Phonology (2026), **Vol. 43**:e11 1–33
doi:10.1017/S0952675726100372

PHONOLOGY

ARTICLE

Learning abstract underlying phonemes

B. Elan Dresher

Department of Linguistics, University of Toronto, Canada
Email: elan.dresher@utoronto.ca

Received: 31 December 2024 UTC; **Revised:** 1 June 2025 UTC; **Accepted:** 31 July 2025 UTC

Keywords: phonological abstractness; opacity; learnability; Inuit dialects; Uyghur; Esimbi; Polish

Abstract

This article discusses cases of phonological abstractness and opacity and shows how they are eminently learnable, given certain assumptions about the innate cognitive endowment that learners bring to acquisition. I argue that opacity is not a learning problem but its solution. I propose that a distinction in patterning between two types of *i* in Inuit dialects is best explained by positing that surface [i] represents a merger of two underlying vowels. The apparently similar case of /i/ in Uyghur is shown to require a different type of solution. A review of contrasting approaches to prefix selection in Esimbi shows that opacity plays no role in evaluating their relative learnability. Though a celebrated case of opacity in Polish appears to have been misanalysed, an abstract analysis can still be motivated to account for alternations in certain lexemes. Uniting these cases is a preference for phonological analyses over diacritics or suppletion.

Contents

1. Introduction **2**
2. Principles of a learning theory for phonology **3**
2.1. The Projection Problem . . . 3
2.2. Assumptions about phonological grammars . . . 4
2.3. Relationship to learning models . . . 8
3. An abstract underlying vowel in Inuit dialects **9**
3.1. Strong and weak [i] . . . 9
3.2. A learning theory for the abstract vowel in Inuit dialects . . . 10
3.3. Assigning features to weak *i* . . . 11
3.4. Weak *i* and opacity . . . 13
4. Abstract /ɯ/ in Uyghur? **14**
4.1. Backness harmony in Uyghur . . . 14
4.2. Learning exceptional harmony in Uyghur . . . 16
5. Emergent Grammar *vs.* Universal Grammar in Esimbi **18**
5.1. Emergent Grammar and Universal Grammar . . . 18
5.2. The basic learnability argument in Esimbi . . . 19
5.3. Further intricacies of the Esimbi vowel system: the A- prefix . . . 22
5.4. A UG analysis of the A- prefix in Esimbi . . . 23
5.5. Other analyses of Esimbi . . . 26
6. Reanalysing opacity in Polish **27**
6.1. The Polish problem in its classical form . . . 27
6.2. A revised view of Polish vowel raising . . . 28
7. Conclusion **29**
References **30**

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

1. Introduction

Compton & Dresher (2011) propose that in some Inuit dialects there is evidence for a covert contrast between /i/ and /ə/ which is neutralised on the surface to [i]. Mayer *et al.* (2022a: 1) reject this sort of analysis in general, suggesting that the covert contrast is not learnable (emphasis added):

> This is an instance of what Kiparsky (1973) called 'the diacritic use of phonological features' (p. 16): an underlying featural contrast is used to condition phonological behavior, despite corresponding to no observable phonetic differences *in the conditioning segments themselves.* […] These analyses therefore make strong claims about both cognitive representations (that speakers have localized the source of the idiosyncrasy to a distinction between two sounds that are identical on the surface) and learning (that there is some learning mechanism that leads to such a representation).

I will argue that the analysis of Compton & Dresher (2011) is not an example of the diacritic use of phonological features, under any reasonable interpretation of what that means. And while I don't see how it is a bad thing to make 'strong claims' about cognitive representations and learning, I will argue against the idea that the burden of proof is on those who propose abstract underlying entities as opposed to those who propose actual diacritics or other non-phonological solutions.

With respect to the sorts of evidence available to learners, Mayer *et al.* (2022a) are not the only phonologists who require 'observable phonetic differences in the conditioning segments themselves' to diagnose an underlying featural contrast. In Esimbi, all vowels in roots are high ([i, ɨ, u]) on the surface; however, some roots take prefixes with [i, u], others with [e, o] and others with [ɛ, ɔ]. Archangeli & Pulleyblank (2015: 10) reject the abstract analysis of Hyman (1988) in which root vowels of various heights all neutralise on the surface to high vowels (emphasis added):

> Assuming that a phonological difference in the roots is the source of the difference in prefix height requires that height distinctions be encoded in roots even though there is no surface evidence – *in the roots* – for the required distinction.

Archangeli & Pulleyblank (2015: 5) propose that their analysis can 'resolve the opacity problem' created by a standard generative approach. OPACITY is a term introduced by Kiparsky (1973) to describe a phonological rule whose structural description is contradicted at the surface. Kiparsky's formulation is given in (1):

(1) *Rule opacity* (Kiparsky 1973: 79)
A rule P of the form A → B / C__D is OPAQUE to the extent that there are surface representations:
   a. A in environment C__D, or
   b. B (derived from A) in environment other than C__D.

Opacity of type (1a) can occur when another rule, Q, applies after P and creates new tokens of A in the environment C___D. This type of interaction is called COUNTERFEEDING because the opposite order, where Q precedes P and creates new inputs to it, is a FEEDING order. It creates the surface appearance that P has UNDERAPPLIED. Opacity of type (1b) arises when a rule Q that applies after P changes the context C___D, making it appear that P has OVERAPPLIED. This is called a COUNTERBLEEDING order, because the opposite order, in which Q removes potential inputs from P, is a BLEEDING order. Over time additional types of opaque rule interactions have been identified (see Baković & Blumenfeld 2024 for a more complete typology), but the essential idea remains that an opaque rule is a rule that appears in some way to be contradicted at the surface.

From the beginning, it has been assumed that opaque rules pose particular learnability problems.[1] This assumption has led to attempts to constrain or completely do away with phonological opacity, or

[1] An early dissenter was Kaye (1974), who observed that some opaque rules are stable over time, and thus presumably not that hard to learn. He observed also that opacity can be beneficial to learners in recovering underlying forms.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

to prefer analyses that do not have it.[2] I will argue that the learnability problem has been misconceived: rule opacity does not pose a learning problem! Rather, opacity is a *solution* within a particular theoretical framework for a learning problem that exists independently of that framework. Given such a framework, I will argue that it is not hard to conceive of a learning mechanism that leads learners to posit abstract underlying representations and opaque rules.

Not all examples of abstractness are the same, and this article will present several case studies that illustrate different aspects of abstractness and opacity. Before looking at specific cases, §2.1 presents a brief schematic depiction of what has been called the Projection Problem for phonology and the basic conditions a learning theory for phonology has to satisfy. §2.2 presents some of the principles that I assume guide phonological learners, and §2.3 puts these principles in the context of phonological learning models. §3 reviews an analysis of Inuit dialects that posits an abstract underlying vowel and shows how it is learnable given the proposed principles of Universal Grammar (UG). Similar abstract analyses have been proposed for Uyghur; §4 shows how, despite some initial similarities, evidence from Uyghur points away from an abstract analysis of the Inuit type. I will argue, however, that the required analysis, though not involving abstract underlying vowels, still involves phonological features rather than diacritics.

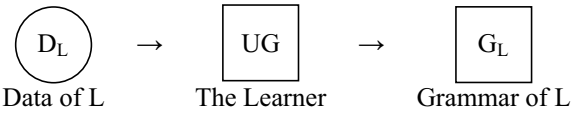
In §5, I compare the learnability of two analyses of Esimbi: an analysis based on Hyman (1988) that posits an opaque rule and abstract underlying vowels, and the analysis of Archangeli & Pulleyblank (2015) that does without opacity. I will argue that the Esimbi case, though more complicated than the Inuit one, has essentially the same sort of abstractness and is learnable in the same way. §6 reviews a well-known example of opacity in Polish which, on further investigation, requires a different analysis, though one that is still phonological and abstract. §7 is a brief conclusion.

2. Principles of a learning theory for phonology

2.1. The Projection Problem

Before delving into abstractness and opacity, I want to emphasise the obvious point that what is easy or hard to learn depends entirely on what learners bring to the task. In generative grammar, what has been called 'the logical problem of language acquisition' (Hornstein & Lightfoot 1981) can be schematically depicted as in (2):

(2) *The Projection Problem*



The diagram illustrates a learner born into a community that speaks a language L, who is exposed to a sampling of data, D_L, from L, and somehow acquires a grammar of L, G_L. What the learner has to do is known as the PROJECTION PROBLEM (Peters 1972; Baker 1979), because the learner is required to PROJECT (i.e., induce) a grammar G_L from D_L. This projection is made possible by the innate set of cognitive principles that learners are equipped with, which can convert D_L into G_L. In generative grammar, these principles have been called UG. Some understand UG in a narrow sense to refer to an innate cognitive endowment that is specific to language. However, it has sometimes been used to refer to the collection of cognitive principles that allow learners to acquire a language, whether these are specific to language or not, and this is the sense in which I will use it. While it is an interesting question how much of UG is particular to language or even phonology, this may not be easy to determine. Even principles

[2]The subject of opacity cannot be separated from the more general issue of phonological abstractness. See Krämer (2012: 38–64), Baković *et al.* (2022), Kenstowicz & Kisseberth (2022) and Wang & Hayes (in press) for discussion of the abstractness debate that arose as a reaction to Chomsky & Halle (1968). For opacity as it relates to serialism, see Baković (2011) and Pruitt (2023). Specific attempts to do away with opaque analyses are discussed in this article.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

that appear to be very specific to phonology, such as the ones presented in §2.2, might, at a higher level of abstraction, turn out to be phonological manifestations of more general principles. For example, it has been shown (Gershman & Niv 2010; Courellis *et al.* 2024) that abstract neural representations arise in all kinds of learning. What is important for my purposes here is that, by definition, UG is not learned, but is rather what the learner brings to the task of acquiring a grammar. In Bayesian terms, UG is the hypothesis space and the set of priors in the learning scenario, leaving open which innate principles are unique to language.

The diagram in (2) abstracts away from many issues surrounding the acquisition and use of language: it is manifestly unequipped to represent the time course of language acquisition, or language acquisition in a multilingual or multidialectal environment. Rather, the diagram is an *idealisation* (not a hypothesis!) that is meant to focus our attention on the two questions in (3):

(3) *Goals of generative grammar*
   a. Descriptive adequacy: What is G_L?
   b. Explanatory adequacy: What is UG?

The first goal of generative grammar (3a) is to determine G_L for each L. A correct grammar of L – one that represents the grammar internalised by speakers of L – achieves DESCRIPTIVE ADEQUACY, in the term used by Chomsky (1965). Some questions phonologists want to answer are given in (4).

(4) *Questions about the nature of grammars*
   a. What are the 'primes' of phonological representations? Do they consist of features, or unary elements, or particles, or gestures?
   b. Are these primes universal or language-particular? Are they fully specified, or do they reflect only contrastive properties?
   c. Do phonological grammars have ordered rules or ranked parallel constraints? What principles govern their interaction?
   d. Does a morpheme have a unique stored representation from which all its surface variants are derived, or multiple representations?
   e. How much phonetic detail is included in lexical representations?

These are all questions that *phonologists* argue about, but they cannot be issues for *learners*. The basic form of each G_L, and of all possible G_Ls, is determined by UG. In Chomsky's (1965) term, a correct theory of UG achieves EXPLANATORY ADEQUACY.[3]

Having to project a grammar from incomplete data requires a type of induction that inevitably involves a POVERTY-OF-THE-STIMULUS problem: without some constraints on possible hypotheses, there are too many potential ways one can generalise from any finite amount of data. Moreover, the data D_L does not by itself tell the learner what the units or operations of mental representation are. UG must be rich enough to bridge the qualitative gap between D_L and G_L. Therefore, when one asks how hard it is for a learner to acquire G_L given exposure to D_L, any answer we give must make some assumptions about UG: what may be hard to learn given UG_1 might be easy to learn given UG_2. Linguists' intuitions as to what might be easy or hard to learn thus presuppose some often unspoken assumptions about the nature of UG. In this article, I will try to make some of these assumptions explicit.

2.2. Assumptions about phonological grammars

I make the basic assumptions in (5):

[3] In the 21st century, a third question has been asked (Chomsky 2004): Why does UG have the properties that it has? Answering this question requires us to go 'beyond explanatory adequacy' to look at biological and evolutionary factors that could have shaped the language faculty. In this article, we will be concerned with the first two goals.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

(5) *Basic assumptions about how UG guides phonological acquisition*

a. Phonemes consist of combinations of smaller units (phonological primes).
b. Interactions between phonemes involve the primes of (5a).
c. Learners have access to the morphological makeup and paradigmatic membership of lexical items.
d. Learners attempt to arrive at a single underlying form for each morpheme.
e. Learners choose underlying forms that allow for the simplest and most regular grammar.
f. Where possible, rules and representations formulated entirely in phonological terms are preferred to those that mention non-phonological (e.g., morphosyntactic, lexical or diacritic) terms.
g. Phonological processes take the form of ordered rules.

Assumptions (5a) and (5b) are common to most theories of phonology. The status of (5c) was at the core of a much-discussed example of what Bloch (1941) called the PARTIAL OVERLAPPING OF PHONEMES. This situation arises when occurrences of the same sound under the same conditions must be assigned to different phonemes. An example involves the word-final devoicing (WFD) of obstruents in German (Trubetzkoy 1939; Moulton 1962), as shown in (6).

(6) *Word-final devoicing in German*

a. 'federation'

	Singular	Plural
UR	/bund#/	/bund+ə/
WFD	bunt	—
SR	[bunt]	[bundə]

b. 'colourful.SG'

	Masculine	Feminine
UR	/bunt#/	/bunt+ə/
WFD	—	—
SR	[bunt]	[buntə]

In (6a), WFD applies to the underlying /d/ of *Bund* /bund#/ 'federation.SG', causing it to surface as [t], like the [t] in (6b) *bunt* 'colourful.M.SG', which derives from an underlying /t/.[4] Thus, two occurrences of phonetic [t] that derive from different phonemes occur in identical phonetic contexts. If we are allowed to take into account the lexical identities of these forms and the evidence of other members of their paradigms, we can see that the final [t] of 'federation' alternates with [d], as in [bundə], whereas the final [t] in 'colourful' is always [t], as in [buntə]. Bloch, like many of his fellow post-Bloomfieldians, assumed that learners may not make use of such evidence when establishing their phonemic systems, and would therefore have no way to know that the final [t]s in these words belong to different phonemes. Therefore, Bloch (1941) proposed that this kind of overlapping should be excluded from phonemic analysis. It then follows that the [t] in both instances of *bunt* would have to be assigned to the phoneme /t/; the relation between [t] and [d] in the forms of 'federation' would have to be accounted for at a higher, morphophonemic, level of analysis.

Halle (1959) showed that this restriction on phonemic analysis results in a loss of a generalisation when applied to Russian; Chomsky (1964) argued that the assumptions undergirding this restriction have no empirical support. Their arguments thus undermined the empirical and psychological bases of what they called the 'taxonomic phoneme' (see further Dresher & Hall 2022; Ladd 2022).

