A contrastivist approach to the emergence of sound inventories

Jade J. Sandstedt
jsandstedt@gmail.com
jsandstedt.hcommons.org

25. May 2019

Contents

Part I
Talk handout

Outline

Topic: problematising the relationship between sound inventories and sets of active features in contrastivist approaches

- **Contrastivist Hypothesis**: a language's features are defined by its lexical contrasts (D. C. Hall 2007, p. 20).

Problem: unclear how the top-down focus of the CH works in language acquisition

Solution: a correlate bottom-up approach

- Features emerge when parsing linguistic input and accumulate in language acquisition
- Sum of features defines a sound inventory
  - reversing the definining relationship between contrasts and features provides a more explicit trajectory
  - speech signal → features → segments → inventories

1 Introduction

1.1 Contrasts, features, and the Contrastivist Hypothesis

Basic observation:

- size/shape of a language's sound inventory ~ active phonological features are correlated (D. C. Hall 2007; Dresher 2009, 2018; Mackenzie 2013, 2016; Iosad 2017)

This is predicted by the Contrasivist Hypothesis in (1).

(1) **Contrastivist Hypothesis** (D. C. Hall 2007, p. 20)
The phonological component of a language *L* operates only on those features which are necessary to distinguish the phonemes of *L* from one another

☞ In other words, the phonological component of your language doesn't have access to all features which in theory could occur in human language

- only those that are crucially required to define your lexical contrasts

1.2 Implementing the Contrastivist Hypothesis

Step 1: Identify contrasts

(2) **Sample of Chewa (N.31) vowel contrasts** (Downing & Mtenje 2017, ch. 3)[1]

a. túm- 'send'
b. gɔ́n- 'sleep'
c. phík- 'cook'
d. tsɛ́k- 'close'

☞ **Acquired contrasts**: {u, ɔ, i, ɛ}

Step 2: Define features

The minimal amount of features required to define any inventory of segments can be determined using the Successive Division Algorithm (SDA) in (3)

(3) **Successive Division Algorithm**

a. Begin with *no* feature specifications: assume all sounds are allophones of a single undifferentiated phoneme.

b. If the set is found to consist of more than one contrasting member, select a feature and divide the set into as many subsets as the feature allows for.

c. Repeat step (b) in each subset: keep dividing up the inventory into sets, applying successive features in turn, until every set has only one member.

The SDA produces contrastive feature hierarchies, wherein segmental relations are modeled via hierarchically nesting featural contrasts within the scope of other features. As described by Iosad (2017, p. 42):

> The hierarchy is essentially a bootstrapping device, which allows the learner to introduce order into the system of phonological contrasts by breaking the phonological space down into more manageable subinventories.

[1] Like many Eastern and Southern Bantu languages, Chewa displays penultimate lengthening as a reflex of predictable stress placement: e.g. [gɔɔ́n-a] 'to sleep' but [gɔ́n-ɛɛ́l-á] 'to sleep on something' (Hyman 2009, Downing & Mtenje 2017). For ease of explication, I do not represent vowel lengthening in this talk.

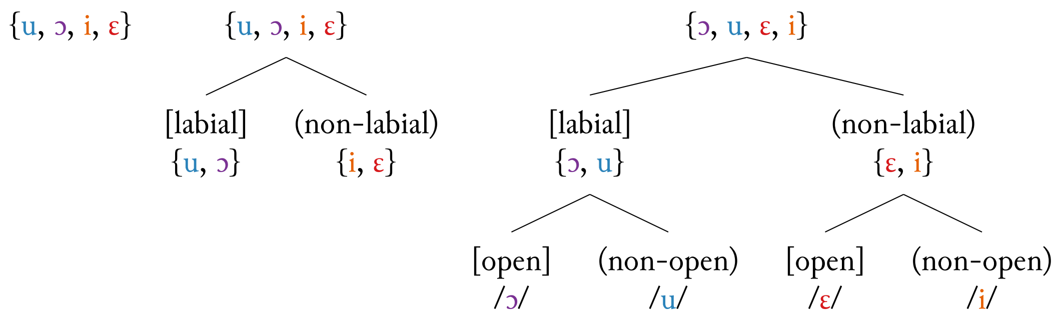

Figure 1: Sample Chewa contrastive feature hierarchy

1.3 The acquisition problem in contrastivist approaches

Contrastivist approaches propose a specific ordering in the acquisition of contrasts and features:

1. Acquire segmental contrasts
2. Define features

Questions:

1. How do language learners acquire contrasts in the absence of features?
2. Once a contrast is acquired, how do language learners select their features?
   - e.g. /ɛ/ vs. /i/ : [open], [close], [ATR], [RTR], something else?
   - e.g. /ɛ/ vs. /ɔ/ : [labial], [dorsal], [coronal], something else?

☞ it is unclear how we go from contrasts → features

As discussed by D. C. Hall & K. C. Hall (2016, p. 4), a complete learning algorithm consistent with the Contrastivist Hypothesis 'would need to elaborate what it takes to identify the presence of a phonemic contrast and how the learner selects the features to assign'.

Crux of the problem:
The top-down focus of contrastivist approaches:

- phonemic inventory → features

* requires serious abstraction

* taking the phonemic inventory as the starting point for defining features is not intuitively the right way around

2 A bottom-up contrastivist approach

Contrastivist insights are not only expressible via top-down limitations on sets of features

- we can capture the same generalisations from the bottom-up following the emergence of features and contrasts

First, we can define the set of contrasts from features using a correlate bottom-up version of the contrastivist hypothesis, as defined in (4).

(4) **Correlate Contrastivist Hypothesis**
The phonemes of a language *L* are equal to the sum of features and feature co-occurrence restrictions which are minimally necessary for the expression of phonological regularities in *L*.

