, Larry (2011). Tone: Is it different? In John A. Goldsmith, Jason Riggle, and Alan C. L. Yu (eds.), *The Blackwell Handbook of Phonological Theory*, 197–238. Wiley-Blackwell.

Hyman, Larry, Katamba, Francis, and Walusimbi, Livingstone (1987). Luganda and the strict layer hypothesis. *Phonology Yearbook*, 4. 87–108.

Hyman, Larry and Katamba, Francis X. (2010). Tone, syntax and prosodic domains in Luganda. In Laura Downing, Annie Rialland, Jean-Marc Beltzung, Sophie Manus, Cdric Patin, and Kristina Riedel (eds.), *Papers from the Workshop on Bantu Relative Clauses*, volume 53 of *ZAS Papers in Linguistics*, 69–98. ZAS Berlin.

Idsardi, William J. (2006). A simple proof that Optimality Theory is computationally intractable. *LI*, 37. 271–275.

Inkelas, Sharon (1995). The consequences of Optimization for underspecification. In Eugene Buckley and S. Iatridou (eds.), *NELS 25*, 287–302.

Ito, Junko, Mester, Armin, and Padgett, Jaye (1995). Licensing and underspecification in Optimality Theory. *LI*, 571–613.

Jardine, Adam and Heinz, Jeffrey (2015). A concatenation operation to derive autosegmental graphs. In *Proceedings of the 14th Meeting on the Mathematics of Language (MoL 2015)*, 139–151. Association for Computational Linguistics, Chicago, USA.

Johnson, C. Douglas (1972). *Formal aspects of phonological description*. Mouton.

Kaplan, Ronald and Kay, Martin (1994). Regular models of phonological rule systems. *Computational Linguistics*, 20. 331–78.

Kay, Martin (1987). Nonconcatenative finite-state morphology. In *Proceedings, Third Meeting of the European Chapter of the Association for Computational Linguistics*, 2–10.

Kenstowicz, Michael (1994). *Phonology in Generative Grammar*. Blackwell Publishing.

Kidima, Lukowa (1990). Tone and syntax in Kiyaka. In Sharon Inkelas and Draga Zec (eds.), *The Phonology–Syntax Connection*, 195–216. Chicago: the University of Chicago Press.

Kidima, Lukowa (1991). *Tone and Accent in KiYaka*. Ph.D. thesis, University of California, Los Angeles.

Kisseberth, C. W. (1984). Digo tonology. In G.N. Clements and John A. Goldsmith (eds.), *Autosegmental Studies in Bantu Tone*, 105–182. Foris Publications.

Kisseberth, Charles and Odden, David (2003). Tone. In Derek Nurse and Gérard Philippson (eds.), *The Bantu Languages*. New York: Routledge.

Kisseberth, Charles W. (1970). On the functional unity of phonological rules. *LI*, 1:3. 291–306. ISSN 00243892.

Kornai, András (1995). *Formal Phonology*. Garland Publication.

Lai, Regine (2012). *Domain Specificity in Learning Phonology*. Ph.D. thesis, University of Delaware.

Lai, Regine (2015). Learnable versus unlearnable harmony patterns. *Linguistic Inquiry*, 46. 425–451.

Laughren, Mary (1984). Tone in Zulu nouns. In G.N. Clements and John A. Goldsmith (eds.), *Autosegmental Studies in Bantu Tone*, 183–235. Foris Publications.

Leben, W. R. (1973). *Suprasegmental phonology*. Ph.D. thesis, Massachussets Institute of Technology.

Macdonell, Arthur (1910). *Vedic grammar*. Trūbner.

McCarthy, John (2010). Autosegmental spreading in Optimality Theory. In John A. Goldsmith, Elizabeth Hume, and W. Leo Wezels (eds.), *Tones and Features: Phonetic and Phonological Perspectives*, 195–222. De Gruyer Mouton.

McCarthy, John J. (1986). OCP effects: gemination and antigemination. *LI*, 17. 207–263.

McMullin, Kevin and Hansson, Gunnar Ólafur (2015). Locality in long-distance phonotactics: evidence for modular learning. In *NELS 44, Amherst, MA*. In press.

McWhorter, John and Good, Jeff (2012). *A Grammar of Saramaccan Creole*. De Gruyter Mouton.

Medvedev, Y. T. (1964). On the class of events representable in a finite automaton. In E. F. Moore (ed.), *Sequential machines – Selected papers*, 215–227. New York: Addison-Wesley.

Meyers, Scott (1997). OCP effects in Optimality Theory. *NLLT*, 15:4. 847–892.