This example shows how learnability considerations have always been at the heart of debates over abstractness in phonology. It also illustrates how assumptions about what resources learners bring to the task of acquiring phonology govern our estimate of how easy or hard underlying forms are to acquire. If learners do not have access to the facts that [bunt] 'federation.SG' is related to [bundə] 'federation-PL' and that [bunt] 'colourful.M.SG' is related to [buntə] 'colourful-F.SG', then it is impossible for them to correctly assign the [t] of any given instance of [bunt] to /d/ or to /t/; if they do have access to this

[4]For the sake of this discussion, I assume that WFD results in the complete neutralisation of /d/ and /t/. There have been arguments that the neutralisations in many cases of this kind are only partial (see Kulikov 2012 for Russian); Charles-Luce (1985) finds, however, that the word-final voicing contrast in German is completely neutralised in some contexts.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

information, then it becomes easy to imagine how they might do this. Following Chomsky and Halle, most analyses in a generative framework have assumed the correctness of (5c).[5]

Assumption (5d) has been assumed by many, but not all, phonological theories, particularly within generative phonology; see Krämer (2012) for a history of underlying representations. For example, the grammar in (6) posits that the morpheme 'federation' has a single underlying representation /bund/. If we deny (5d), we could assign it two underlying representations, /bunt/ for when the stem is word final and /bund/ for when it is followed by a suffix. But such an analysis fails to capture the transparent generalisation that both [bunt#] and [bund-] follow predictably from /bund/ and WFD. Hyman (2018: 595), who argues in support of this assumption, calls it the MORPHOPHONEMIC PRINCIPLE: 'one should derive allomorphs from the same UR, wherever possible and "motivated".' Sometimes, no plausible phonological account can be constructed, as in cases of suppletive allomorphy, and learners may be unable to arrive at a single underlying form. Bermúdez-Otero (2018a) argues that underlying representations adhering to (5d) are required in order to distinguish between regular phonology and suppletive allomorphy.[6] Sanders (2003) and O'Neill (2024), who oppose (5d), call it the PRINCIPLE OF LEXICAL MINIMISATION. It is clear from all the above writers that the utility and validity of this principle depend on what other machinery one attributes to UG. Sanders (2003) and O'Neill (2024) posit different theories of phonology and morphology from the one adopted here.

Assumption (5e) is standard in theories that adopt (5d). Since the evaluation measure proposed by Chomsky & Halle (1968), what 'the simplest and most regular grammar' is and how learners arrive at it has been a central issue in phonological learnability (see the introduction by Katzir *et al.* 2021 to a special issue of the *Journal of Language Modelling* on simplicity in grammar learning and the papers therein). My focus is to show how learners guided by the proposed UG principles could plausibly arrive at abstract representations and opaque rules in the cases to be discussed. See §2.3 for some proposals for how simplicity might be computed by learners in a general way.

Assumption (5f) expresses the idea that phonological analyses preferably employ phonological terms rather than non-phonological terms. Of course, there are morphologically governed alternations that may look phonological but which do not have a purely phonological solution. However, it may not be evident to learners at first what sort of case they are dealing with. Therefore, even in such cases, logic requires learners to *first* look for a phonological solution, and only resort to non-phonological terms when none is available. Otherwise, they risk overlooking a phonological solution when there is one, contrary to (5f).

Diacritics are here considered to be non-phonological by definition; in practice, however, whether a term is classified as phonological or diacritic depends on one's theory. For example, Nespor & Vogel (2007) consider the prosodic hierarchy and its constituents (the phonological phrase *φ,* the intonational phrase *I,* etc.) to be native to the phonology, but Scheer (2008) considers them to be diacritics. For our purposes, I consider a 'diacritic' to be an *ad hoc* symbol that has no phonological/phonetic interpretation in itself, but which has phonological consequences. For example, if one were to distinguish two types of *i* in a language by designating them i_1 and i_2, then the numerical subscripts are diacritics. I understand 'the diacritic use of phonological features' (Kiparsky 1973) to mean that features are being used in an *ad hoc,* arbitrary way that has no connection to their purported phonological/phonetic interpretation. Conversely, 'the phonological use of diacritic features' (Kiparsky 1973: 18) refers to diacritics that function as phonological features; according to (5f), learners would posit actual phonological features rather than diacritics that have the same effect.

Finally, I assume that underlying representations are related to surface representations by means of ordered rules (5g), as in classical generative grammar. I adopt this model for concreteness and do not intend to exclude refinements such as Autosegmental Phonology (Goldsmith 1976), Metrical Phonology (Liberman & Prince 1977; Halle & Vergnaud 1987) or levels as in Lexical Phonology and Morphology (Kiparsky 1982a,b, 1985). Many phonologists assume rather that phonological processes

[5] What items count as part of a 'paradigm' is an important question that I will not pursue here.

[6] Where exactly the boundary between regular phonology and suppletive allomorphy is remains a subject of investigation.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

take the form of ranked constraints, as in various versions of Optimality Theory (OT; Prince & Smolensky [1993] 2004). Some of the basic assumptions carry over to versions of OT that allow for synchronic opaque interactions, including Stratal OT (Kiparsky 2000; Rubach 2003; Bermúdez-Otero 2018b; Jaker & Kiparsky 2020) and Harmonic Serialism (McCarthy 2007, 2008; Elfner 2016; Rasin 2022; Pruitt 2023).[7]

There are differing views as to the nature of the basic primes that constitute phonemes. What form they take is an empirical issue; for concreteness, I will assume that they are binary features as in (7). Privative (unary) elements or structures have also been proposed; see van 't Veer *et al.* (2023) for a recent survey. Many of the learnability considerations discussed here carry over to these other representations.

(7) *Basic assumptions about phonological primes*
   a. The phonological primes are FEATURES (Halle 1962).
   b. Features are binary and language-particular.
   c. Each feature has a MARKED and an UNMARKED value, determined on a language-specific basis (Rice 2003, 2007). By convention, the marked value of a feature F is assigned to [+F], and the unmarked value to [-F].
   d. On a language-specific basis, either both values of a feature may be active, or marked values may be more active than unmarked values, which can serve as defaults and may be more or less inert.

By 'language-particular' (7b), I mean that I accept the arguments by Mielke (2008), Samuels (2011) and others that features are not universal, but are constructed by learners based on the evidence supplied by their language. Positing that features are EMERGENT in this sense does not lessen the need for a rich set of innate UG principles; on the contrary, if individual features are not innately supplied it is all the more important to have principles that constrain how many features there can be and what properties they have. I assume that perceptual and physical systems put constraints on what a possible feature is (Heffner *et al.* 2019). I also assume that constraints are imposed by considerations of phonological contrast (Cowper & Hall 2014; Dresher 2014; Hall 2023), as discussed further below.

It should be emphasised that the kind of markedness assumed in (7) is language-particular, in the spirit of Trubetzkoy (1939); see further Rice (2003, 2007), Battistella (2022: 229–231) and Dresher & Hall (2022: 373–374). On this conception, just as each feature has a positive and a negative pole, it also has a marked and an unmarked one. This kind of markedness is a formal property of feature theory and is therefore not subject to the critiques that markedness is a prime example of 'substance abuse' (Hale & Reiss 2000; Reiss 2018) or vagueness (Haspelmath 2006; Martins 2017). Assumption (7d) is based on the observation by Rice (2007) that susceptibility to assimilation or deletion is characteristic of unmarked elements, which tend to be targets of phonological processes, whereas segments with marked features tend to be triggers of processes.

I further assume that features are assigned in accord with Contrastive Hierarchy Theory (CHT), whose main tenets are shown in (8):

(8) *Contrastive Hierarchy Theory* (Dresher *et al.* 1994; Dresher 2009)
   a. *The Successive Division Algorithm* (SDA; Dresher 1998, 2003, 2009): Contrastive features are assigned by successively dividing the inventory until every phoneme has been distinguished.
   b. *Variability of feature ordering*: Features and feature ordering are language-particular and thus can vary over space and time.
   c. *The Contrastivist Hypothesis* (Hall 2007): The phonological component of a language L operates only on the contrastive features derived by the SDA.

[7] Jurgec & Wang (in press) show that parallel OT with extended input faithfulness constraints can account for certain types of opaque interactions.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

These principles of CHT constrain the number of features that can be posited for any given inventory (Dresher 2018). Since the ordering of features is language-particular, learners need a way to determine which features are contrastive and how they are ordered. According to (8c), only contrastive features can be active; therefore, by hypothesis, an active feature must be contrastive. In CHT, then, an important source of evidence for learners is PHONOLOGICAL ACTIVITY, which can be defined as in (9).

(9) *Phonological activity* (adapted from Clements 2001: 77)
A feature can be said to be ACTIVE if it plays a role in the phonological computation; that is, if it is required for the expression of phonological regularities in a language, including both static phonotactic patterns and patterns of alternation.

CHT together with the assumptions in (5) and (7) posits that phonology is serial and rule-based, and that only contrastive features are phonologically active. The analyses I will present presuppose this framework, but it is possible that their essential elements can also arise in other theories.

Dresher (1981) considers how learners might arrive at abstract grammars. He suggests that 'the most common phonetic rules involve the assimilation of one feature to a feature in its environment' and proposes that a rule of the general form in (10) is a highly-valued rule that learners would be drawn to construct.

(10) *Basic template for assimilation* (Dresher 1981)

$$\begin{bmatrix} \alpha_1 F_1 \\ \vdots \\ \alpha_n F_n \end{bmatrix} \rightarrow [\alpha_j F_j] \Big/ \begin{array}{c} __ \text{ X } [\alpha_j F_j] \\ \textit{or} \\ [\alpha_j F_j] \text{ X } __ \end{array}$$

According to (10), if a segment $S = [\alpha_1 F_1 \ldots \alpha_n F_n]$ takes on a certain feature value $[\alpha_j F_j]$ in the presence of another segment T, the learner will suppose, unless there is evidence to the contrary, that T also bears the same specification and is the cause of the change in S. A similar (perhaps too strong) constraint has been proposed by Danesi (2022) under the name of the No Ex Nihilo Hypothesis:

(11) *No Ex Nihilo Hypothesis* (Danesi 2022: 192)
Phonological computation cannot manipulate primes that are absent from the representation of the target and the trigger.

2.3. Relationship to learning models

The principles presented in §2.2 are not intended to provide a step-by-step acquisition procedure or computational algorithm. As stated in §2.1, they are intended to address the logical problem of language acquisition; in particular, they are an attempt to provide a set of learning principles that can lead a learner to the abstract representations that I argue for in the subsequent sections. Thus, I abstract away from early or intermediate stages of learning and look at what crucial principles of UG a learner supplied with all the relevant data would need to arrive at the proposed grammars.

Various models have been presented that aim at aspects of learning that the principles in §2.2 do not address (see Jarosz 2019 and Belth 2023a for reviews). Some of these models are claimed to be able to learn some types of opaque rules and/or abstract underlying representations. To the extent that they do not conflict with the principles assumed here, they can be viewed as complementing the current proposal.

In a series of publications (Rasin & Katzir 2016, 2020; Rasin *et al.* 2018, 2021; Rasin 2022), Roni Katzir, Ezer Rasin and their collaborators have argued in favour of a learning approach that follows the principle of minimum description length (MDL; Rissanen 1978; Rissanen & Ristad 1992). They

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

show that it is capable of learning at least some types of opacity in both constraint-based and rule-based phonology.[8]

MDL enforces a type of locality in phonology, and locality is the main principle of the Parsimonious Local Phonology (PLP) learner proposed by Belth (2023a). He shows that PLP can learn abstract underlying forms and certain opaque rules. This approach complements principles such as (10) and (11), which focus on the relation between the target and the context of a rule; PLP focuses on the locality of the context (the 'X' in (10)), but not on the relation between the target and the context. Belth (2026) extends the approach to long-distance tier-based interactions (cf. Nevins's [2010] notion of relativised locality).

Yang (2005, 2016) has proposed the Tolerance Principle, which posits the maximum number of exceptions that a rule can withstand before a learner is required to revise or abandon it. Belth (2023a,b) incorporates it into his learning model, and nothing prevents it from being a part of the UG outlined in §2.2.

Chandlee *et al.* (2018) characterise opaque interactions computationally in terms of input strictly local (ISL) maps (Chandlee 2014). They argue that 'the various types of interactions which researchers have identified as opaque all share the computational property of being ISL', which comes with strong learnability results (Chandlee *et al.* 2018: 172). They conclude that such opaque interactions 'are not necessarily more complex than single processes'.

Computational learning models within OT have been shown to be able to learn some types of opaque interactions. One influential approach uses Maximum Entropy Grammar (Goldwater & Johnson 2003; Hayes & Wilson 2008); see O'Hara (2017), Nazarov & Pater (2017) and Wang & Hayes (in press) for applications.

Though the models mentioned in this section may differ conceptually and empirically from my approach, the fact that different types of learning models succeed in learning opaque interactions, when supplied with appropriate principles, shows that opacity is not the stumbling block to learnability that it has often been supposed to be.

In the following sections, I will show how the principles in §2.2 apply to several cases of abstractness. In each case, I will argue that abstractness and opacity pose no special problems that would require the learning theory to depart from the normal phonological principles that apply to transparent rules. Moreover, I will contend that the proposed alternative non-opaque analyses are not any easier to learn, and that learnability considerations do not militate against abstract phonological representations and opaque rules in these cases.

3. An abstract underlying vowel in Inuit dialects

3.1. Strong and weak [i]

Many Inuit dialects make a distinction between 'strong *i*' and 'weak *i*'.[9] For example, in Barrow (North Alaskan) Inupiaq (Kaplan 1981b: 81–82), the suffixes *-lu* 'and a N', *-nik* 'OBL.PL', and *-tun* 'like a N' follow a stem whose last vowel is *u*, as in (12a).

(12) *Barrow Inupiaq palatalisation after strong* i (Kaplan 1981b)

	Stem	*Gloss*	'and a N'	'N-OBL.PL'	'like a N'
a.	iɣlu	'house'	iɣlu-lu	iɣlu-nik	iɣlu-tun

[8]The application of MDL to grammar may antedate the cited references. According to Kiparsky (2022: 39), the Sanskrit grammar of Pāṇini (*ca.* 500 BCE?) 'seeks solely to extract all grammatical regularities, rigorously guided by the twin imperatives of complete coverage and the principle of Minimum Description Length'.

[9]I use the term 'Inuit dialects' to refer to all languages in Alaska, Canada and Greenland that descend from the Inuit branch of PIY.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

b.	iki	‘wound’	iki-ʎu	iki-ɲik	iki-sun
	saviɣ	‘knife’	saviɣ-ʎu	saviɣ-ɲik	savik-sun
	qimmiq	‘dog’	qimmiʁ-ʎu	qimmi-ɲik	qimmi-sun
c.	ini	‘place’	ini-lu	ini-nik	ini-tun
	kamik	‘boot’	kamig-lu	kamig-nik	kamik-tun
	aiviq	‘walrus’	aiviʁ-lu	aiviʁ-nik	aiviq-tun

These suffixes all begin with an alveolar consonant which is palatalised after some *i*, as shown in (12b): *l* becomes *ʎ*; *n* becomes *ɲ* and *t* becomes *s*, which is the usual Inuit palatalisation (actually assibilation) of *t*. This is called ‘strong *i*’ in the literature. After other *i* (‘weak *i*’) as in (12c), there is no palatalisation, and the suffixes appear as they do after *u* in (12a) (and also after *a*, not shown here). Following Kaplan (1981b), Compton & Dresher (2011) propose that whereas strong *i* is underlyingly /i/, weak *i* derives from an underlying vowel that is distinct from /i/.[10]

This abstract analysis reflects the historical provenance of strong and weak *i* from Proto-Inuit-Yupik (PIY) as reconstructed by Fortescue *et al.* (1994): strong *i* mostly derive from PIY *i, as in (13b), and weak *i* mostly derive from a fourth vowel reconstructed as PIY *ə, as in (13c).[11] Cognates of these words in Central Alaskan Yupik (CAY) dialects, which retain a phonetic contrast between /i/ and /ə/, usually have [i] corresponding to strong *i* and [ə] corresponding to weak *i* (13).[12] Child learners of modern Inuit dialects have no access to the PIY origins of these words or to the CAY cognates. The abstract analysis is not motivated by PIY or by CAY, but by the synchronic data that learners have access to.