☞ features → phonemes

- not features ↚ phonemes

Same basic road map

- Observe phonological patterns
  - tsɛ́k- vs. phík-
- Generalise a feature for any observed distinction
  - [F]
    - drawing insights from emergent (abstract) and substance-free feature theories (cf. Mielke 2008; Cristà, Seidl & Francis 2011; Iosad 2017)
- Presence/absence of a feature mechanically defines a contrast
  - [F] /a/ vs. (non-F) /b/
- Realised according to language-particular rules of phonetic implementation
  - [F] /ɛ/ vs. (non-F) /i/
  - 'phonetics does not determine phonological destiny', nor vice versa (cf. D. C. Hall 2014; Anderson 1981; Hale & Reiss 2000)

☞ speech signal → features → contrasts

2.1 Micro-cue model of acquisition

Phonological acquisition:

- based on insights from Westergaard's (2009, 2013, 2014) model of micro-cues

Acquisition via micro-cues

Assumptions:

- children are sensitive to fine linguistic distinctions
- generalise small pieces of abstract linguistic structure ('micro-cues') when parsing linguistic input
  - e.g. a cue for OV word order is generalised as $_{\text{VP}}$[DP V]

☞ Micro-cues accumulate in the course of language acquisition
  - the sum of which defines the linguistic grammar

Adapting the micro-cue model to phonology

Westergaard's micro-cue acquisition model can be adapted to phonology in the following way. When parsing phonological input, language learners posit:

- emergent, privative features and feature co-occurrence restrictions
  - e.g. [F], [G], and *[F, G]

Emerging inventories:

- micro-cues accumulate in the course of language acquisition
  - defining a set of contrasts, as in (5)

(5) **Segment inventory defined by [F], [G], and *[F, G] cues**

Micro-cues	Phonemes
[F]	/a/
[ ]	/b/
[G]	/c/
*[F, G]	*/d/

A working phonological learning algorithm

When parsing phonological input:

1. For every observed contrast/alternation
   - generalise a unique representational micro-cue in the form of a feature [F]
   - e.g. [phí̱k-] vs. [tsɛ̱́k-] = [F]

2.2 Chewa test case: distinctions in lexical meaning and phonological behaviour

Chewa vowel contrasts and alternations

The data in (6) evidence root vowel alternations and height harmony via vowel lowering of /i/→[ɛ] following non-high vowels.

(6) **Chewa contrasts and height harmony**[2]

HIGH	phí̱k-il-	'cook'-APPL.	tú̱m-il-	'send'-APPL.
MID	tsɛ̱́k-ɛl-	'close'-APPL.	gɔ̱́n-ɛ́l-	'sleep'-APPL.

Generalising Chewa representational micro-cues

The accurate generalisation/acquisition of the patterns in (6) requires the language learner to observe three important distinctions, outlined in (7) below.

(7) **Generalising Chewa vocalic representational micro-cues**

	Patterns	Surface generalisations			Micro-cue
a.	tsɛ̱́k-ɛl- phí̱k-il-	[ɛ] [open]	vs. vs.	[i] (non-open) contrasts/harmony	[open]
b.	tú̱m-il- phí̱k-il-	[u] [labial]	vs. vs.	[i] (non-labial) contrasts	[labial]
c.	gɔ̱́n-ɛ́l- tú̱m-il-	[ɔ] [labial, open]	vs. vs.	[u] [labial] contrasts	[labial, (open)]

Defining inventories using a set of emergent features

The set of [open]/[labial] representational micro-cues define the following sound inventory in (8).

[2] Chewa verbs display additional affixes outside the harmony domain (i.e. the verbal prefix [ku-] and the so-called 'final vowel suffix' [-a]). For simplicity of presentation, I do not represent these affixes here, which are always extra-harmonic/non-harmonising: e.g. [ku-tsɛ̱́k-ɛl-a], not *[kɔ-tsɛ̱́k-ɛl-ɛ].

(8) **Chewa vowels defined by [open] and [labial] cues**

Micro-cues	Phonemes
[open]	/ɛ/
[ ]	/i/
[labial]	/u/
[open, labial]	/ɔ/

Reinterpretation of the issues

Working in this way from the bottom-up turns the issues on their head, providing much more satisfactory answers to how contrastivist ideals pair with language acquisition.

Questions:

1. ~~How do language learners acquire contrasts in the absence of features?~~
    - contrasts are mechanically defined by the emergent set of features
2. ~~Once a contrast is acquired, how do language learners select their features?~~
    - features predictably emerge when parsing linguistic input

3 Conclusions

A bottom-up approach to the emergence of sound inventories:

1. recapitulates contrastivist insights
    - upper bound on number of features
    - correlation between sound inventories ~ features
    - predicts a minimalist representational architecture
2. provides an explicit account of how representations are acquired

The broader picture

We have explored the minimal amount of work we must do when acquiring a language to get the bare minimum of representational structure we need to do phonology, but we've done so in about as simplified/idealised conditions as possible – using only:

- Small, symmetric inventory and only two features
    - /ɛ, i, ɔ, u/ and [open]/[labial]

What about?:

- asymmetric inventories (feature co-occurrence restrictions)

- interpretation of different kinds of segmental phonological patterns
  - e.g. locality variation (hierarchical organisation of features)
  - e.g. grouped behaviour of features (geometric grouping into classes)
- the formalisation of representational micro-cues
- extended handout online at `jsandstedt.hcommons.org`

☞ Sandstedt (2018, ch. 2–3)

Part II
Handout appendix: The broader picture

4 Asymmetric inventories and feature co-occurrence restrictions

Asymmetric inventories are indicative of restrictions on the co-occurrence of some two features:

- positively evidenced by *phonological inactivity*
  - e.g. when some [G]-specified segment is an invalid recipient of [F]-harmony
- negatively evidenced by lack of contrasts
  - e.g. asymmetric [F] /a/, [ ] /b/, [G] /c/, but *[F, G] */d/

4.1 Asymmetric Chewa height contrasts

Chewa displays asymmetric height distinctions, as illustrated in (9):

- High /i, u/
- Mid /ɛ, ɔ/
- Low /a/

For simplicity's sake, we will ignore labial contrasts from now on.

(9) **Chewa height contrasts**

High	phị́k-il-	'cook'-appl.
Mid	tsɛ̣́k-ɛl-	'close'-appl.
Low	vạ́l-il-	'get dressed'-appl.
	*vạ́l-ɛl-	

Phonetics ≠ phonology
Although /a/ is articulatorily/acoustically 'non-high' – it is not obvious that it's specified for any 'non-high' (i.e. [open]) feature.