Mohri, Mehryar (1997). Finite-state transducers in language and speech processing. *Computational Linguistics*, 23:2. 269–311.

Monier-Williams, Monier (1899). *A Sanskrit-English Dictionary*. Great Britain: Oxford University Press. Maintained online at http://www.sanskrit-lexicon.uni-koeln.de/.

Moreton, Elliott (2008). Analytic bias and phonological typology. *Phonology*, 25. 83–127.

Moreton, Elliott and Pater, Joe (2012). Structure and substance in artificial-phonology learning. part i: Structure. *Language and Linguistics Compass*, 6. 686—701.

Nevins, Andrew (2010). *Locality in Vowel Harmony*. MIT Press.

Odden, David (1986). On the role of the Obligatory Contour Principle in phonological theory. *Lg*, 62:2. 353–383.

Padgett, Jaye (1995). Partial class behavior and nasal place assimilation. In K. Suzuki and D. Elzinga (eds.), *Proceedings of the 1995 Southwestern Workshop on Optimality Theory*.

Pater, Joe (2004). Austronesian nasal substitution and other NC effects. In John McCarthy (ed.), *Optimality Theory in Phonology: A Reader*, 271–289. Oxford and Malden, MA: Blackwell.

Payne, Amanda (2014). Dissimilation as a subsequential process. In *NELS 44*.

Prince, Alan and Smolensky, Paul (1993). Optimality Theory: Constraint interaction in generative grammar. *Rutgers University Center for Cognitive Science Technical Report*, 2.

Prince, Alan and Smolensky, Paul (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing.

Ringen, Catherine O. and Vago, Robert M. (1998). Hungarian vowel harmony in Optimality Theory. *Phonology*, 393–416.

Rogers, James and Hauser, Marc D. (2010). The use of formal language theory in studies of artificial language learning: a proposal for distinguishing the differences between human and nonhuman animal lerners. In Harry van der Hulst (ed.), *Recursion and Human Language (Studies in Generative Grammar [SGG]) 104*. de Gruyter.

Rogers, James, Heinz, Jeffrey, Fero, Margaret, Hurst, Jeremy, Lambert, Dakotah, and Wibel, Sean (2013). Cognitive and sub-regular complexity. In *Formal Grammar*, volume 8036 of *Lecture Notes in Computer Science*, 90–108. Springer.

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

Rose, Sharon and Walker, Rachel (2004). A typology of consonant agreement as correspondence. *Lg*, 80. 475–531.

Roundtree, S. Catherine (1972). Saramaccan tone in relation to intonation and grammar. *Lingua*, 29. 308–325.

Ryan, Kevin (2015). Attenuated spreading in sanskrit retroflex harmony. Ms., Harvard University.

Safir, Ken (1982). Nasal spreading in Capanahua. *LI*, 13. 689–694.

Schein, Barry and Steriade, Donca (1986). On geminates. *LI*, 17. 691–744.

Shih, Stephanie and Inkelas, Sharon (2014). A subsegmental correspondence approach to contour tone (dis)harmony patterns. In John Kingston, Claire Moore-Cantwell, Joe Pater, and Robert Staubs (eds.), *Proceedings of the 2013 Meeting on Phonology (UMass Amherst)*, Proceedings of the Annual Meetings on Phonology. LSA.

Staubs, Robert (2014). *Computational modeling of learning biases in stress typology*. Ph.D. thesis, University of Massachusetts, Amherst.

Suzuki, Keiichiro (1998). *A typological investigation of dissimilation*. Ph.D. thesis, University of Arizona, Tucson.

Walker, Rachel (1998). *Nasalization, neutral segments, and opacity effects*. Ph.D. thesis, University of California, Santa Cruz.

Whitney, William D. (1889). *Sanskrit grammar*. Oxford University Press.

Wiebe, Bruce (1992). *Modelling Autosegmental Phonology with Multi-Tape Finite State Transducers*. Master's thesis, Simon Fraser University.

Wilson, Colin (2003). Analyzing unbounded spreading with constraints: marks, targets, and derivations. Unpublished ms.

Wilson, Colin (2006a). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science*, 30. 945–982.

Wilson, Colin (2006b). Unbounded spreading is myopic. Paper presented at the workshop Current Perspectives on Phonology.

Yip, Moira (2002). *Tone*. Cambridge University Press.

Yli-Jyrä, Anssi (2013). On finite-state tonology with autosegmental representations. In *Proceedings of the 11th International Conference on Finite State Methods and Natural Language Processing*, 90–98. Association for Computational Linguistics.