(13) *Historical origins of strong and weak* i (Fortescue *et al.* 1994)

	Barrow	*Gloss*	*PIY*	*CAY*
a.	iɣlu	‘house’	*əŋlu	ŋəllu, nəllu
b.	iki	‘wound’	*əki	əkiq
	saviɣ	‘knife’	*caviɣ	tʃavik
	qimmiq	‘dog’	*qikmiʁ	qikmiq
c.	ini	‘place’	*ənə	ənа (base ənə)
	kamik	‘boot’	*kaməɣ	kamək-
	aiviq	‘walrus’	*ajvəʁ	azvəq

3.2. A learning theory for the abstract vowel in Inuit dialects

How hard is it, in Inuit dialects, to acquire an underlying vowel that has no distinct surface realisation? It depends on the contents of UG. If UG limits the learner to underlying representations that exist as such at the surface, then there is no path to acquiring an abstract vowel. But I know of no evidence for this limitation. It would perhaps have some rationale if UG treated phonemes as undecomposable primes. If that were the case, one could perhaps argue that it would be difficult, if not impossible, for learners exposed only to allophones of /i/, /u/, and /a/ to posit a fourth phoneme, say /ə/, in lexical representations. But the theory adopted here assumes that phonemes are composed of features, as in (7). If so, then it is these primes that are the material of phonological computation. I will argue that *this* kind of computation can easily lead to abstract phonemes.

[10] I follow the literature in assuming that there is full surface identity between weak *i* and strong *i*, or, in some contexts and dialects, between weak *i* and /a/ or /u/ (see further §3.4). Kaplan (1981b) assumes this to be the case for North Alaskan Inupiaq. I do not know of a phonetic study of that dialect, but phonetic studies of West Greenlandic, which also distinguishes between strong and weak *i*, have not found any phonetic differences between them. Wood (1971: 59) writes that the reflexes of abstract /ə/ ‘have no separate identity in speech’, a view supported by Hagerup (2011: 5).

[11] For PIY, Fortescue *et al.* (1994) use the term Proto-Eskimo. I have changed their ⟨ʁ⟩ to IPA ⟨ʁ⟩ and ⟨y⟩ to ⟨j⟩.

[12] Kaplan (1981b: 85–86) writes that correspondences are not always as expected, because of additional diachronic changes in some dialects. For example, in Inupiaq, some *i* descended from PIY *ə have become palatalising (i.e., strong) *i* (Kaplan 1981a).

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Recall that one of the 'strong claims' of the Compton & Dresher (2011) analysis, according to Mayer *et al.* (2022a: 1), is 'that speakers have localized the source of the idiosyncrasy to a distinction between two sounds that are identical on the surface'. In this case, it is very clear from (12) that some stem-final *i* are what cause palatalisation of the suffix-initial consonants, which does not occur in their absence.[13] It is not so obvious what the palatalising feature is; for now, let's call it [+P]. Thus, *i* in (12b) carries this feature, and stem-final *u* in (12a) and *i* in (12c) do not. Further, by our assumption in (10), learners prefer to posit a rule that assigns [+P] in the context of [+P], suggesting that palatalisation assigns [+P] to the palatalised consonants. Identifying one feature [P] for both cause and effect is simpler than identifying two different features [P] and [Q], if palatalisation is not the result of the assimilation of a feature. Adapting the template in (10) to our case, we suppose that learners would posit a rule like (14) to account for palatalisation.[14]

(14) *Palatalisation of alveolars* (modified from Kaplan 1981b: 85)

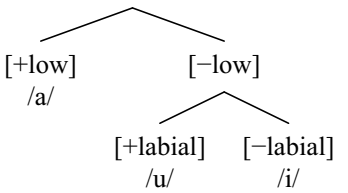
$$\begin{bmatrix} \text{C} \\ \text{+coronal} \\ \text{–continuant} \end{bmatrix} \rightarrow [\text{+P}] \; / \; \begin{matrix} \text{V} \\ [\text{+P}] \end{matrix} \; (\text{C}) \; __$$

Whereas stem-final *u* in (12a) clearly lacks [+P] and the stem-final *i* in (12b) clearly has it, (12c) presents conflicting signals: stem-final *i* is phonetically the same as *i* in (12b), hence [+P]; but the suffixes that follow it suggest that it is not [+P]. The learning theory tells the learner how to resolve this conflict. In classical generative phonology, the resolution occurs via a derivation: though stem-final *i* in (12c) has [+P] at the surface, it must lack it underlyingly and at least until after palatalisation has applied.

3.3. Assigning features to weak i

How does a learner assign features to weak *i*? The same way that features are assigned to every other phoneme. Let us consider first Inuit dialects with three underlying vowels, /i, a, u/. These dialects have completely lost any contrast between PIY *i and *ə. Interestingly, none of these dialects have palatalisation after /i/. Compton & Dresher (2011) propose that this is because the Inuit-Yupik feature hierarchy has [low] ordered over [labial] (notated [low] ≫ [labial]), as shown in (15).[15] As only two contrastive features are needed in a three-vowel system, /i/ in these dialects has no contrastive palatalising feature; it's the unmarked vowel. Compton & Dresher (2011: 223) observe that /i/ in these dialects is the epenthetic vowel inserted to satisfy phonotactic constraints in adapting loanwords: 'This use of /i/ as the default vowel supports its unmarked status in the inventory of three-vowel dialects.'

(15) *Contrastive hierarchy for three-vowel Inuit dialects:* [low] ≫ [labial]
Compton & Dresher (2011: 223)



[13] In raising the localisation issue, Mayer *et al.* (2022a) may have been influenced by backness harmony in Uyghur, the subject of their own analysis, in which it is not so straightforward to pinpoint the source of the backness feature (see further §4). This is one important way that Inuit dialects differ from Uyghur: Barrow Inupiaq learners would have abundant data that shows that palatalisation is associated with some preceding *i*, and not with consonants or anything else.

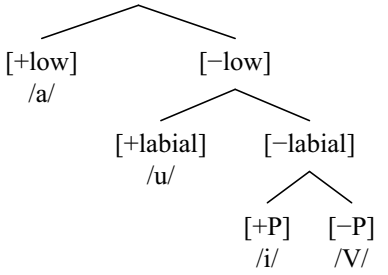
[14] Kaplan (1981b) considers /l/ and /n/ to be [+continuant], which makes the rule more complicated. I assume that they are [−continuant] in Inupiaq (see Mielke 2008: 56–77 on language-specific values of [continuant] for laterals and nasals). As shown in (12b), a stem-final consonant can intervene between strong *i* and the suffix, hence the (C) in (14).

[15] I have changed Compton & Dresher's feature notation from [F] ~ (non-F) to [+F] ~ [−F], where the first term in each represents the marked value and the second term the unmarked value.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Now consider dialects which have retained a contrast between strong and weak *i*, that is, between /i/ and a fourth vowel, which for now we can call /V/. By assumption (5f), the contrast between /i/ and the fourth vowel /V/ requires a phonological feature, which must be the palatalising feature [P]. Several considerations converge to the conclusion that this feature should be added to the hierarchy as shown in (16).

(16) *Contrastive hierarchy for four-vowel Inuit dialects:* [low] ≫ [labial] ≫ [P]
Compton & Dresher (2011: 221)



In support of (16), the fourth vowel merges with /i/ (in most cases; some other outcomes are discussed in the next section), suggesting that its features apart from [P] are the same as those of /i/. Second, there is no evidence that the fourth vowel is either [+low] or [+labial]. Third, we will see evidence that the fourth vowel is unmarked in these four-vowel dialects, supporting its lack of marked vowel features.

What is [P]? Compton & Dresher (2011) proposed that [P] is [coronal], following studies that argue that V-place [coronal] causes palatalisation of consonants (Clements 1976, 1991a; Hume 1994), even those with C-place [coronal] (see Kochetov 2011 for a review). If strong /i/ is [+coronal], then weak *i* must be [−low, −labial, −coronal], that is, a non-low, non-rounded, non-front vowel we could call /ə/ or /ɨ/.

Kaplan (1981b) followed Chomsky & Halle (1968) in proposing that the palatalising feature is [+high]. This view is supported by Lahiri & Evers (1991) and Lahiri (2018), who argue that the palatalisation of /l/ to [ʎ] and /n/ to [ɲ] amounts to the change of [−high] to [+high].[16] Suppose we assume that the third feature is [high]. On this view, weak /i/ must be [−low, −labial, −high], a mid non-rounded vowel which again could be considered to be /ə/ (or perhaps /e/).

In our case, it doesn't much matter what we call the palatalising feature. Whether feature theory has C-place and V-place tiers is determined by UG; it's not something the learner has to figure out. Nor does it matter what exactly we call the fourth vowel. In CHT, segments are underspecified in the phonological component, and phonetic detail is added in the phonetics; but the fourth vowel has no independent phonetic realisation. There is no harm in identifying it with a particular segment like [ə] or [ɨ] or [e], but we should keep in mind that its actual identity is the set of contrastive features it is assigned, which strictly speaking have no phonetic realisation (being underspecified).[17]

We can now address another critique of the abstract analysis by Mayer *et al.* (2022a: 1), that it is an instance of 'the diacritic use of phonological features'. As discussed in §2.2, I understand this expression to refer to the use of phonological features in an *ad hoc*, non-phonological manner. However, there is nothing *ad hoc* or arbitrary about the features employed in this analysis: rather, phonological features are motivated by phonological processes and contrasts.

[16] As to /t/ to [s], Kaplan (1981b) proposes that adding [+high] to /t/ gives [tʃ], which then becomes [s]. The palatalisation of /t/ is [s] in many Inuit dialects, presumably by a similar mechanism.

[17] The assumption that underlying phonemes are underspecified may simplify the procedure for assigning them features, but it is not crucial to the argument that abstract phonemes are learnable. If, contrary to CHT, all underlying phonemes had to be fully specified for some set of features, I assume that UG would provide a mechanism for doing so.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

3.4. Weak i *and opacity*

I have now presented the outline of a learning theory that can lead an Inuit learner to identify weak *i* with an underlying set of features that can be written /ə/ (with the above caveats). To complete the story, let's consider sample derivations with strong and weak *i*. In (17a), palatalisation applies to /l/ that follows underlying (strong) /i/. In (17b), palatalisation does not apply to /l/ that follows underlying /ə/ (weak *i*). The neutralization of /ə/ to [i] must follow the application of palatalisation. The relative ordering of palatalisation and /ə/ → [i] makes the former rule opaque.

(17) *Sample derivations with strong and weak* i

a. 'and a wound'

UR	/iki+lu/
Palatalisation	ikiʎu
/ə/ → [i]	—
SR	[ikiʎu]

b. 'and a place'

UR	/inə+lu/
Palatalisation	—
/ə/ → [i]	inilu
SR	[inilu]

The examples in (12c) make palatalisation an example of the type of opacity defined in (1a): palatalisation appears to have underapplied in (17b) in *ini-lu* (**ini-ʎu*), and similarly in *ini-nik* (**ini-ɲik*) and *ini-tun* (**ini-sun*). The underapplication is only apparent, however; at the point in the derivation where palatalisation applies, the relevant forms are /inə-lu/, /inə-nik/ and /inə-tun/, and the context for palatalisation is not met.

Does this opacity thereby make the palatalisation rule hard to learn? No! In our learning scenario, learners have *already* acquired palatalisation (14), as well as the underlying contrast between /i/ and /ə/. Ordering palatalisation before /ə/ → [i] – that is, creating opacity – is a *solution* to the problem of conflicting signals sent by weak *i*.

Even if it is agreed that the abstract analysis is learnable, one might ask whether it should be preferred to a diacritical analysis. What would be wrong with positing a single /i/ in palatalising dialects, and marking the so-called weak *i* as exceptionally not triggering palatalisation? Or alternatively, marking palatalising /i/ with a diacritic, say /iP/, that triggers palatalisation? Is it not somehow simpler, more conservative, closer to the null hypothesis, to posit a diacritic than an abstract phoneme?

No! The impression that a diacritic solution is the default option is an illusion based on thinking of phonological computation as accessible to conscious introspection. In our case, we have an example of palatalisation of consonants triggered by an *i*, a cross-linguistically very common process that is uncontroversially mediated by a phonological feature. As this account is available in our case – indeed, preferred, given the learning assumptions we are following – it follows that the burden of proof is rather on those who wish to argue that this approach does not apply here.

Compton & Dresher (2011) survey different Inuit dialects that distinguish strong and weak *i*. In some dialects, palatalisation is no longer productive and is limited only to the change of /t/ to [s]. In such dialects, the explanatory work done by positing a fourth vowel is limited, whatever analysis we adopt. Other dialects display more robust evidence for a fourth vowel, and in those dialects there are indeed significant differences between these approaches. Thus, Kaplan (1981b: 31) writes that although all Inupiaq dialects have a phonetic three-vowel system, 'there are strong arguments' based on phonological processes for the underlying existence of a fourth vowel.

Learners of North Alaskan Inupiaq would have several reasons in addition to palatalisation to distinguish two types of [i] (Kaplan 1981b: 116–149). Unlike strong *i*, weak *i* (which Kaplan transcribes as *ï*) alternates with *a* in certain contexts, as before the copula *-u-* in (18b); compare the strong *i* in (18a), which does not alternate. In other contexts, weak *i* but not strong *i* alternates with *u*, as well as with ∅; the latter alternation is evidence for either epenthesis or syncope, depending on the analysis.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Table 1. *Uyghur vowel inventory (Hall & Ozburn 2019).*

	Front		Back	
	Unrounded	Rounded	Unrounded	Rounded
High	i	y		u
Mid	e	ø		o
Low	æ		ɑ	

(18) *Barrow Inupiaq strong and weak* i *before a vowel*

a. *Strong* i*: iki-*

/iki -u- -tuq/ → [ikiuɹuq], *[ikauɹuq]
wound -COP- -PART.3SG

'It is a wound.'

b. *Weak* i*: aŋguti-*

/aŋutə -u- -tuq/ → [aŋutauɹuq], *[aŋutiuɹuq]
man -COP- -PART.3SG

'It is a man.'