/a/ is both a non-trigger (9) and non-target (10):

- failing to trigger harmonic lowering:
  - e.g. /vạ́l-il-/ → [vạ́l-il-], not *[vạ́l-ɛl-]
- failing to undergo harmonic raising:
  - e.g. /chịnga-il-/ → [chịnga-il-], not *[chịngɘ-il-]

- /a/ prohibits lowering harmony from spreading across itself:
  - e.g. /wɛ̱lam-il-/ → [wɛ̱lam-il-], not *[wɛ̱lam-ɛl-]

(10) **Chewa neutral blocking /a/ with applicative and causative [-il, -its] suffixes**

High	chi̱nga-il-	'welcome someone'	lu̱ngam-its-	'make righteous'
Non-high	wɛ̱lam-il-	'bend'	pɔ̱lam-il-	'stoop'
	*wɛ̱lam-ɛl-		*pɔ̱lam-ɛl-	

The neutral harmony patterns in (9–10):

- provide multiple diagnostics evidencing an *[open, low] co-occurrence restriction
  - generalised in the same way as featural contrasts (11)

(11) **Generalising Chewa height representational micro-cues**

	Patterns	Surface generalisations			Micro-cue
a.	tsɛ̱́k-ɛl- phí̱k-il-	[ɛ] [open]	vs. vs.	[i] (non-open) contrasts/harmony	[open]
b.	vá̱l-il- phí̱k-il-	[a] [low]	vs. vs.	[i] (non-low) contrasts	[low]
c.	wɛ̱lam-il- *wɛ̱lɐm-ɛl-	[a] [low]	vs. vs.	*[ɐ] *[open, low] contrasts	*[open, low]

The combination of the generalised representational micro-cues in (11)

- define Chewa assymetric height contrasts in (12)

(12) **Segment inventory defined by [open], [low], and *[open, low] cues**

Micro-cues	Phonemes
[open]	/ɛ/
[ ]	/i/
[low]	/a/
*[open, low]	*/ɐ/

5 Relation between features

5.1 Hierarchical organisation of features

Inherent dichotomy in phonological heuristics

The set of features and their interrelation are informed by distinct diagnostics

Set of features:

- informed by *phonological activity*
  - e.g. contrasts, alternations, etc.
- represented as a set of micro-cues: e.g. [F], [G], *[F, G]

Relation between features:

- informed by *phonological visibility*
  - e.g. in/visibility of [G]-specified segments to [F]-processes
- represented via ordered/geometric scope asymmetries in contrastive feature hierarchies (Fig. 2)

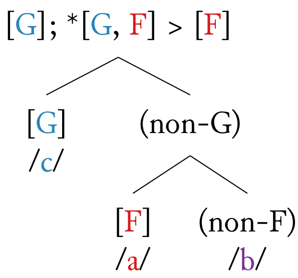

Figure 2: [G] and [F] scope asymmetries in a contrastive hierarchy

The feature hierarchy architecture

In Sandstedt (2018), I model the organisation and relation between phonological features using a version of Contrastive Hierarchy Theory (CHT; Dresher, Piggott & Rice 1994; D. C. Hall 2007; Dresher 2003, 2009; Mackenzie 2013, 2016; Iosad 2017) which incorporates emergent features and feature-nodes.

Key idea: While the existence and co-occurrence of features in a language is evidenced by its phonological activity (i.e. contrasts and alternations)

- the language-particular organisation and shape of feature classes is informed by phonological visibility (i.e. locality asymmetries such as harmony transparency vs. blocking).

6 Contrastive feature hierarchies

In CHT, segmental and featural relations are defined via hierarchically nesting featural contrasts within the scope of other featural contrasts.

- A simple abstract example is provided in Fig. 3
  - assumes the ordered set of micro-cues [G]; *[G, F] > [F]
    - producing three segments [G] /c/, [F] /a/, and non-G/non-F /b/

Figure 3: Classes and sub-classes in a privative contrastive feature hierarchy

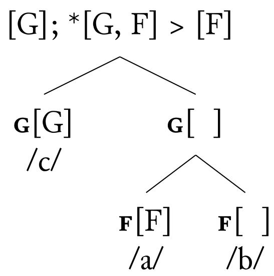

/c/	/a/	/b/
G[G]	G[ ]	G[ ]
∅	F[F]	F[ ]

Sandstedt's (2018) version of CHT assumes nesting relationships between a language's featural contrasts which depend on feature-specific nodes F and G.

- e.g. [F]-contrasts in Fig. 3 are a sub-distinction of non-G segments.

Contrastivity for the feature [F] is defined by bearing an F feature-node

- distinguishing the (contrastive) set of segments F[(F)] /a, b/ from non-contrastive under-specified (F-node-less) /c/ segments.

Sub-inventories of the contrastive set /a, b/ are distinguished by feature specifications:

- the marked or dominant sub-inventory /a/ is specified [F]
- the unmarked or recessive (non-F) /b/ sub-inventory is non-specified for [F]

6.1 The Successive Division Algorithm

Feature inventories and sub-inventories are derived by the Successive Division Algorithm (SDA). A slightly simplified version of the revised Successive Division Algorithm from Sandstedt (2018, 42) is provided in (13).

(13) **Successive Division Algorithm**

a. The input (I) to the algorithm is one or more ordered feature and feature co-occurrence micro-cues.

b. If I is found to contain a feature, then it is divided into two (non-empty) sub-inventories: a marked set M, to which is assigned F[F], and its unmarked complement set $\bar{M}$, to which is assigned F[ ], obeying *[F, G] co-occurrence restrictions.

c. M and $\bar{M}$ are then treated as the input to the algorithm; the process continues until all feature cues are divided

The SDA consists of three important components:

1. Representational micro-cues are hierarchically divided into binary-branching feature classes
   - hierarchical organisation of features
2. Each sub-inventory is associated with an emergent feature-node
   - geometric grouping into classes
3. The relative hierarchical ordering of features is cross-linguistically variable
   - emergent or cross-linguistically varying phonological classes

Sub-inventories of each feature contrast (e.g. in Chewa [open] /ɛ, ɔ/ vs. non-open /i, u/) are differentiated by feature non-/specifications

- the marked (dominant) class is assigned a feature-node F as well as a privative feature specification [F]
  - e.g. OPEN[open] /ɛ, ɔ/
- the unmarked (recessive) class bears an empty or bare node F[ ] and is non-specified for the feature
  - e.g. OPEN[ ] /i, u/

6.2 Feature scope asymmetries

Following the representations in Fig. 3, harmony locality can be accounted for in a traditional way by assuming strictly local feature spreading between [F]-specified and non-specified segments.

- Feature-nodes provide the landing sites for harmonic spreading (Fig. 4; cf. Avery & Rice 1989, Odden 1994)

Fig. 4 illustrates feature spreading between harmony triggers and targets across transparent or non-contrastive, underspecified segments.

☞ Note that the feature specifications and order of feature nodes in Fig. 4 are defined by the specifications and order of featural divisions in Fig. 3.

For an exploration of the full typology of vowel harmony behaviour types predicted by this framework, see Sandstedt (2018, §3,4).

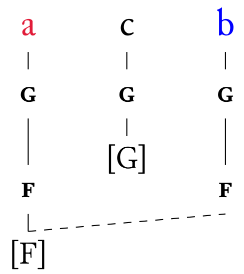

Figure 4: Local [F]-spreading between contrastively specified triggers and non-specified targets

6.3 Locality variation

The hierarchical organisation of features is informed chiefly by phonological visibilty asymmetries

- e.g. locality variation in harmony processes

Harmony neutrality variation in Bantu

Chewa (N.31) and Ndendeule (N.101) have similar low, mid, and high contrasts (Downing & Mtenje 2017; Ngonyani 2004).[3]

- cf. (14)

(14) **High, mid, and low vowel contrasts in Chewa and Ndendeule**

High	phík-il-	'cook'-APPL.
Mid	tsɛ́k-ɛl-	'close'-APPL.
Low	vál-il-	'get dressed'-APPL.
	*vál-ɛl-	

(a) Chewa

yib-il-	'steal'-APPL.
yemb-el-	'sing'-APPL.
kang-il-	'push'-APPL.
*kang-el-	

(b) Ndendeule

Low vowels in both Chewa and Ndendeule fail to trigger lowering (14)

- e.g. Chewa /vál-il-/ → [vál-il-], not *[vál-ɛl-]
- e.g. /kang-il-/ → [kang-il-], not *[kang-el-]

Variation in /a/-visibility

Low vowels are non-triggers of lowering in both Chewa and Ndendeule, but:

- display variation in lowering *visibility* in target (non-initial) positions (15)
  - Ndendeule /a/ is both *inactive* (non-trigger) and *invisible* (non-target)

[3]Ndendeule also displays additional tongue root /e, o/ vs. /ɛ, ɔ/ contrasts on mid vowels. For simplicity's sake, we will ignore these contrasts in the following analysis.

* i.e. harmony transparency: /ko̲βal-il-/ → [ko̲βal-el-], not blocked *[ko̲βal-il-]
- Chewa /a/ is *inactive* (non-trigger) but *visible* (blocking target)
  * i.e. neutral blocking: /pɔ̲lam-il-/ → [pɔ̲lam-il-], not transparent *[pɔ̲lam-ɛl-]

(15) **Variation in /a/-in/visibility in Bantu height harmony**

a) Ndendeule (N.101) transparent /a/:

ka̲ng-il-	'push'	inactive	[a…i]
ko̲βal-el-	'stumble'	invisible	[o…a…e], not *[**o**…a…**i**]

b) Chewa (N.31) neutral blocking /a/:

vá̲l-il-	'get dressed'	inactive	[a…i]
pɔ̲lam-il-	'stoop'	visible	[ɔ…a…i], not *[**ɔ**…a…**ɛ**]

Locality variation as feature scope asymmetries

As illustrated in Fig. 6, variation in /a/-visibility is captured in contrastive hierarchy by feature ranking asymmetries

- Ndendeule assumes [low] > [open]
  - '[low] is *outside* the scope of [open]'
  - /a/ is therefore *non-contrastively underspecified* for [open] (i.e. invisibly ∅)
- Chewa assumes [open] > [low]
  - '[low] is *within* the scope of [open]'
  - /a/ is *contrastively non-specified* for [open] (i.e. visibly **OPEN**[ ])

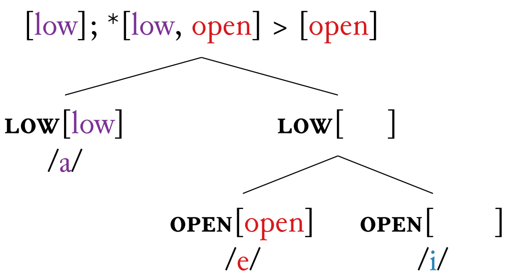

/a/	/e/	/i/
LOW[low] ∅	LOW[ ] OPEN[open]	LOW[ ] OPEN[ ]

(a) Ndendeule: [low] > [open]

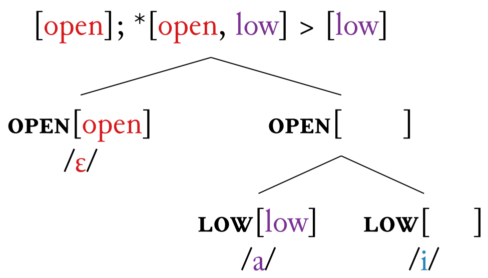

/a/	/ɛ/	/i/
OPEN[ ] LOW[low]	OPEN[open] ∅	OPEN[ ] LOW[ ]

(b) Chewa: [open] > [low]

Figure 6: Alternative contrastive feature hierarchies of Bantu vowels

In the following sections I illustrate how the asymmetries in feature underspecification in Fig. 6 result in the observed locality differences in Ndendeule and Chewa.