Now the advantages of an abstract (i.e., phonological) analysis over a diacritic one are more evident. The various kinds of phonological activity that the two *i* enter into are all consistent with strong *i* being more marked than weak *i*: palatalisation suggests that strong *i* has an active feature that weak *i* lacks; and the fact that weak *i* is easily assimilated and either deleted or epenthetic suggests that its features are relatively unmarked, as per (7d). Indeed, in Yupik, which retains an overt four-vowel system /i, a, u, ə/, the schwa does not have the same status as the other vowels. According to Kaplan (1990: 147), it 'cannot occur long or in a cluster with another vowel', instead undergoing dissimilation or assimilation when adjacent to full vowels. Miyaoka (2012: 38) writes that Yupik schwa 'frequently occurs as an epenthetic vowel (EV) to break up an unallowable consonant cluster'. This is exactly the behaviour of weak *i* in Inupiaq, which follows from the abstract analysis. On a diacritic analysis, however, there is no reason to expect that a non-palatalising *i* should also be subject to assimilation and deletion or epenthesis, unless we ascribe to the diacritic the properties of an unmarked phonological feature, essentially adopting the abstract analysis.

4. Abstract /ɯ/ in Uyghur?

The critique by Mayer *et al.* (2022a) of the abstract analysis of weak *i* in Inuit dialects is in an article on backness harmony in Uyghur. Mayer *et al.* (2022a) are opposed to what they call 'the traditional analysis' of Uyghur (Lindblad 1990; Hahn 1991a,b), which posits a covert contrast between front /i/, which appears on the surface, and an abstract unobserved back /ɯ/. I agree that this analysis is not correct for Uyghur, but not because all such abstract analyses are incorrect. Rather, the evidence available to learners is crucially different in the two cases, and therefore, learners of Uyghur will arrive at a different analysis. In this section, I briefly review Uyghur backness harmony with a view to showing why it requires a different, though still phonological, solution.

4.1. Backness harmony in Uyghur

The vowel inventory of Uyghur is shown in Table 1.[18]

[18] See Hall & Ozburn (2019) and Dresher *et al.* (2024) for further discussion of the general approach I follow here.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Uyghur displays backness harmony, whereby the vowels of suffixes harmonise with the final vowel of a stem, as in (19) and (20). Some consonants also participate in this harmony: the dative front *-gæ* (or *-kæ*) in (19c) alternates with back *-qa* (or *-ʁa*) in (20c).

(19) *Front harmonising stems*

a. yn-dyr ‘appear-CAUS’
b. køl-lær ‘lake-PL’
c. øj-gæ ‘home-DAT’

(20) *Back harmonising stems*

a. sun-dur ‘break-CAUS’
b. jol-lɑr ‘road-PL’
c. ot-qɑ ‘fire-DAT’

The vowels *i* and *e* have no back counterparts (**ɯ, *ɤ*). They can co-occur with both front and back vowels, and transparently allow backness harmony to affect following suffixes, as in (21) and (22).[19]

(21) *Front harmonising stems*

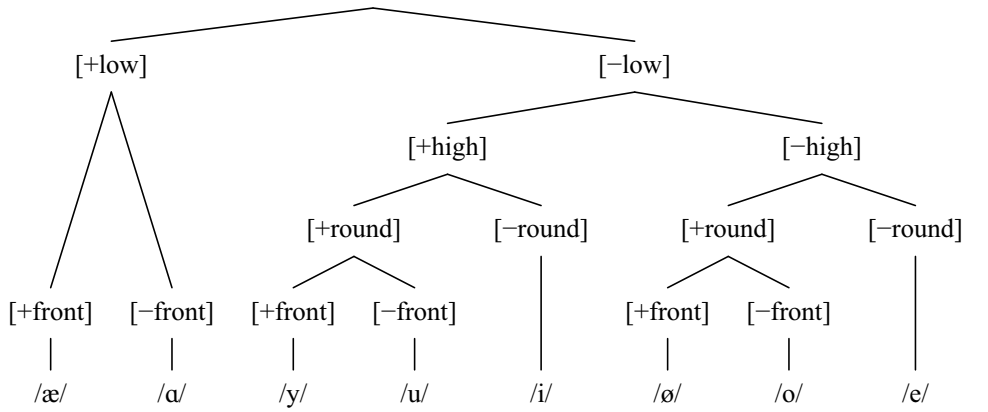
a. mømin-gæ ‘believer-DAT’
b. jyz-imiz-gæ ‘face-our-DAT’

(22) *Back harmonising stems*

a. ɑmil-ʁɑ ‘element-DAT’
b. jol-imiz-ʁɑ ‘road-our-DAT’

Given the data up to here, when Uyghur [i] co-occurs with back vowels, one might suppose that it is the surface realisation of an abstract back /ɯ/. In Uyghur, however, there is another way to interpret the varying signals sent by surface [i]: rather than representing two underlying vowels, one could also conclude that there is a single /i/ that is neutral with respect to backness harmony. This interpretation is supported by the fact that, unlike the other front vowels in Uyghur, /i/ and /e/ lack a back counterpart, raising the possibility that they have no contrastive backness feature at all. If we generate contrastive features hierarchically, we find that there are many feature orderings in which /i/ and /e/ do not receive a specification for [front/back]. An example is (23).[20]

(23) *An Uyghur vowel feature hierarchy:* [low] ≫[high] ≫ [round] ≫ [front]



On this interpretation, which is not available in Inuit dialects, the harmonic [front] feature originates from the non-neutral stem vowels that bear a contrastive [front] feature. However, there is more to it than that.

Stems containing only /i/ and /e/ are of two types: some take front suffixes, as in (24), and some take back suffixes, as in (25).

[19]The vowel /e/ has a restricted distribution.

[20]A more detailed analysis is required to determine what the feature hierarchy of Uyghur actually is. Note that the layout of Table 1, in which the front–back distinction is superordinate to the unrounded–rounded distinction, is not consistent with any analysis in which /i, e/ are neutral (unspecified) with respect to backness.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

(24) *Front harmonising stems with [i]*

a. tʃekin-dyr ‘retreat-CAUS’

b. bilim-gæ ‘knowledge-DAT’

(25) *Back harmonising stems with [i]*

a. siʁiʃ-dur ‘contain-CAUS’

b. sinip-qɑ ‘classroom-DAT’

These examples cast doubt on the analysis whereby /i/ is neutral, and appear to favour the abstract analysis that posits an underlying /ɯ/ in (25). Additional data, however, puts these forms into a wider context which shows other types of disharmony between stems and suffixes. Mayer *et al.* (2022a) and Mayer *et al.* (2022b) adduce several cases, some of which are shown in (26). These include: (26a) disharmony with *k*, a [+front] consonant that is expected to take *-lær* (cf. (19c)), and (26b) the converse disharmony with [−front] *q*, which should take *-lar* (cf. (20c)); (26c) disharmony following the non-neutral vowel *a*; and (26d) variation in the suffix following *sowet*.

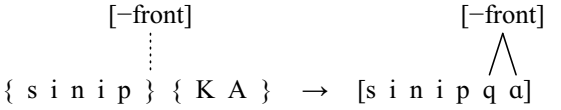
(26) *Other examples of Uyghur stem–suffix disharmony*

a.	etnik-lar	‘ethnic.groups-PL’	front *k* with back suffix
b.	tætqiq-lær	‘research-PL’	back *q* with front vowels
c.	tɑʃ-tæk	‘stone-like’	invariant front suffix
d.	sowet-lær/lɑr	‘Soviet-PL’	varying suffix

These examples show that the source of a harmonizing [±front] feature may be associated with a stem and not with a specific vowel or consonant. That is, the feature may be FLOATING, or on a separate autosegmental tier. In the analysis of Hall & Ozburn (2019), the contrastive feature on non-neutral vowels and consonants is privative PALATAL that is a dependent of V-Place; lacking that feature, they surface as back by default. Similarly, Mayer (2021: 121) proposes that [+back] is unmarked in Uyghur. Dresher *et al.* (2024) use a formalisation in terms of phonological events, features, and precedence drawn from Raimy (2000), Papillon (2020) and Idsardi (2022).

Here, I will assume that a contrastive autosegmental [±front] feature associates with vowels and consonants that are contrastive for this feature and can also be lexically associated in floating fashion with the end of a stem, as in the example in (27). The vowels of the stem *sinip* have no contrastive [front] feature, and the end of the stem has a floating [−front] feature that associates to the vowel and back consonant of the suffix (K and A represent a back consonant and low vowel unspecified for backness; braces represent morpheme boundaries), yielding *qa*.

(27) *Uyghur floating feature: sinip-qa* ‘classroom-DAT’



4.2. Learning exceptional harmony in Uyghur

In comparison to some other examples of abstractness, there is nothing ‘abstract’ about the harmonising feature [±front], which is observed in the speech signal. Nor are there any opaque processes in the data surveyed in the previous section: suffixes transparently receive a [+front] or [−front] feature in the presence of certain stems.[21] Rather, the challenge for learners, as Mayer *et al.*’s (2022a) remark suggests, is to localise the source of the [±front] feature that is affecting the suffixes. In examples like those in (19) and (20), it is clear that the source of the feature is in the vowel of the stem, but in other examples, the source is not so clear. In examples like (25), it is apparent that the source *cannot* be the stem vowels, which seem to have the opposite value of the feature affecting the suffix. So where is the feature coming from?

[21] Uyghur also exhibits a much-discussed process whereby /æ/ and /ɑ/ raise or reduce to [i]. This process does involve opacity, but I cannot pursue it here (see the cited references for discussion and analyses).

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Learners following the learning theory outlined in §2 would see that the backness of alternating suffixes is dependent on the stems they attach to. According to the template in (10), they would assume that the [±front] feature observed on the suffix originates in the preceding context. Where there is an overt vowel or consonant with the harmonising feature, they would locate the source there. If there is no such overt segment, and absent evidence that such a segment has been deleted or altered, the next hypothesis would be that the feature is associated with the stem itself. In the version of the theory assumed here, this association takes the form of a floating [±front] feature. Invariant suffixes like *-tæk* in (26c) have their own prespecified [±front] feature, which can spread to following suffixes (Mayer 2021: 188).

Mayer *et al.* (2022a), following Mayer (2021), propose an OT analysis using lexical diacritics. On this analysis, when the source of harmony is an overt vowel, harmony is governed by the AGREE constraints in (28), a standard OT implementation of feature spreading or agreement. But roots with neutral vowels are marked with an index indicating whether they induce front or back harmony, and these roots are subject to the indexed HARMONISE constraints in (29).[22]

(28) *Constraints driving agreement* (Mayer 2021: 29)

a. VAGREEBACK: Assess one violation if a front suffix is attached to a root whose final harmonising vowel is back.

b. VAGREEFRONT: Assess one violation if a back suffix is attached to a root whose final harmonising vowel is front.

(29) *Constraints driving harmony with neutral vowel roots* (Mayer 2021: 49)

a. HARMONISEBACK$_i$: A harmonising suffix must be back if the immediately preceding morpheme bears the index *i*.

b. HARMONISEFRONT$_i$: A harmonising suffix must be front if the immediately preceding morpheme bears the index *i*.

Mayer (2021: 140) writes:

> A covert contrast analysis requires that learners intuit the existence of a phonemic category that has no clear phonetic correlates, while an analysis using lexically indexed constraints requires only that they learn the idiosyncratic harmonizing behavior of certain roots.

This analysis is an example of 'the phonological use of diacritic features' (Kiparsky 1973: 18), because the diacritic here (the index) is functioning exactly like the phonological feature [±back] or [±front]. Indeed, the index will at some point have to be converted into one of these feature values to actually accomplish the harmony. The learning challenge in the case of neutral and otherwise exceptional morphemes is to learn which stems (or roots) induce what kind of harmony, absent an overt harmonising vowel or consonant. Once learners have met this challenge, the question is how they encode that information in their grammar. That depends on UG.

According to the learning assumptions set out in this article, they would prefer to do so phonologically, using a phonological feature to trigger harmony in exceptional cases just as it does in regular cases. Concern about an abstract */ɯ/ in Uyghur is misplaced: the same learning theory that leads a learner to posit abstract /ə/ in Inuit dialects leads a Uyghur learner to posit that /i/ is not contrastively specified for [±front] and that there are morpheme-level [±front] features. We would need a convincing reason to complicate the learning theory, as well as the phonology, to require two entirely different systems to accomplish backness harmony: one using features for when the source of harmony is overt, and a second using indices for when the source of harmony is not overt. As Dresher *et al.* (2024: 4) conclude, '[t]o all appearances, Uyghur harmony and disharmony are phonological. Lacking evidence to the contrary, we maintain that phonological representations are phonological.'

[22] Presumably, the index *i* has a different value in (29a) than it has in (29b). Mayer *et al.* (2022a) replace (29a) with *FRONTSUFFIX to reflect the markedness of [−back].

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

5. Emergent Grammar *vs.* Universal Grammar in Esimbi

Archangeli & Pulleyblank (2015) argue for 'a largely Emergent Grammar in phonology' and against a particular conception of UG. One of their examples concerns prefix allomorphy in Esimbi, a Bantoid language of Cameroon.[23] They choose this example because it has been the subject of analyses that appeal to abstractness and opaque rule interactions. They suggest that such analyses are problematic because of their opacity and propose an alternative that, in their words, 'resolves the pattern without opacity concerns'.

In this section, I will show that Archangeli & Pulleyblank's (2015) characterisation of Emergent Grammar (EG) is far too weak and is consistent with any conceivable analysis of Esimbi.[24] Turning to their analysis, I will show that, in some respects, it is not very different from the abstract one, and in other respects, it is far less insightful. I will argue that the fact that their analysis has no opacity in the technical sense is of no significance for its relative learnability compared to the abstract analysis, which is learnable by the same principles that apply to the Inuit and Uyghur cases.

5.1. Emergent Grammar and Universal Grammar

As Archangeli & Pulleyblank tell it, the EG hypothesis competes with the UG hypothesis in these terms: Do infants learn grammar 'due to an innate capability specific for language, the Universal Grammar hypothesis (UG), or are they simply the abilities that infants use to learn about all aspects of their world, the Emergent Grammar hypothesis (EG)?' They propose that language learners make use of basic cognitive principles, listed in (30), which are not special to language:

(30) *Principles of Emergent Grammar* (Archangeli & Pulleyblank 2015: 1)
- a. Ability to create categories
- b. Ability to attend to frequency
- c. Ability to generalize and create a symbolic system

That's it! Recall that these principles have to determine what phonological grammars are like; they have to answer the questions in (4). The principles in (30) are indeed fundamental and are presupposed by every learning theory, including the one adopted in this article. But that means that they are compatible with any conceivable theory of phonology. They are thus far too weak to support the acquisition of phonology, which must be based on innate principles that will guide learners to a specific type of phonological grammar. Whatever the exact nature of the phonological grammar, overcoming the poverty of the stimulus requires a contentful theory of UG.

Archangeli & Pulleyblank single out opacity as a problem that UG analyses have and that their EG analysis does not. I will show, however, that opacity is a red herring that obscures the real learning problem.

Esimbi has an unusual restriction on root vowels and an interesting set of prefix alternations. Only three vowels can occur in Esimbi roots, and they are all high: [i, ɨ, u]. Moreover, if a root contains more than one vowel, they must be identical. Esimbi prefixes have more vowels, as shown in (31).