6.4 Harmony grammar

The basic insights of Bantu height harmony can be captured by the simple licensing principle in (23)

- adapted from Iosad (2017, pp. 52–54) and Walker (2005) – inspired by Nevins (2010)

(16) LICENSE(NON-INITIAL-V–OPEN, [open]):
'Non-initial vowels which are contrastive for [open] should be associated with [open]'

The licensing principle in (23) specifies:

1. what positions harmonise
2. for what feature.

LICENSE(NON-INITIAL-V–OPEN, [open]) motivates [open]-spreading from initial to non-initial syllables in both Chewa (17) and Ndendeule (18).

☞ Note that the feature specifications and order of feature nodes in Fig. 17 are defined by the specifications and order of featural divisions in Fig. 6b.

(17) **Chewa height harmony as privative [open]-spreading**

(a) High harmony as [open] non-spreading (b) Non-high harmony as [open] spreading

Ndendeule displays the reversed ranking of [open] and [low] features, but displays the same essential harmony patterns (18).

(18) **Ndendeule height harmony as privative [open]-spreading**

(a) High harmony as [open] non-spreading (b) Non-high harmony as [open] spreading

Transparency vs. neutral blocking

Ndendeule transparency in (19a):

- displays a [low] > [open] ranking
- /a/ is therefore non-contrastively underspecified for [open]
  - i.e. lacks any **OPEN** node and is therefore invisible to [open] spreading

Chewa neutral blocking in (19b):

- displays a [open] > [low] ranking
- /a/ is therefore contrastively non-specified for [open]
  - i.e. has a **OPEN** node and is therefore visible to [open] spreading
    - but harmony would derive an illicit *[open, low] output – blocking harmony

(19) **Ndendeule transparency and Chewa neutral blocking**

(a) [low] > [open] transparency (b) [open] > [low] neutral blocking

This derivational approach illustrates the crucial insights of how a simple licensing account (23) combined with representations (Fig. 6) defined by a version of the Successive Division Algorithm (13) which takes an ordered set of privative representational micro-cues implies a well-articulated, concise, and typologically accurate vowel harmony grammar. In the final section, I provide a much more explicit and detailed illustration of how this approach can be translated to ranked constraints.

7 Formalising contrastivist approaches

This appendix illustrates in greater detail how this contrastivist approach can be formalised in a more explicit grammatical model. This presentation follows closely the exploration of tongue root harmony patterns in Standard and Ifẹ Yoruba in Sandstedt (2018, §3.2).

7.1 Formalising contrastivist representations

We have seen how the acquisition of sound inventories and active phonological features may proceed from the generalisation of representational micro-cues in the form of feature labels [F] / [G] and prohibited feature co-occurrence restrictions *[F, G]. Ultimately the grammatical formalisation of these micro-cues could take different forms, and I do not commit to any one system at this time. In this section, I provide a brief illustration of one way this approach may be grammatically implemented. Regardless of the model, the important insight is that phonological operations are limited by representational generalisations. Mackenzie & Dresher (2004) and Mackenzie (2013, 2016) have devised an algorithm which captures this idea using faithfulness and feature co-occurrence constraints to convert contrastive feature hierarchies to non-derivational constraint rankings. I provide a simplified privative interpretation of this approach in (20). In this section, I outline only the broad points of this method; for finer details, see Sandstedt (2018) and Mackenzie (2013, 2016).

(20) **Converting a contrastive hierarchy to a constraint ranking** (adapted from Mackenzie 2013, p. 305)

a. Select a faithfulness constraint Max$[F]_i$, where $[F]_i$ is the highest ordered contrastive feature for which Max$[F]_i$ has not yet been ranked. Rank Max$[F]_i$ below any Max[F] constraints ranked by prior application of step (a) and above all other Max[F] constraints.

b. Above this faithfulness constraint, rank any co-occurrence constraints of the form *[ϕ, F_i] where ϕ consists of features ordered higher than $[F]_i$ and where contrastive specification of ϕ is prohibited in segments specified for $[F]_i$. If there are more contrastive features, go to (a); otherwise, end.

In the way of an example, we can convert the three-way height contrast in Ndendeule which includes features [low], [open], and the co-occurrence restriction *[low, open] in the following way. According to the first step of the algorithm in (20), the ranking of faithfulness constraints mirrors the ordering of features in contrastive feature hierarchies. In Fig. 6a, the highest ordered feature is [low]; therefore the highest ranked faithfulness constraint will be Max[low]. Step (b) in the algorithm refers to feature co-occurrence restrictions which define inventory asymmetries in the distribution of [low] and higher ordered features. Since [low] has broadest scope/highest rank in Ndendeule, this step does not apply. After this, we have a second ordered feature [open], and the process therefore repeats. Step (a) dictates that Max[open] should be ranked after the previously ordered Max[low] constraint. Step (b) requires any [open] and higher (i.e. [low]) feature co-occurrence restrictions to be ranked above Max[open]. In Ndendeule, *[low, open] co-occurrence is not permitted – excluding the feature [open] within segments specified for [low]. A co-occurrence constraint *[low, open] is therefore ordered above Max[open]. At this point, there are no other features to be ordered, and the algorithm terminates. The final constraint ranking in this example is summarised in (21).

(21) Max[low], *[low, open] » Max[open]

This combination of faithfulness and markedness achieves the featural asymmetries defined by the contrastive feature hierarchy in Fig. 6a. As illustrated in the tableau in (22), the ranking in (21) will map any fully specified inputs to contrastively specified outputs. The ordering of feature-nodes in autosegmental representations in (22) mirrors the order of features in contrastive feature hierarchies (see Fig. 6a). In (22), the input is fully specified for both features, however according to the contrastive hierarchy in Fig. 6a, [low] segments are non-contrastive for [open]. The faithful [low, open] /a/ candidate in (a) therefore violates the high ranking *[low, open] co-occurrence constraint. Deleting [open] specifications on [low] segments as in candidate (b) produces a contrastively specified segment according to Fig. 6a and is optimal, violating only the low ranked Max[open] constraint. Candidate (c) violates the high-ranked Max[low] constraint and is eliminated.