(31) *Esimbi prefix vowels*

i		u
e		o
ɛ		ɔ
	a	

[23] Although Esimbi is often classified as Tivoid, it has no known close relatives (Stallcup 1980b: 139), and its Tivoid affiliation has not been adequately demonstrated (Larry Hyman, p.c.).

[24] See Archangeli & Pulleyblank (2022) for a more detailed account of the EG theory; they do not revisit the Esimbi case there.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

The two prefixes we will be looking at first each have three allomorphs whose choice depends on the root they attach to. Roots with [i] are shown in (32), and roots with [u] are shown in (33); the numerals 3 and 9 represent noun classes.

(32) *Allomorphs of the SG 9 and SG 3 prefixes co-occurring with root vowel [i]*

a. *The SG 9 prefix*		b. *The SG 3 prefix*	
[i-] with 'goat':	i-bi	[u-] with 'end':	u-tili
[e-] with 'antelope':	e-kibi	[o-] with 'tail':	o-ki
[ɛ-] with 'animal':	ɛ-nyimi	[ɔ-] with 'grain':	ɔ-simi

(33) *Allomorphs of the SG 9 and SG 3 prefixes co-occurring with root vowel [u]*

a. *The SG 9 prefix*		b. *The SG 3 prefix*	
[i-] with 'fish':	i-su	[u-] with 'fire':	u-wusu
[e-] with 'bird':	e-nunu	[o-] with 'ear':	o-tu
[ɛ-] with 'hippo':	ɛ-fumu	[ɔ-] with 'hand':	ɔ-bu

The patterns in (32) and (33) are general: prefixes alternate among either [i, e, ɛ] or [u, o, ɔ].[25] The examples show that roots with [i] and [u] come in three types: those that occur with prefixes [i, u]; those that occur with [e, o]; and those that occur with [ɛ, ɔ]. There are also roots with [ɨ], as shown in (34). There are only two types of such roots; prefixes [i, u] are missing.

(34) *Allomorphs of I-, U- prefixes co-occurring with root vowel [ɨ]*

a. *The SG 9 prefix*		b. *The SG 3 prefix*	
[e-] with 'cane rat':	e-bɨ	[o-] with 'spear':	o-tɨ
[ɛ-] with 'place':	ɛ-tlɨ	[ɔ-] with 'broom':	ɔ-bɨ

The generalizations we have observed up to here are:

- A given prefix is always front or always round.
- The height of a prefix depends on the root it attaches to.

At first sight, these prefix alternations appear puzzling. Many languages have height harmony, where the height of a prefix vowel must match the height of the root vowel. But in Esimbi, all root vowels are high!

5.2. The basic learnability argument in Esimbi

Building on work by Kenneth Stallcup (see Stallcup 1980a,b), Hyman (1988) proposes that the restriction of roots to high vowels at the surface masks the fact that in underlying representations there is a wider variety of vowels that trigger height harmony in the prefixes. For root vowels that surface as [i] and [u], the vowels of the prefixes mirror the height of the underlying root vowels. The vowels that surface as [ɨ] evidently derive from two central vowels: a mid vowel that co-occurs with prefixes [e-] and [o-], which Hyman represents as /ə/, and a low vowel that co-occurs with [ɛ-] and [ɔ-], represented as /a/. The ingredients of his analysis are summarised in (35).

[25] There is a third prefix pattern [o, ɛ~ɔ, a] which we will look at in §5.3.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

(35) *The analysis of Hyman (1988)*

a. I- prefixes [i, e, ɛ] are specified [−back]; U- prefixes [u, o, ɔ] are [+round].

b. [i] roots derive from /i/, /e/ or /ɛ/; [u] roots derive from /u/, /o/ or /ɔ/; [ɨ] roots derive from /ə/ or /a/.

c. *Prefix height assimilation (PHA)*: Prefixes assimilate to the height features of the roots.

d. *Root vowel raising (RVR)*: Following PHA, all root vowels become [+high].

Some sample derivations that illustrate this analysis are given in (36). Note that RVR makes PHA opaque: after the root vowels raise, we can no longer see the rationale for why the prefixes have the height that they do.[26]

(36) *Hyman (1988): Sample derivations*

	'a goat'	'a spear'	'a grain'
	SG 9 + *bi*	SG 3 + *tɨ*	SG 3 + *simi*
UR	/[−bk] + bi/	/[+rd] + tə/	/[+rd] + sɛmɛ/
PHA	i-bi	o-tə	ɔ-sɛmɛ
RVR	—	o-tɨ	ɔ-simi
SR	[i-bi]	[o-tɨ]	[ɔ-simi]

Archangeli & Pulleyblank (2015: 10–11) are critical of this analysis; they write:

> Assuming that a phonological difference in the roots is the source of the difference in prefix height requires that height distinctions be encoded in roots even though there is no surface evidence—in the roots—for the required distinction.

The suggestion here – not stated explicitly – is that learners may not have evidence for positing height distinctions in the lexical forms of the roots. Moreover, they observe that the analysis results in surface opacity and propose alternative generalisations that 'resolve the opacity problem', without explaining why opacity is a problem.

Archangeli & Pulleyblank (2015) present an analysis with no opacity and no abstract underlying vowels that are different from the surface vowels. In their analysis, the words for 'goat', 'antelope' and 'animal' (32) all have the same high vowel in their lexical representations: /bi/, /kibi/ and /nyimi/, and similarly for all the other root vowels [u] and [ɨ]. So how do they account for the fact that roots take prefixes of different heights: [i-bi], [e-kibi], [ɛ-nyimi]? Here is how they do it: in the lexical entry of each root, they indicate – in terms of features – what sort of prefix the root prefers. Some sample representations are shown in (37).

(37) *Root lexical representations* (Archangeli & Pulleyblank 2015: 7–9)

a. *Set A*		b. *Set B*		c. *Set C*	
$\{bi_{HI,\ ATR}\}$	'goat'	{ki}	'tail'	$\{simi_{LO,\ RTR}\}$	'grain'
		{tɨ}	'spear'	$\{bi_{LO,\ RTR}\}$	'broom'
$\{sumu_{HI,\ ATR}\}$	'thorn'	{tu}	'ear'	$\{zu_{LO,\ RTR}\}$	'snake'

In (37), the notation indicates that /bi/ 'goat' in Set A prefers a prefix that is [high] and [ATR], and /simi/ 'grain' in Set C prefers a prefix that is [low] and [RTR]. Set B roots have no specified preference; Archangeli & Pulleyblank propose that they take the unmarked prefixes, which in their analysis are [e-] and [o-]. The reasons why prefix selection is framed as preferences, why two features are needed for Sets A and C, and why Set B preferences are left unstated, have to do with the realisations of the third prefix, discussed in §5.3.

Comparing this analysis to one such as Hyman's, Archangeli & Pulleyblank (2015: 10–11) write (emphasis added):

[26] In (36), the root vowels are represented as segments. Determining their features has to wait until we discuss the third prefix.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

> *The issue of surface opacity* […] is a non-issue under this analysis. *The problem* derives from assuming that patterns such as these are entirely phonological. […]
>
> Emergent Grammar recognizes all types of generalizations that the learner might make. Among these are generalizations over sets of lexical items that are arbitrary based on their surface forms […]. *It is the recognition of such lexical generalizations co-existing with phonological generalizations that eliminates opacity as an issue in Esimbi prefix selection.*

Archangeli & Pulleyblank's analysis and discussion raise some questions:

- Why is opacity an issue? What is the issue? What is gained by eliminating it?
- How do we draw the line between phonological and lexical generalisations when the latter are stated entirely in terms of phonological features?
- Are features easier to learn if they are represented as small capital subscripts?

I mention the notation because I think it obscures some similarities between the UG and the EG analyses. The UG analysis of a word like *simi* 'grain', which takes [ɛ, ɔ] prefixes, attributes to the root vowel two different feature specifications: one is underlying, whose height (and tongue root) feature is spread to the prefix, and the other is for the root after raising, as shown in (38a). The EG analysis of this word *also* attributes to the root vowel two different feature specifications: one is for the root vowel and the other is for the prefix, as in (38b).[27]

(38) *UG vs. EG analyses of /simi/*

a.	*UG analysis*	/sɛmɛ/ → [simi]
	Lexical features	[front, non-high, RTR]
	Surface features	[front, high, ATR]
b.	*EG analysis*	/simi/$_{\text{LO, RTR}}$
	Prefix features	[low] and/or [RTR]
	Lexical features	[front, high, ATR]

What does the learner have to discover in order to arrive at either the UG or the EG analysis? The learner has to recognise that, though the root [simi] has vowels that appear as [high, ATR], it carries somewhere, in some fashion, [non-high] and [RTR] specifications. What sort of principles will enable Esimbi learners to discover this? Esimbi learners must have a learning mechanism that directs them to *look for a solution in terms of features associated with the roots to account for the height of the prefixes.*

Consider first the UG analysis: Do we have a learning theory that does this? Yes, we do! In the learning theory outlined in §2.2, the first two assumptions in (5) are that segments are analysed into features and that segments interact through features. Recall also our template for a highly valued rule in (10). Applied to our case, if learners see the form [ɔ-simi] 'SG 3-grain', where the height features of the prefix (but not its roundness) are determined by the root, they will attribute those height features to the root.

Like weak *i* in Inuit dialects, root vowels are sending conflicting signals. In (38a), the surface form of the root vowel [i] in [simi] signals that it is high, but the prefix vowel [ɔ] signals that it is not high. The conflict is resolved by assigning both [high, ATR] and [non-high, RTR] to the root vowel [i]. UG tells the learner how to accommodate these contradictory specifications. In derivational generative phonology, the accommodation takes the form of a derivation with ordered rules, as in (36). [−high, RTR] is assigned to the underlying representation and spreads to the prefix by PHA; the root vowel then raises to [+high, ATR] by RVR.

Is there a learning theory for the analysis presented in Archangeli & Pulleyblank (2015)? The EG as presented there would have to be greatly enriched beyond the general abilities in (30) in order to arrive at their analysis of Esimbi. We don't know what such a theory would look like when fully spelled out.

[27] Just to be clear, the prefix features in (38b) are preferences, not requirements; the reason for this will become evident in the next section. Nor are they equivalent to floating features. As we have not yet decided on the features of /ɛ, ɔ/, for now I will follow Archangeli & Pulleyblank (2015) in specifying them as [non-high] and [RTR].

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

In terms of the representations it ends up with, it is not very different – so far! – from the abstract UG analysis. However, whereas the UG theory posits a single unified mechanism of feature transmission for transparent and opaque situations, the EG theory posits two entirely different mechanisms, thereby requiring more, not less, phonology-specific innate machinery, or so it appears. Be that as it may, Archangeli & Pulleyblank (2015) have not shown that there is a UG opacity 'issue' that needs to be resolved, or how EG resolves it.

Some questions about the analysis have yet to be answered:

- Why do Archangeli & Pulleyblank (2015) write that roots 'prefer' certain prefixes? Why not 'select' or 'assign'?
- Why are the Set B prefix preferences left unmarked?
- Finally, what exactly are the phonological features of Esimbi vowels?

To answer these questions, we have to look at a third Esimbi prefix, which Hyman (1988) calls the A- prefix. This prefix considerably complicates the analysis of Esimbi vowel phonology. And here we will see a considerable difference in the adequacy of the UG analysis compared to the EG analysis.

5.3. Further intricacies of the Esimbi vowel system: the A- prefix

As Hyman (1988) observes, the A- prefix appears one height degree lower than the vowels of the I- and U- prefixes which occur with the same root. For example, the SG 3 U- prefix in (39a) has the allophones shown. In (39b) are the same roots with the PL 6 A- prefix: [u] in (39a) corresponds to [o] in (39b); [o] corresponds to [ɔ] and [ɔ] corresponds to [a].

(39) *Allomorphs of the SG 3 and PL 6 prefixes co-occurring the same root vowel*

a. *The SG 3 U- prefix*

[u-] with 'end':	u-tili
[o-] with 'spear':	o-tɨ
[ɔ-] with 'hand':	ɔ-bu

b. *The PL 6 A- prefix*

[o-] with 'end':	o-tili
[ɔ-] with 'spear':	ɔ-tɨ
[a-] with 'hand':	a-bu

Unlike the I- and U- prefixes, the A- prefix does not consistently appear as front or round, nor does it consistently receive these features from the root. As shown in (40), A- appears as [o] with /i, u/; as [ɔ] with /o, ə/; as [ɛ] with /e/ and as [a] with /ɛ, a, ɔ/. Hyman (1988) writes that the frontness and rounding of the A- prefix are presumably determined by secondary processes.

(40) *Realisations of the A- prefix* (Hyman 1988: 259)

	Root vowel	*A- prefix*
a.	/i, u/	[o-]
b.	/ə, o/	[ɔ-]
c.	/e/	[ɛ-]
d.	/ɛ, a, ɔ/	[a-]

We will be mainly concerned with the A- prefix specifications that have the effect of lowering the height assigned by the root vowels by one step. This is not a trivial problem! Let's first look at Archangeli & Pulleyblank's approach, which is very different from any UG account. They do not characterise the three prefixes in terms of features, but rather as sets of allomorphs. Their analysis is summed up in Table 2.

Recall that Archangeli & Pulleyblank (2015) specify the Set A roots, which Hyman considers to be underlyingly high, as preferring prefixes that are [high] and [ATR] (e.g., $bi_{HI,\ ATR}$ 'goat'). The highest prefixes in the front and round classes, [i] and [u], are both [high] and [ATR]. The highest prefix in the non-high class, [o], is [ATR] but not [high]. This is why prefix selection is framed as preferences: Set A

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Table 2. *Descriptive summary of Esimbi prefixes (Archangeli & Pulleyblank 2015: 6; adapted from their Table 8)*

	Prefix class		
Root	Front	Round	Non-high
Set A	i	u	o
Set B	e	o	ɛ *or* ɔ
Set C	ɛ	ɔ	a

prefers prefixes that are [high] and [ATR], but will accept only one of these if necessary. In Archangeli & Pulleyblank's (2015) analysis, [o] is the only non-high prefix that is [ATR]; [ɛ, ɔ, a] are all [RTR].

The Set C roots, which Hyman considers to be underlyingly /ɛ, ɔ, a/, prefer prefixes that are [low] and [RTR] (e.g., $zu_{LO,\ RTR}$ 'snake'). In Archangeli & Pulleyblank's analysis, [ɛ, ɔ] are the only front and round prefixes that are [RTR] (though not [low]). In the non-high prefix, [ɛ, ɔ] are [RTR] but lose to [a] which is also [low].

Now consider the Set B roots, which Hyman considers to be underlyingly /e, o, ə/. In Archangeli & Pulleyblank's analysis, these roots have no specified preferences (e.g., {ki} 'tail') but take the 'unmarked' member of each prefix, stipulated to be the middle member. Now we can see that they are forced to this solution because there is no way to state the Set B preferences in a positive way. One would need to say that Set B prefers [non-high] and [ATR] to account for the front and round prefixes, but this selects the wrong vowel [o] in the non-high prefix. The non-high allomorphs that need to be selected are [ɛ, ɔ], which are [RTR]; but then the front and round prefixes will select the wrong vowels [ɛ] and [ɔ].