(22) **[open]-specifications are eliminated among [low]-segments**

a LOW [low] OPEN [open]	MAX[low]	*[low, open]	MAX[open]
a. a LOW [low] OPEN [open]		*!	
☞ b. a LOW [low] OPEN [open]			*
c. a LOW [low] OPEN [open]	*!		

This example provides a sample schematic of one simple way in which representational micro-cues can be grammatically implemented to produce representations consistent with contrastive feature hierarchies. The key insight is that there must be some explicit grammatical architecture defining a language's active phonological features and feature classes which interacts transparently (in derivational terms) with phonological processes. Contra Reiss' (2017, p. 26) claim that 'phonology doesn't "care" about contrast, because it has no mechanism by which to do so,' a version of the SDA incorporating contrast-defining representational micro-cues provides just such a mechanism.

7.2 Formalising a representational harmony grammar

The nuts and bolts of a harmony mechanism are really quite simple. To generalise the grammatical application of any harmony system, we need to specify at least two components: 1) what positions harmonise and 2) for what feature. For instance, the basic insights of Bantu height harmony can be captured by the simple licensing principle in (23) – adapted from Iosad (2017, pp. 52–54) and Walker (2005) – which states that non-initial vowels which are contrastive for the harmony feature [open] should be associated with [open] where possible. I assume a non-initial vowel may satisfy this rule by being specified for [open] or by local feature spreading.

(23) License(Non-initial-V–**open**, [open]):
'Non-initial vowels which are contrastive for [open] should be associated with [open]'

The licensing principle in (23) captures the basic insight of progressive (left-to-right) height harmony via vowel lowering. It dictates that non-initial, contrastively non-open segments are 'needy' in the sense of Nevins (2010) and will seek out [open]-feature specifications to copy from. Where there is no [open]-source to copy from, the harmony procedure comes up empty handed and no change occurs – resulting in 'high' harmony: for example, underlyingly non-open /phí̱k-il-/ → [phí̱k-il-] 'cook'-APPL. in Chewa (14a). On the other hand, if a local [open] feature is available, it spreads – resulting in lowering harmony: for example, /tsɛ̱́k-il-/ → [tsɛ̱́k-ɛl-] 'close'-APPL.

This licensing account could be grammatically implemented in a wide variety of ways, and I do not commit to any one framework. In the way of an illustration, building on Mackenzie's (2013, 2016) method of implementing contrastive hierarchies in OT constraint rankings, we could translate the harmony licensing principle in (23) to the more explicit [open]-licensing constraint defined in (24); cf. similar approaches in Downing & Mtenje (2017, ch. 4) and Harris (1994, 1997). OPEN→[open] requires segments which are contrastive for [open] to be specified for [open]. Combined with Max/Dep faithfulness constraints, this licensing constraint motivates feature spreading over feature insertion/deletion. I assume OPEN→[open] competes with a lower ranked DepLink[open] constraint, which effectively penalises harmonic spreading. Specifically, Dep/MaxLink constraints require input–output correspondents to preserve autosegmental associations (i.e. don't insert or delete association lines); for further discussion, see Morén (2001), Blaho (2008), and Iosad (2017). This approach provides an explicit representational account of how asymmetries in [open]/non-open specifications drives lowering harmony in Bantu, limited by contrastive hierarchy representations.

(24) **Licensing and faithfulness constraints motivating height harmony**

a. OPEN→[open]: Segments contrastive for [open] must be associated with [open].

b. DepLink[open]: If an output segment x_o is linked to $[\text{open}]_o$, then its input correspondent x_i must be linked to $[\text{open}]_i$.

c. Dep[open]: Assign a violation mark for any instance of [open] in an output that does not have an input correspondent.

d. Max[open]: Assign a violation mark for any instance of [open] in an input that does not have an output correspondent.

According to the contrastive hierarchy method, the basic grammatical machinery in height harmis exactly the same in Chewa and Ndendeule, which differ only in their representations of low vowels. I assume any faithfulness/markedness constraints introduced by Mackenzie's (2013, 2016) algorithm in section 7.1 – e.g. the MAX[open] constraint from (22) – are ranked above the licensing/faithfulness constraints defined in (24). For simplicity's sake, the example below in (25) concerns only [open] feature specifications. The tableau in (25) illustrates the analysis with a root-initial [open] vowel in [tsɛ́k-ɛl-] 'close'-APPL. The tableau demonstrates how [open] faithfulness and licensing constraints in combination motivate height harmspreading from initial to non-initial vowels.

(25) **[open]-licensing in a constraint hierarchy**

tsɛ́k il OPEN OPEN | [open]	MAX[open]	DEP[open]	OPEN→[open]	DEPLINK[open]
a. tsɛ́k il OPEN OPEN | [open]			*!	
b. tsɛ́k il OPEN OPEN ǂ [open]	*!		**	
c. tsɛ́k il OPEN OPEN | ¦ [open] [open]		*!		*
☞ d. tsɛ́k il OPEN OPEN | ⋰ [open]				*

The candidate in (a) [tsɛ́k-il-] is faithful but not optimal since it bears a segment contrastive for [open] which is not associated with [open]. It thus violates OPEN→[open] which motivates harmony. The surface-harmonic candidate (b) [tsík-il-] illustrates that OPEN→[open] cannot be satisfied by deleting the [open] feature since doing so violates the high-ranked MAX[open] constraint introduced by Mackenzie's (2013, 2016) algorithm in section 7.1. In a similar way, high-ranked DEP[open] motivates feature spreading over feature insertion, as illustrated by candidate (c). Candidate (d) [tsɛ́k-ɛl-] is most optimal since it obeys both of the aforementioned faithfulness and licensing constraints by spreading [open] from the initial to non-initial vowel, violating only the low-ranked DepLink[open] constraint.

In summary, the principal components of harmony processes are grammatically very simple. Assuming that faithfulness/markedness constraints introduced by the SDA have highest rank produces all typologically common categorical harmony behaviours (e.g. triggers, targets, transparent segments, blockers, etc.) as emergent effects of inventory-defining features and feature co-occurrence restrictions.