Let's reflect for a moment on how Archangeli & Pulleyblank (2015) specify the variants of each prefix. Each prefix consists of a list of allomorphs, as in Table 2. Why these particular allomorphs? In this analysis, there is no reason why the heights of the front and round allomorphs correspond to each other. And it appears to be a coincidence that the non-high allomorphs are exactly one step below them. Thus, the EG analysis has no explanation for why the prefix allomorphs are the way they are.

By contrast, here is what Hyman (1988: 260) considered to be the criterion for a successful analysis of the Esimbi prefixes (emphasis added):

> To repeat our aim, an analysis must be sought that correctly and *insightfully captures the vowel height relations* between the prefixes I- and U- *vs.* those of the prefix A-, which are one step lower.

The EG analysis abandons this aim: it simply stipulates the vowel heights of each prefix. So let's look at a UG analysis in the spirit of Hyman (1988).

5.4. A UG analysis of the A- prefix in Esimbi

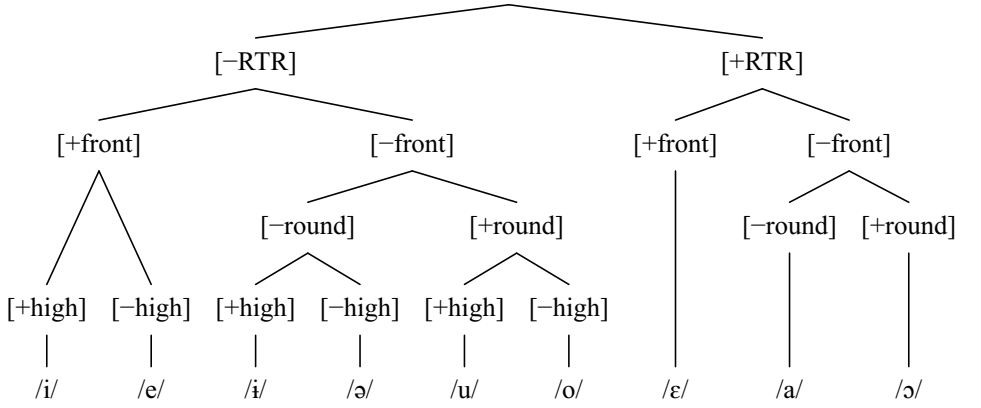
What features must the A- prefix have to produce a vowel that is one step lower than the root vowel? There are theories which can do this very elegantly; for example, the Particle Phonology of Schane (1984) or versions of Element Theory (Kaye *et al.* 1985; Backley 2011). In such theories, we could say that the A- prefix literally consists of an A element; adding one A to the root vowel lowers it by one step. Maybe this approach is right! But here I will continue to follow the feature theory set out in (7).

Hyman (1988) makes a number of proposals to account for the behaviour of the A- prefix. The one I find most compelling is that /a/ is characterised by [−ATR] (which, for purposes of this discussion, is equal to [+RTR]). The idea is that [−ATR] can interact with height features in particle phonology style, lowering each height by one step. While I agree with the spirit of Hyman's analysis, I want to modify his specifications slightly. I assume the active feature is [+RTR]. Following Archangeli & Pulleyblank (2015) and for reasons of contrast, I consider /ɛ, ɔ/ to also be [+RTR]. Finally, I would like to include /ɨ/; though it is a derived vowel, it consists entirely of contrastive features, so is part of the contrastive

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

inventory at some point. A possible contrastive feature hierarchy for this system is given in (41a); the feature specifications so generated are displayed in (41b).

(41) a. *Esimbi vowel hierarchy:* [RTR] ≫ [front] ≫ [round] ≫ [high]



b. *Esimbi vowel features generated by (41a)*

	/i/	/e/	/ɛ/	/ɨ/	/ə/	/a/	/u/	/o/	/ɔ/
[RTR]	−	−	+	−	−	+	−	−	+
[front]	+	+	+	−	−	−	−	−	−
[round]				−	−	−	+	+	+
[high]	+	−		+	−		+	−	

Now that we have an explicit account of the vowel features of Esimbi, we can make the analysis outlined in (35) more precise, as in (42).

(42) *The analysis of Hyman (1988) in terms of the features in (41)*

a. *Prefixes*
   i. I- prefixes [i, e, ɛ] are specified [+front]
   ii. U- prefixes [u, o, ɔ] are [−front, +round]
   iii. the A- prefix [o, ɔ, ɛ, a] is [+RTR]

b. *Roots*
   i. [i] roots derive from /i/, /e/, or /ɛ/
   ii. [u] roots derive from /u/, /o/, or /ɔ/
   iii. [ɨ] roots derive from /ə/ or /a/

c. *Prefix height assimilation*: Prefixes that are unspecified for height and/or [RTR] assimilate to those features of the roots.

d. *Root vowel raising*: Following PHA, all root vowels become [−RTR, +high].

The specifications of the I- and U- prefixes (42a) are slightly revised: I- is [+front] rather than [−back], and U- is [−front] in addition to [+round], to conform to the feature hierarchy in (41a). Following Hyman (1988), I assume that the A- prefix is specified only [+RTR]. The underlying root vowels in (42b) are as before.

In the current analysis, the height of a vowel is regulated by two features, [±high] and [±RTR]; therefore, both PHA in (42c) and RVR in (42d) must involve both those features. The output of RVR is a vowel that is [−RTR, +high]: underlying /i, u/ do not change; /e, o, ə/ change to [+high] and /ɛ, ɔ, a/ change to [−RTR] and add [+high].

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

According to PHA, prefixes unspecified for [RTR] and/or [high] assimilate to the root values of those features. In particular, the I- and U- prefixes receive: [−RTR, +high] from /i, u/, yielding [i, u]; [−RTR, −high] from /e, o, ə/, yielding [e, o] and [+RTR] from /ɛ, ɔ, a/, yielding [ɛ, ɔ].

The A- prefix is specified [+RTR]. The result of combining it with [−RTR, −high] from the root vowels /e, o, ə/ is a [+RTR] vowel with front/round features as specified in (44a) and (44b).[28] When [+RTR] combines with root vowels /i, u/, the expected result is [+RTR, +high] (*[ʊ], looking ahead to (44b)). The actual result is [o]. As Hyman (1988: 267) explains, the pronunciation of [ɪ, ʊ] can be very close to [e, o]; in several languages, these sounds have merged. Moreover, a constraint against [+RTR, +high] is cross-linguistically very common (Archangeli & Pulleyblank 1994). Thus, we have the principle in (43), a kind of STRUCTURE PRESERVATION (Kiparsky 1985) whereby vowels specified [+RTR, +high] are interpreted as [−RTR, −high].

(43) *Structure Preservation (SP): Interpretation of [+RTR, +high]*
When [+high] combines with a vowel specified [+RTR], the vowel is interpreted as the phonetically nearest [−RTR] vowel, which is [−RTR, −high].

The front/round features of the A- prefix are derived following the principles in (44).

(44) *Front/round assignment (FRA): Deriving the front/round features of the A- prefix*
a. When the root vowel is [−high, +front] (/e/), the [+RTR] prefix is assigned [+front].
b. Otherwise, when the root vowel is [±high] (/i, u, ə, o/), the [+RTR] prefix is assigned [−front, +round].
c. When the root vowel has no height feature (/ɛ, a, ɔ/), the [+RTR] prefix is assigned [−front, −round].

To illustrate the above analysis, for roots with underlying [±high] vowels, we have the outcomes in (45) for the A- prefix.[29]

(45) *Derivations of the A- prefix: roots with underlying [±high] vowels*
a. *Roots with underlying high vowels:* /[+RTR]+tili/ 'PL 6-end'

	Prefix		*Root*
UR	/[+RTR]	+	t [−RTR, +front, +high] lV/
(42c) PHA	[+RTR, +high]		—
(44b) FRA	[+RTR, +high, −front, +round]		—
(42d) RVR	—		—
(43) SP	[−RTR, −high, −front, +round]		—
SR	[o-tili]		

b. *Roots with underlying /e/:* /[+RTR]+ke/ 'PL 6-tail'

	Prefix		*Root*
UR	/[+RTR]	+	k [−RTR, +front, −high]/
(42c) PHA	[+RTR, −high]		—
(44a) FRA	[+RTR, −high, +front]		—
(42d) RVR	—		k [−RTR, +front, +high]
SR	[ɛ-ki]		

[28]These root vowels have a [−high] feature that does not co-occur with the underlying [+RTR] vowels. Either this feature remains in the A- prefix, resulting in new allophones that are indistinguishable from underlying /ɛ, ɔ/, or else changing [−RTR] to [+RTR] operates in structure-preserving fashion, in which case the [−high] feature is deleted.

[29]For roots with two vowels, as in (45a), I specify the features for the first vowel and represent the second one as V, where it is understood that the second vowel must be identical to the first. Alternately, we can suppose that vowel features of the root are attached to the root, as in the case of the floating features in Uyghur.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

c. *Roots with underlying /ə, o/:* /[+RTR]+tə/ ‘PL 6-spear’

	Prefix		*Root*
UR	/[+RTR]	+	t [−RTR, −front, −round, −high]/
(42c) PHA	[+RTR, −high]		—
(44b) FRA	[+RTR, −high, −front, +round]		—
(44b) RVR	—		t [−RTR, −front, −round, +high]
SR		[ɔ-tɨ]	

When the A- prefix occurs with the [+RTR] root vowels /ɛ, a, ɔ/, it simply remains [+RTR]; in this case, it receives the default features [−front] and [−round] (44c), resulting in [a]. A sample derivation is shown in (46).

(46) *Deriving the A- prefix with [+RTR] roots:* /[+RTR]+bɔ/ ‘PL 6-hand’

	Prefix		*Root*
UR	/[+RTR]	+	[+RTR, −front, +round]/
(42c) PHA	—		—
(44c) FRA	[+RTR, −front, −round]		—
(42d) RVR	—		[−RTR, −front, +round, +high]
SR		[a-bu]	

As we saw for the I- and U- prefixes, when the A- prefix combines with a root, learners receive conflicting signals; in this case, the conflict is compounded:

- in a root like /tu/ ‘ear’, the root vowel signals that it is [+high];
- the U- prefix in [o-tu] signals that the root vowel is [−RTR, −high];
- the A- prefix in [ɔ-tu] signals yet a different vowel height, [+RTR].

The simplest solution, and the one most consistent with the UG assumed here, is a version of the one adopted by Hyman (1988): the I- and U- prefixes show the underlying height of the root, and the A- prefix lowers it. Three sets of features are in play:

- the surface features on the root vowels after raising;
- the underlying height/RTR features assigned to the root that show up in the I- and U- prefixes;
- and the [+RTR] feature in the representation of the A- prefix.

5.5. Other analyses of Esimbi

Other writers have discussed the Esimbi prefix alternations from various points of view: Clements (1991b), Odden (1991) and Halle (1995) adopt the basic analysis of Hyman (1988), with different theories of features and feature geometry; Walker (1999) and Kalinowski (2009) employ an OT framework and van der Hulst (2018: 261–263) uses Radical CV Phonology. Though the details of these analyses differ, they all, as do Archangeli & Pulleyblank (2015), attribute the height (and tongue root features, for some) of the prefixes to features associated with the roots.[30] Halle (1995: 40) concludes his discussion with the following observation about the learnability of an abstract analysis of Esimbi:

> Esimbi speakers can therefore readily reconstitute the feature values of the underlying stem vowel: those for [back] and [round] surface on the stem vowel, those for [low] and [high] on the prefix. Thus, from an informational point of view the situation is only slightly less transparent than that encountered in a language where prefixes copy all features of the stem vowel.

[30]Kalinowski (2009) identifies a fourth prefix, E-, in addition to the three assumed by other analyses. She discusses additional data not accounted for by other studies, but does not propose a solution.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

6. Reanalysing opacity in Polish

One of the best-known cases of an opaque rule interaction in the classical generative phonology literature involves the interaction of *o*-raising and WFD in Polish (Kenstowicz & Kisseberth 1979; Bethin 1992; Kenstowicz 1994: 74–78). More recent analyses have argued that these accounts are wrong. In this section, I will review this controversy and agree that a reanalysis of the Polish data is called for. Nevertheless, the fact that a particular case of opaque interaction turns out to be an incorrect analysis cannot be taken as an argument against all such analyses in general. I will argue that the best synchronic analysis remains an abstract phonological one along the lines of the other cases discussed in this article.

6.1. The Polish problem in its classical form

The Polish opaque interaction became well known due to its inclusion in a problem set by Kenstowicz & Kisseberth (1979: 72–73). Some of the essential forms, as presented there, are shown in (47).

(47) *The Polish data* (following Kenstowicz & Kisseberth 1979)

	Singular	*Plural*	*Gloss*		*Singular*	*Plural*	Gloss
a.	klup	klubi	'club'	e.	ʃum	ʃumi	'noise'
b.	trup	trupi	'corpse'	f.	dom	domi	'house'
c.	ʒwup	ʒwobi	'crib'	g.	ʒur	ʒuri	'soup'
d.	kot	koti	'cat'	h.	bur	bori	'forest'

Solving this problem set requires the two rules in (48), applying in the order given.

(48) *Rules required to solve the problem set in (47)*

a. *o*-Raising: $/o/ \rightarrow [+high] / __ \begin{bmatrix} C \\ \begin{matrix}+voice \\ -nasal\end{matrix} \end{bmatrix} \#$

b. Word-final devoicing (WFD): $[-sonorant] \rightarrow [-voice] / __ \#$

The two rules in (48) are not particularly complicated, though vowel raising before a final voiced consonant is very rare, perhaps unique.[31] Nevertheless, their opaque interaction makes the problem surprisingly difficult if one encounters it on a time-limited exam. The trick is to realise that an underlying *o* raises before an underlying word-final voiced consonant, as seen clearly in the case of a non-nasal sonorant, as in (47h). However, when the final voiced consonant is an obstruent, as in (47c), the raised vowel only appears when the consonant has been devoiced, as in the singular, and is not raised when the final obstruent is overtly voiced, as in the plural. The opaque interaction of these rules is shown in the derivations in (49).

(49) *Polish opaque derivation*

UR	a. /klub#/	b. /trup#/	c. /ʒwob#/	d. /kot#/
o-Raising	—	—	ʒwub	—
WFD	klup	—	ʒwup	—
SR	[klup]	[trup]	[ʒwup]	[kot]

[31] As Kenstowicz (1994: 78) puts it, 'the Polish raising rule cannot readily be assigned to a natural phonological category. It is not a rule of assimilation, weakening, strengthening, enhancement, and so on. Polish raising typifies what Jan Baudouin de Courtenay (1845–1929) called a *sinnlos* ("irrational") alternation.' Nor does it adhere to the template in (10).