7.3 Locality variation via feature scope asymmetries

Harmony locality variation is illustrated by the tableaux in (26, 27) which incorporate all the relevant components outlined above in an illustration of Ndendeule transparency and Chewa neutral blocking low vowels using the constraints outlined earlier in (22). This analysis shows that the neutrality (non-harmonisation) of [low] segments in both varieties is a product of the lack of permitted [low, (open)] contrasts. That is, the prohibition against *[low, open] co-occurrence which defines the Chewa and Ndendeule contrastive feature hierarchies in Fig. 6 prohibits harmonic spreading to [low] vowels, resulting in [open]-disharmonic Chewa [vál-il-] 'get dressed'-APPL. and Ndendeule [kang-il-] 'push'-APPL. (14). Both languages share the same set of active harmonising and non-harmonising vowels (e.g. [open] /ɛ/ or /e/ ~ non-open /i/ vs. neutral /a/) but differ in the visibility of their non-harmonising [low] vowels, resulting in blocked [open]-harmony patterns in Chewa (e.g. [pɔlam-il-], not *[pɔlam-ɛl-] 'stoop'-APPL.) vs. transparent skipping in Ndendeule (e.g. [kobal-el-], not *[kobal-il-] 'stumble'-APPL.). These differences in the visibility of [low] segments to [open]-harmony is a simple, predicted effect of the relative scope of [open] and [low] features in contrastive feature hierarchies (see Fig. 6).

In the way of illustration, consider the tableau in (26) which illustrates the representational and grammatical constraints driving long-distance harmony in Ndendeule. First, any faithfulness/markedness constraints introduced by Mackenzie's (2013, 2016) algorithm are ranked high. The tableau in (26) therefore includes the Max[low], *[low, open] » Max[open] constraints introduced in (22). For the sake of space, inactive high-ranked Dep[open] is not shown in (26, 27) below. In Ndendeule, [low] has broader scope while [open] has narrower scope. Given the broad scope of the [low] feature and co-occurrence restriction against *[low, open] vowels, [low] vowels are underspecified for [open] – lacking any **OPEN** node. In (26), the faithful candidate (a) is disharmonic – violating the harmony licensing constraint OPEN→[open] since it has an **OPEN** node which is unassociated with an [open] specification. The harmonic candidate (b) satisfies the harmony licensing constraint by [open] spreading. Given that [low] vowels are underspecified for the harmony feature, they are not visible to local [open] spreading – resulting in transparency. Finally, candidate (c) illustrates that [open]-harmony cannot be satisfied by [open] deletion.

(26) **Transparent [low] vowels in Ndendeule: [low]; *[low, open] > [open]**

	Max[low]	*[low, open]	Max[open]	OPEN→[open]	DepLink[open]
ko bal il LOW LOW LOW OPEN [low] OPEN [open]					
a. ko bal il LOW LOW LOW OPEN [low] OPEN [open]				*!	
☞ b. ko bal il LOW LOW LOW OPEN [low] OPEN [open] (dashed link to il's OPEN)					*
c. ko bal il LOW LOW LOW OPEN [low] OPEN [open] (delinked)			*!	**	

The situation in Chewa is quite similar. The harmony mechanism is identical; Ndendeule and Chewa differ only in their inventory-defining faithfulness/markedness constraints introduced by the SDA, which follows from the alternative feature orderings in Fig. 6. Specifically, in Chewa the harmony feature [open] has broadest scope and [low] has narrower scope, resulting in the reversed order of high-ranking Max[open], *[open, low] » Max[low] constraints in (27). In contrast to Ndendeule above where [low] vowels are non-contrastively underspecified for [open] and therefore lack any **open**-node, Chewa [low] vowels are *contrastively non-specified* for [open]. In other words, [low] vowels are within the scope of [open] and therefore have an **open** feature-node in (27) below. This predicts that [low] vowels in Chewa should be associated with [open] according to open→[open]. Furthermore, since the **open** node defines harmony visibility, this account predicts low vowels should behave non-transparently with respect to [open] spreading in Chewa.

These predictions are born out in (27). In (27) the faithful candidate (a) is optimal – despite not being harmonic – since [open]-spreading would result in an illicit *[open, low] output, as in candidate (b). Following Mackenzie's (2013, 2016) algorithm, each feature is introduced by a Max constraint which outranks the open→[open] licensing constraint. A violation of *[open, low] introduced by [open]-harmony can therefore not be circumvented by deleting [low] as in candidate (c). Since feature-nodes define harmony spreading/landing sites, the fact that low vowels fail to undergo [open]-harmony in Chewa halts the process from spreading further downstream to the non-initial [open]-contrastive vowel /i/ – causing neutrally blocked surface patterns (i.e. blocked [pɔ̲lam-il-], not transparent *[pɔ̲lam-ɛl-]).

(27) ***[open, low] neutral blocking in Chewa: [open], *[open, low] > [low]**

	Max[open]	*[open, low]	Max[low]	open→[open]	DepLink[open]
pɔ lam il; OPEN OPEN OPEN; [open]; LOW LOW; [low]					
☞ a. pɔ lam il; OPEN OPEN OPEN; [open]; LOW LOW; [low]				**	
b. pɔ lam il; OPEN OPEN OPEN; [open]; LOW LOW; [low]		*!			**
c. pɔ lam il; OPEN OPEN OPEN; [open]; LOW LOW; [low]			*!		**

This sketch of Ndendeule and Chewa neutral harmony patterns illustrates how the representational insights of the contrastive hierarchy method can be incorporated into an explicit grammatical model.

☞ The important insight is that constraints on phonological representations introduced in the course of phonological acquisition limit whatever grammatical mechanism spreads harmony between segments contrastive for the harmony feature.

Consistent with other contrastive hierarchy methods, this approach posits that narrow variation in phonological visibility is predicted by variable feature ordering, as illustrated by the above contrasting neutral blocking and transparency patterns in Chewa and Ndendeule, respectively.