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

6.2. A revised view of Polish vowel raising

The forms in (47) make for a great exam question, so difficult that it appears that Polish learners did not all pass the test. According to Buckley (2001) and Sanders (2003), there are exceptions in both directions to *o*-raising as formulated in (48a): some vowels followed by an underlying voiced consonant do not raise, as in (50a)–(50d), and some vowels followed by an underlying voiceless consonant do raise, as in (50e)–(50h).[32]

(50) *Exceptions to vowel raising* (Buckley 2001; Sanders 2003)

	UR	*SR*	*Gloss*		*UR*	*SR*	*Gloss*
a.	/batog/	[batok]	'cowhide'	e.	/stop/	[stup]	'foot.GEN.PL'
b.	/xowd/	[xowt]	'homage'	f.	/kop/	[kup]	'heap.GEN.PL'
c.	/tor/	[tor]	'rail'	g.	/robot/	[robut]	'work.GEN.PL'
d.	/xol/	[xol]	'lobby'	h.	/tsnot/	[tsnut]	'virtue.GEN.PL'

To the lexical exceptions, Sanders (2003) adds experimental evidence showing that the *o~u* alternation is not productive, and he joins Buckley (2001) in inferring that *o*-raising as formulated in (48a) is not a synchronic rule in the grammar of Modern Polish. Sanders (2003: 57) concludes that 'thus, there is no need to analyze this case of opacity with a synchronic grammar'. Sanders (2003: 14) assumes a strong form of lexicon optimisation (Prince & Smolensky [1993] 2004) without lexical minimisation; on this approach, even predictable information is stored in the lexicon in order to make underlying representations as close as possible to surface representations. Thus, he derives stems with alternating *o~u* (actually *ɔ~u*) from more than one underlying form; for example, *ʒwup* 'crib' would derive from underlying /ʒwup/ and *ʒwɔbi* 'cribs' would derive from underlying /ʒwɔb-/.

Strong lexicon optimisation directly contradicts the morphophonemic principle in (5d) and related assumptions, and treats many predictable phonological alternations as suppletion. In the theory assumed here, the *o~u* alternation has to get some representation in a synchronic grammar.

Buckley (2001) proposes that Polish retains a rule raising a back vowel before a word-final consonant, but that the voicing of that consonant is no longer relevant. He proposes rather that the vowel that alternates between [o] and [u] cannot be either /o/ or /u/, neither of which alternate, but must be a third vowel different from both of them. He proposes that this vowel, which he designates /O/, is underspecified for [high]. We thus have the contrasts shown in (51).

(51) *Underspecified /O/ in Polish* (Buckley 2001)

a.	Stable /u/	[+high]	e.g., *krul-a ~ krul*	'king'	/krul/
b.	Stable /o/	[−high]	e.g., *gol-a ~ gol*	'goal'	/gol/
c.	Alternating /O/	[high]	e.g., *mol-a ~ mul*	'moth'	/mOl/

The merit of this proposal is that it posits three different vowels to account for three different behaviours. It also reflects the historical developments that led to the current grammar. As recounted in the sources (see Sanders 2003: 57–85), there was no historic raising of *o* before a word-final voiced consonant. Rather, at an early stage, *o or *ɔ *lengthened* before a word-final voiced consonant, perhaps as a compensatory lengthening for the loss of a final jer. The lengthened *oː* and short *o* subsequently diverged in height, as [oː] *vs.* [ɔ]. Later, long *o* shortened and merged with *u to *u*. So at a relevant point, there were indeed three non-low back vowels: /ɔ, o, u/.

Buckley (2001) cites some precedents for his analysis (Kiparsky 1993; Buckley 1994; Inkelas *et al.* 1994). However, assuming CHT, I do not see a way to derive an /O/ underspecified for height that would be in contrast with both [−high] /o/ and [+high] /u/. An /O/ unspecified for height can occur in

[32] I follow Sanders's (2003) transcription, except for now I retain ⟨o⟩ from Kenstowicz & Kisseberth (1979) where Sanders and others have ⟨ɔ⟩; Buckley (2001) uses Polish orthography.

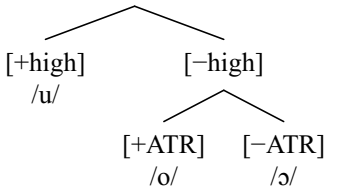
https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

ARCHIPHONEMIC fashion in contexts with a reduced number of contrasts, such as can arise, for example, in unstressed positions (Spahr 2014) or in desinences (Dyck 1995). But this is not such a case.

I would slightly revise Buckley's (2001) account. Following Sanders (2003) and others, I represent the non-alternating non-high vowel as /ɔ/. What learners would know is that the alternating vowel surfaces as either [ɔ] or [u] but is underlyingly distinct from both of them, because the latter do not alternate. A contrasting vowel that surfaces as either [ɔ] or [u] would most economically have features close to both of these vowels (cf. the 'intermediate' segment of Wang & Hayes in press). Such a vowel is /o/, which can be distinguished from /ɔ/ by [±ATR] and from /u/ by [±high].

A partial feature hierarchy for the Polish non-low back vowels is given in (52), and (53) shows sample derivations. When /o/ becomes [+high], the only destination is /u/. Elsewhere, unraised [−high] /o/ must become [−ATR].

(52) *Partial contrastive hierarchy for Polish vowels:* [high] ≫ [ATR]



(53) *Polish derivations*

UR	a. /krul#/	b. /gɔl#/	c. /mol+a/	d. /mol#/
o → [+high]/__C#	—	—	—	mul
o → [−ATR]	—	—	məla	—
SR	[krul]	[gɔl]	[məla]	[mul]

7. Conclusion

There have been many attempts, from the taxonomic phoneme to the present day, to put limits on abstractness in phonology, because it has been believed that abstract grammars are difficult or impossible to learn. If that were indeed the case, then there would be good reason to restrict abstractness. I have tried to show on the contrary that learners will be guided to acquire abstract analyses if their innate cognitive endowment (UG) has certain principles. The required principles are not particularly outlandish, but rather have underpinned much work in mainstream phonology, generative or otherwise.

There are, of course, generalisations that are not learned, as has been argued for the Polish example in §6.2. In the other cases reviewed here, however, the non-abstract analyses attribute to learners essentially the same generalisations as the abstract analysis. All accounts of Esimbi, for example, assume that learners have discovered that the height of the prefixes is determined by features that are somehow associated with the roots. Opacity in these cases is thus not a learning problem, but is rather a concise way in the UG framework of reconciling conflicting signals in the data.

Lahiri (2002, 2020) has called attention to two types of PERTINACITY in phonology. In one type, output forms remain the same despite underlying changes in the grammar; in the other, phonetic realisations change but the grammar remains the same. Abstractness and opacity attest to a third type of pertinacity, that of the phonology itself, which perseveres undeterred by complexity and conflicting signals. These are inherent in the very nature of phonology, which is about how phonemes interact and affect each other.

Acknowledgements. I would like to thank Norbert Hornstein and Bill Idsardi for illuminating discussions of basic issues and for comments on an earlier draft. I am grateful, too, for comments from Daniel Currie Hall, Larry Hyman, Paul Kiparsky and audiences at the MOT^{h} Phonology/Phonetics Workshop at McGill University in April 2024 and the Workshop on Pertinacious

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Phonology & Morphology at Ettington Park, Stratford-upon-Avon, in September 2024. This article has also benefitted from the insightful reviews of three anonymous reviewers and an Associate Editor of this journal.

Competing interests. The author declares no competing interests regarding the publication of this article.

References

Archangeli, Diana & Douglas Pulleyblank (1994). *Grounded phonology*. Cambridge, MA: MIT Press.
Archangeli, Diana & Douglas Pulleyblank (2015). Phonology without universal grammar. *Frontiers in Psychology* **6**, article no. 1229, 12 pp.
Archangeli, Diana & Douglas Pulleyblank (2022). *Emergent phonology*. Berlin: Language Science Press.
Backley, Phillip (2011). *An introduction to Element Theory*. Edinburgh: Edinburgh University Press.
Baker, C. L. (1979). Syntactic theory and the projection problem. *LI* **10**, 533–581.
Baković, Eric (2011). Opacity and ordering. In John A. Goldsmith, Jason Riggle & Alan C. L. Yu (eds.) *The handbook of phonological theory*, 2nd edition. Oxford: Wiley-Blackwell, 40–67.
Baković, Eric & Lev Blumenfeld (2024). A formal typology of process interactions. *Phonological Data and Analysis* **6**, article no. 3, 43 pp.
Baković, Eric, Jeffrey Heinz & Jonathan Rawski (2022). Phonological abstractness in the mental lexicon. In Anna Papafragou, John C. Trueswell & Lila R. Gleitman (eds.) *The Oxford handbook of the mental lexicon*. Oxford: Oxford University Press, 11–32.
Battistella, Edwin (2022). The Prague School: Nikolai Trubetzkoy and Roman Jakobson. In Dresher & van der Hulst (2022), 221–241.
Belth, Caleb (2023a). A learning-based account of local phonological processes. *Phonology* **40**, 1–33.
Belth, Caleb (2023b). Towards a learning-based account of underlying forms: a case study in Turkish. In Tim Hunter & Brandon Prickett (eds.) *Proceedings of the Society for Computation in Linguistics (SCiL) 2023*. Amherst, MA: Association for Computational Linguistics, 332–342.
Belth, Caleb (2026). A learning-based account of phonological tiers. *LI* **57**, 229–265.
Bermúdez-Otero, Ricardo (2018a). In defence of underlying representations: Latin rhotacism, French liaison, Romanian palatalization. *Probus* **30**, 171–214.
Bermúdez-Otero, Ricardo (2018b). Stratal phonology. In Hannahs & Bosch (2018), 100–134.
Bethin, Christina (1992). *Polish syllables: the role of prosody in phonology and morphology*. Columbus, OH: Slavica.
Bloch, Bernard (1941). Phonemic overlapping. *American Speech* **16**, 278–284.
Breit, Florian, Bert Botma, Marijn van 't Veer, & Marc van Oostendorp (eds.) (2023). *Primitives of phonological structure*. Oxford: Oxford University Press.
Buckley, Eugene (1994). Default feature prespecification: the two [i]'s of Kashaya. *NELS* **24**, 17–30.
Buckley, Eugene (2001). Polish o-raising and phonological explanation. Presented at the Annual Meeting of the Linguistic Society of America, Washington, DC, January 2001.
Chandlee, Jane (2014). *Strictly local phonological processes*. PhD dissertation, University of Delaware.
Chandlee, Jane, Jeffrey Heinz & Adam Jardine (2018). Input Strictly Local opaque maps. *Phonology* **35**, 171–205.
Charles-Luce, Jan (1985). Word-final devoicing in German: effects of phonetic and sentential contexts. *JPh* **13**, 309–324.
Chomsky, Noam (1964). *Current issues in linguistic theory*. The Hague: Mouton.
Chomsky, Noam (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
Chomsky, Noam (2004). Beyond explanatory adequacy. In Adriana Belletti (ed.) *Structures and beyond: the cartography of syntactic structures*. Oxford: Oxford University Press, 104–131.
Chomsky, Noam & Morris Halle (1968). *The sound pattern of English*. New York: Harper and Row.
Clements, G. N. (1976). Palatalization, linking or assimilation? *CLS* **12**, 96–109.
Clements, G. N. (1991a). Place of articulation in consonants and vowels: a unified theory. *Working Papers of the Cornell Phonetics Laboratory* **5**, 77–123.
Clements, G. N. (1991b). Vowel height assimilation in Bantu languages. *BLS* **17**, 25–63.
Clements, G. N. (2001). Representational economy in constraint-based phonology. In T. Alan Hall (ed.) *Distinctive feature theory*. Berlin: Mouton de Gruyter, 71–146.
Compton, Richard & B. Elan Dresher (2011). Palatalization and 'strong *i*' across Inuit dialects. *Canadian Journal of Linguistics* **56**, 203–228.
Courellis, Hristos S., Juri Minxha, Araceli R. Cardenas, Daniel L. Kimmel, Chrystal M. Reed, Taufik A. Valiante, C. Daniel Salzman, Adam N. Mamelak, Stefano Fusi & Ueli Rutishauser (2024). Abstract representations emerge in human hippocampal neurons during inference. *Nature* **632**, 841–849.
Cowper, Elizabeth & Daniel Currie Hall (2014). *Reductīo ad discrīmen*: where features come from. *Nordlyd* **41**, 145–164.
Danesi, Paolo (2022). *Contrast and phonological computation in prime learning: raising vowel harmonies analyzed with emergent primes in Radical Substance Free Phonology*. PhD dissertation, Université Côte d'Azur.
Dresher, B. Elan (1981). On the learnability of abstract phonology. In C. L. Baker & John J. McCarthy (eds.) *The logical problem of language acquisition*. Cambridge, MA: MIT Press, 188–210.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Dresher, B. Elan (1998). On contrast and redundancy. Presented at the Annual Meeting of the Canadian Linguistic Association, University of Ottawa, May 1998.
Dresher, B. Elan (2003). Contrast and asymmetries in inventories. In Anna-Maria di Sciullo (ed.) *Asymmetry in grammar, volume 2: morphology, phonology, acquisition*. Amsterdam: Benjamins, 239–257.
Dresher, B. Elan (2009). *The contrastive hierarchy in phonology*. Cambridge: Cambridge University Press.
Dresher, B. Elan (2014). The arch not the stones: universal feature theory without universal features. *Nordlyd* **41**, 165–181.
Dresher, B. Elan (2018). Contrastive Hierarchy Theory and the nature of features. *WCCFL* **35**, 18–29.
Dresher, B. Elan & Daniel Currie Hall (2022). Developments leading toward generative phonology. In Dresher & van der Hulst (2022), 372–395.
Dresher, B. Elan, Daniel Currie Hall, William J. Idsardi & Eric Raimy (2024). Uyghur disharmony without diacritics (or, phonological representations are phonological). Presented at the Annual Meeting of the Linguistic Society of America, New York, January, 2024. Poster and handout available online at https://incl.pl/lsa2024.
Dresher, B. Elan & Harry van der Hulst (eds.) (2022). *The Oxford history of phonology*. Oxford: Oxford University Press.
Dresher, B. Elan, Glyne L. Piggott & Keren D. Rice (1994). Contrast in phonology: overview. *Toronto Working Papers in Linguistics* **13**, iii–xvii.
Dyck, Carrie (1995). *Constraining the phonology–phonetics interface, with exemplification from Spanish and Italian dialects*. PhD dissertation, University of Toronto.
Elfner, Emily (2016). Stress–epenthesis interactions in Harmonic Serialism. In John J. McCarthy & Joe Pater (eds.) *Harmonic Grammar and Harmonic Serialism*. Sheffield: Equinox, 261–300.
Fortescue, Michael, Steven A. Jacobson & Lawrence D. Kaplan (1994). *Comparative Eskimo dictionary with Aleut cognates*. Fairbanks, AK: Alaska Native Language Center.
Gershman, Samuel J. & Yael Niv (2010). Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*, **20**, 251–256.
Goldsmith, John A. (1976). *Autosegmental phonology*. PhD dissertation, Massachusetts Institute of Technology.
Goldwater, Sharon & Mark Johnson (2003). Learning OT constraint rankings using a Maximum Entropy model. In Jennifer Spenader, Anders Eriksson & Östen Dahl (eds.) *Proceedings of the Stockholm Workshop on Variation within Optimality Theory*. Stockholm: Department of Linguistics, Stockholm University, 111–120.
Hagerup, Asger (2011). *A phonological analysis of vowel allophony in West Greenlandic*. Master's thesis, Norwegian University of Science and Technology.
Hahn, Reinhard F. (1991a). Diachronic aspects of regular disharmony in Modern Uyghur. In William G. Boltz & Michael C. Shapiro (eds.) *Studies in the historical phonology of Asian languages*. Amsterdam: Benjamins, 68–101.
Hahn, Reinhard F. (1991b). *Spoken Uyghur*. Seattle, WA: University of Washington Press.
Hale, Mark & Charles Reiss (2000). 'Substance abuse' and 'dysfunctionalism': current trends in phonology. *LI* **31**, 157–169.
Hall, Daniel Currie (2007). *The role and representation of contrast in phonological theory*. PhD dissertation, University of Toronto.
Hall, Daniel Currie (2023). Contrast and content in phonological features: substance use in moderation. In Breit *et al.* (2023), 108–130.
Hall, Daniel Currie & Avery Ozburn (2019). When is derived [i] transparent? a subtractive approach to Uyghur vowel harmony. *NELS* **49**, 89–102.
Halle, Morris (1959). *The sound pattern of Russian: a linguistic and acoustical investigation*. The Hague: Mouton.
Halle, Morris (1962). Phonology in generative grammar. *Word* **18**, 54–72.
Halle, Morris (1995). Feature geometry and feature spreading. *LI* **26**, 1–46.
Halle, Morris & Jean-Roger Vergnaud (1987). *An essay on stress*. Cambridge, MA: MIT Press.
Hannahs, S. J. & Anna R. K. Bosch (eds.) (2018). *The Routledge handbook of phonological theory*. New York: Routledge.
Haspelmath, Martin (2006). Against markedness (and what to replace it with). *JL* **42**, 25–70.
Hayes, Bruce & Colin Wilson (2008). A maximum entropy model of phonotactics and phonotactic learning. *LI* **39**, 379–440.
Heffner, Christopher C., William J. Idsardi & Rochelle S. Newman (2019). Constraints on learning disjunctive, unidimensional auditory and phonetic categories. *Attention, Perception, & Psychophysics* **81**, 958–980.
Hornstein, Norbert & David Lightfoot (1981). Introduction. In Norbert Hornstein & David Lightfoot (eds.) *Explanation in linguistics: the logical problem of language acquisition*. London: Longman, 9–31.
van der Hulst, Harry (2018). *Asymmetries in vowel harmony: a representational account*. Oxford: Oxford University Press.
Hume, Elizabeth (1994). *Front vowels, coronal consonants and their interaction in nonlinear phonology*. New York: Garland.
Hyman, Larry M. (1988). Underspecification and vowel height transfer in Esimbi. *Phonology* **5**, 255–273.
Hyman, Larry M. (2018). Why underlying representations? *JL* **54**, 591–610.
Idsardi, William J. (2022). Underspecification in time. *Canadian Journal of Linguistics* **67**, 670–682.
Inkelas, Sharon, C. Orhan Orgun & Cheryl Zoll (1994). Exceptions and static phonological patterns: cophonologies vs. prespecification. Ms, University of California, Berkeley.
Jaker, Alessandro & Paul Kiparsky (2020). Level ordering and opacity in Tetsǫ́t'ınné: a Stratal OT account. *Phonology* **37**, 617–655.
Jarosz, Gaja (2019). Computational modeling of phonological learning. *Annual Review of Linguistics* **5**, 67–90.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Jurgec, Peter & Yixin Wang (in press). Input prevocalic faithfulness and opacity. To appear in *LI*. Published online September 2025; forthcoming in print. https://doi.org/10.1162/ling.a.56.
Kalinowski, Cristin (2009). Multidirectional vowel harmony in Esimbi. *BLS* **35**, 147–155.
Kaplan, Lawrence D. (1981a). On Yupik–Inupiaq correspondences for *i*: a case of Inupiaq innovation. *Études/Inuit/Studies* **5**, 85–89.
Kaplan, Lawrence D. (1981b). *Phonological issues in North Alaskan Inupiaq*. Fairbanks, AK: Alaska Native Language Center.
Kaplan, Lawrence D. (1990). The language of the Alaskan Inuit. In Dirmid R. F. Collis (ed.) *Arctic languages: an awakening*. Paris: UNESCO, 131–158.
Katzir, Roni, Timothy J. O'Donnell & Ezer Rasin (2021). Introduction to the special issue on simplicity in grammar learning. *Journal of Language Modelling* **9**, 1–4.
Kaye, Jonathan (1974). Opacity and recoverability in phonology. *Canadian Journal of Linguistics* **19**, 134–149.
Kaye, Jonathan, Jean Lowenstamm & Jean-Roger Vergnaud (1985). The internal structure of phonological elements: a theory of charm and government. *Phonology Yearbook* **2**, 305–328.
Kenstowicz, Michael J. (1994). *Phonology in generative grammar*. Cambridge, MA: MIT Press.
Kenstowicz, Michael J. & Charles W. Kisseberth (1979). *Generative phonology: description and theory*. Orlando, FL: Academic Press.
Kenstowicz, Michael J. & Charles W. Kisseberth (2022). Phonological derivation in early generative grammar. In Dresher & van der Hulst (2022), 419–439.
Kiparsky, Paul (1973). Phonological representations. In Osamu Fujimura (ed.) *Three dimensions of linguistic theory*. Tokyo: Taikusha, 3–136.
Kiparsky, Paul (1982a). From Cyclic Phonology to Lexical Phonology. In Harry van der Hulst & Norval Smith (eds.) *The structure of phonological representations (part I)*. Dordrecht: Foris, 131–175.
Kiparsky, Paul (1982b). Lexical morphology and phonology. In Seok Yang (ed.) *Linguistics in the morning calm: selected papers from SICOL-1981*. Seoul: Hanshin, 3–91.
Kiparsky, Paul (1985). Some consequences of Lexical Phonology. *Phonology Yearbook* **2**, 85–138.
Kiparsky, Paul (1993). Blocking in non-derived environments. In Sharon Hargus & Ellen Kaisse (eds.) *Studies in lexical phonology*. San Diego, CA: Academic Press, 277–313.
Kiparsky, Paul (2000). Opacity and cyclicity. *The Linguistic Review* **17**, 351–365.
Kiparsky, Paul (2022). Pāṇini. In Dresher & van der Hulst (2022), 38–63.
Kochetov, Alexei (2011). Palatalization. In Marc van Oostendorp, Colin J. Ewen, Elizabeth V. Hume & Keren Rice (eds.) *The Blackwell companion to phonology*, vol. 3, chapter 71. Oxford: Wiley-Blackwell, 1666–1690.
Krämer, Martin (2012). *Underlying representations*. Cambridge: Cambridge University Press.
Kulikov, Vladimir (2012). *Voicing and voice assimilation in Russian stops*. PhD dissertation, University of Iowa.
Ladd, D. Robert (2022). Mid-century American phonology: the post-Bloomfieldians. In Dresher & van der Hulst (2022), 356–371.
Lahiri, Aditi (2002). Pertinacity in representation and change. Presented at the Workshop on Pertinacity, Konstanz, July 2002.
Lahiri, Aditi (2018). Predicting universal phonological contrasts. In Larry M. Hyman & Frans Plank (eds.) *Phonological typology*. Berlin: De Gruyter Mouton, 229–272.
Lahiri, Aditi (2020). Converging evidence for morpho-phonological pertinacity: diachrony and experimental psycholinguistics. Presidential lecture presented at the Annual Meeting of the Philological Society, online, June 2020. Available at https://www.youtube.com/watch?v=JhJvgqrFLV8.
Lahiri, Aditi & Vincent Evers (1991). Palatalization and coronality. In Carole Paradis & Jean-François Prunet (eds.) *The special status of coronals: internal and external evidence*. New York: Academic Press, 79–100.
Liberman, Mark & Alan Prince (1977). On stress and linguistic rhythm. *LI* **8**, 249–336.
Lindblad, Vern M. (1990). *Neutralization in Uyghur*. Master's thesis, University of Washington.
Martins, Pedro Tiago (2017). There is no place for markedness in biologically-informed phonology. In Bridget D. Samuels (ed.) *Beyond markedness in formal phonology*. Amsterdam: Benjamins, 219–231.
Mayer, Connor (2021). *Issues in Uyghur backness harmony: corpus, experimental, and computational studies*. PhD dissertation, University of California, Los Angeles.
Mayer, Connor, Travis Major & Mahire Yakup (2022a). Are neutral roots in Uyghur really neutral? Evaluating a covert phonemic contrast. In Peter Jurgec, Liisa Duncan, Emily Elfner, Yoonjung Kang, Alexei Kochetov, Brittney K. O'Neill, Avery Ozburn, Keren Rice, Nathan Sanders, Jessamyn Schertz, Nate Shaftoe & Lisa Sullivan (eds.) *Proceedings of the 2021 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
Mayer, Connor, Adam McCollum & Gülnar Eziz (2022b). Issues in Uyghur phonology. *Language & Linguistics Compass* **16**, article no. e12478, 25 pp.
McCarthy, John J. (2007). *Hidden generalizations: phonological opacity in Optimality Theory*. London: Equinox.
McCarthy, John J. (2008). The serial interaction of stress and syncope. *NLLT* **26**, 499–546.
Mielke, Jeff (2008). *The emergence of distinctive features*. Oxford: Oxford University Press.
Miyaoka, Osahito (2012). *A grammar of Central Alaskan Yupik (CAY)*. Berlin: De Gruyter Mouton.
Moulton, William G. (1962). *The sounds of English and German*. Chicago: University of Chicago Press.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press

Nazarov, Aleksei & Joe Pater (2017). Learning opacity in Stratal Maximum Entropy Grammar. *Phonology* **34**, 299–324.
Nespor, Marina & Irene Vogel (2007). *Prosodic phonology: with a new foreword*. 2nd edition. Berlin: Mouton de Gruyter.
Nevins, Andrew (2010). *Locality in vowel harmony*. Cambridge, MA: MIT Press.
Odden, David (1991). Vowel geometry. *Phonology* **8**, 261–289.
O'Hara, Charlie (2017). How abstract is more abstract? learning abstract underlying representations. *Phonology* **34**, 325–345.
O'Neill, Paul (2024). Morphologization and the boundary between morphology and phonology in the Romance languages. In Mark Aronoff (ed.) *Oxford research encyclopedia of linguistics*. Oxford: Oxford University Press, not paginated.
Papillon, Maxime (2020). *Precedence and the lack thereof: precedence-relation-oriented phonology*. PhD dissertation, University of Maryland, College Park.
Peters, Stanley (1972). The projection problem: how is a grammar to be selected? In Stanley Peters (ed.) *Goals of linguistic theory*. Englewood Cliffs, NJ: Prentice-Hall, 171–188.
Prince, Alan & Paul Smolensky ([1993] 2004). *Optimality Theory: constraint interaction in generative grammar*. Oxford: Blackwell. Originally published in 1993 as technical report no. 2 of the Rutgers Center for Cognitive Science.
Pruitt, Kathryn (2023). Serialism and opacity in phonological theory. *Annual Review of Linguistics* **9**, 497–517.
Raimy, Eric (2000). *The phonology and morphology of reduplication*. Berlin: De Gruyter Mouton.
Rasin, Ezer (2022). Shifting interactions and countershifting opacity: a note on opacity in Harmonic Serialism. *LI* **53**, 836–851.
Rasin, Ezer, Iddo Berger, Nur Lan & Roni Katzir (2018). Learning phonological optionality and opacity from distributional evidence. *NELS* **48**, 269–282.
Rasin, Ezer, Iddo Berger, Nur Lan, Itamar Shefi & Roni Katzir (2021). Approaching explanatory adequacy in phonology using Minimum Description Length. *Journal of Language Modelling* **9**, 17–66.
Rasin, Ezer & Roni Katzir (2016). On evaluation metrics in Optimality Theory. *LI* **47**, 235–282.
Rasin, Ezer & Roni Katzir (2020). A conditional learnability argument for constraints on underlying representations. *JL* **56**, 745–773.
Reiss, Charles (2018). Substance Free phonology. In Hannahs & Bosch (2018), 425–452.
Rice, Keren (2003). Featural markedness in phonology: variation. In Lisa Cheng & Rint Sybesma (eds.) *The second* Glot International *state-of-the-article book: the latest in linguistics*. Berlin: Mouton de Gruyter, 387–427.
Rice, Keren (2007). Markedness in phonology. In Paul de Lacy (ed.) *The Cambridge handbook of phonology*. Cambridge: Cambridge University Press, 79–98.
Rissanen, Jorma (1978). Modeling by shortest data description. *Automatica* **14**, 465–471.
Rissanen, Jorma & Eric Sven Ristad (1992). Language acquisition in the MDL framework. In Eric Sven Ristad (ed.) *Language computations*, Number 17 in DIMACS Series in Discrete Mathematics and Theoretical Computer Science. Providence, RI: AMS, 149–166.
Rubach, Jerzy (2003). Polish palatalization in derivational optimality theory. *Lingua* **113**, 197–237.
Samuels, Bridget D. (2011). *Phonological architecture: a biolinguistic perspective*. Oxford: Oxford University Press.
Sanders, Nathan (2003). *Opacity and sound change in the Polish lexicon*. PhD dissertation, University of California, Santa Cruz.
Schane, Sanford A. (1984). The fundamentals of Particle Phonology. *Phonology Yearbook* **1**, 129–155.
Scheer, Tobias (2008). Why the prosodic hierarchy is a diacritic and why the interface must be direct. In Jutta Hartmann, Veronika Hegedüs & Henk van Riemsdijk (eds.) *Sounds of silence: empty elements in syntax and phonology*. Amsterdam: Elsevier, 145–192.
Spahr, Christopher (2014). A contrastive hierarchical account of positional neutralization. *The Linguistic Review* **31**, 551–585.
Stallcup, Kenneth L. (1980a). A brief account of nominal prefixes and vowel harmony in Esimbi. In Luc Bouquiaux (ed.) *L'expansion bantou*, vol. 2. Paris: Société d'Études Linguistiques et Anthropologiques de France, 435–441.
Stallcup, Kenneth L. (1980b). Noun classes in Esimbi. *Southern California Occasional Papers in Linguistics* **8**, 139–153.
Trubetzkoy, Nikolai S. (1939). Grundzüge der Phonologie. *Travaux du Cercle linguistique de Prague* **7**, 1–271.
van 't Veer, Marijn, Bert Botma, Florian Breit & Marc van Oostendorp (2023). The structure and content of phonological primitives. In Breit *et al.* (2023), 1–36.
Walker, Rachel (1999). *Esimbi vowel height shift: implications for faith and markedness*. Ms, University of Southern California.
Wang, Yang & Bruce Hayes (in press). Learning phonological underlying representations: the role of abstractness. To appear in *LI*. Published online September 2025; forthcoming in print. https://doi.org/10.1162/ling_a_00553.
Wood, Sidney (1971). A spectrographic study of allophonic variation and vowel reduction in West Greenlandic Eskimo. *Working Papers of the Lund University Phonetics Laboratory* **4**, 58–94.
Yang, Charles (2005). On productivity. *Linguistic Variation Yearbook* **5**, 265–302.
Yang, Charles (2016). *The price of linguistic productivity: how children learn to break the rules of language*. Cambridge, MA: MIT Press.

https://doi.org/10.1017/S0952675726100372 Published online by Cambridge University Press