8 Conclusions

Goal:
Well-articulated theory about the trajectory and nature of the emergence of features and contrasts

☞ speech signal → features → segmental contrasts

Feature emergence:

- informed by *phonological activity*

Feature relations:

- informed by *phonological visibility*

Emergent feature geometry:
Feature geometric insights are captured via a version of Contrastive Hierarchy Theory

- incorporating a version of the SDA (13) which takes representational micro-cues as its input

Representational harmony grammar:
Contrastive feature hierarchy insights accurately predict locality variation in harmony systems

- eschewing grammatically specific constraints/parameters

References

Anderson, Stephen R. 1981. Why phonology isn't 'natural'. *Linguistic Inquiry* 12(4). 493–539.

Avery, Peter & Keren Rice. 1989. Segmental structure and coronal underspecification. *Phonology* 6(2). 179–200.

Blaho, Sylvia. 2008. *The syntax of phonology: A radically substance-free approach*. University of Tromsø PhD thesis.

Cristà, Alejandrina, Amanda Seidl & Alexander L. Francis. 2011. Phonological features in infancy. In G. Nick Clements & Rachid Ridouane (eds.), *Where do phonological features come from?: Cognitive, physical and developmental bases of distinctive speech categories* (Language Faculty and Beyond: Internal and External Variation in Linguistics 6), 303–26. Amsterdam: John Benjamins.

Downing, Laura J. & Al Mtenje. 2017. *The phonology of Chichewa*. Oxford: Oxford University Press.

Dresher, B. Elan. 2003. The contrastive hierarchy in phonology. *Toronto Working Papers in Linguistics* 20. 47–62.

Dresher, B. Elan. 2009. *The contrastive hierarchy in phonology*. Cambridge: Cambridge University Press.

Dresher, B. Elan. 2018. Contrastive Hierarchy Theory and the nature of features. In Wm. G. Bennett, Lindsay Hracs & Dennis Ryan Storoshenko (eds.), *Proceedings of the 35th West Coast Conference on Formal Linguistics*, 18–29. Somerville, MA: Cascadilla Proceedings Project.

Dresher, B. Elan, Glyne Piggott & Keren Rice. 1994. Contrast in phonology: Overview. *Toronto Working Papers in Linguistics* 14. iii–xvii.
Hale, Mark & Charles Reiss. 2000. 'Substance abuse' and 'dysfunctionalism': Current trends in phonology. *Linguistic Inquiry* 31(1). 157–69.
Hall, Daniel Currie. 2007. *The role and representation of contrast in phonological theory*. Toronto: University of Toronto PhD thesis.
Hall, Daniel Currie. 2014. On substance in phonology. In Laura Teddiman (ed.), *2014 CLA Conference Proceedings*. Online: http://cla-acl.ca/wp-content/uploads/Hall-2014.pdf. Canadian Linguistic Association.
Hall, Daniel Currie & Kathleen Currie Hall. 2016. Marginal contrasts and the contrastivist hypothesis. *Glossa* 1(1). 1–23.
Harris, John. 1994. Monovalency and opacity: Chicheŵa height harmony. *UCL Working Papers in Linguistics* 6. 509–47.
Harris, John. 1997. Licensing inheritance: An integrated theory of neutralization. *Phonology* 14(3). 315–70.
Hyman, Larry M. 2009. Penultimate lengthening in Bantu: Analysis and spread. *UC Berkeley Phonology Lab Annual Report*. 195–209.
Iosad, Pavel. 2017. *A substance-free framework for phonology: An analysis of the Breton dialect of Bothoa*. Vol. 2 (Edinburgh Studies in Theoretical Linguistics). Edinburgh: Edinburgh University Press.
Mackenzie, Sara. 2013. Laryngeal co-occurrence restrictions in Aymara: Contrastive representations and constraint interaction. *Phonology* 30(2). 297–345.
Mackenzie, Sara. 2016. Consonant harmony in Nilotic: Contrastive specifications and Stratal OT. *Glossa* 1(1). 1–38.
Mackenzie, Sara & B. Elan Dresher. 2004. Contrast and phonological activity in the Nez Perce vowel system. In Pawel M. Nowak, Corey Yoquelet & David Mortensen (eds.), *Proceedings of the twenty-ninth annual meeting of the Berkeley Linguistics Society*, 283–94. Berkeley: Berkeley Linguistics Society.
Mielke, Jeff. 2008. *The emergence of distinctive features*. Oxford: Oxford University Press.
Morén, Bruce. 2001. *Distinctiveness, coercion, and sonority: A unified theory of weight*. London: Routledge.
Nevins, Andrew. 2010. *Locality in vowel harmony*. Cambridge, MA: MIT Press.
Ngonyani, Deo. 2004. Vowel harmony in Kindendeule and Chingoni languages of Tanzania. *Utafiti* 5(1). 9–112.
Odden, David. 1994. Adjacency parameters in phonology. *Language* 70(2). 289–330.
Reiss, Charles. 2017. Contrast is irrelevant in phonology: A simple account of Russian /v/ as /V/. In Bridget D. Samuels (ed.), *Beyond markedness in formal phonology* (Linguistik Aktuell/Linguistics Today, 241), 23–46. Amsterdam: John Benjamins.
Sandstedt, Jade J. 2018. *Feature specifications and contrast in vowel harmony: The orthography and phonology of Old Norwegian height harmony*. Online: https://jsandstedt.hcommons.org/publications-presentations/. University of Edinburgh PhD thesis.
Walker, Rachel. 2005. Weak triggers in vowel harmony. *Natural Language & Linguistic Theory* 23(4). 917–89.

Westergaard, Marit. 2009. *The acquisition of word order: Micro-cues, information structure and economy* (Linguistik Aktuell/Linguistics Today, 145). Amsterdam: John Benjamins.

Westergaard, Marit. 2013. The acquisition of linguistic variation: Parameters vs. micro-cues. In Terje Lohndal (ed.), *In search of universal grammar: From Old Norse to Zoque*, 275–98. Amsterdam: John Benjamins.

Westergaard, Marit. 2014. Linguistic variation and micro-cues in first language acquisition. *Linguistic Variation* 14(1). 26–45.