


Radical substance-free phonology and feature learning

DAVID ODDEN 

University of Washington, Seattle, Washington, USA

david.oddn@gmail.com

Abstract

This article argues that phonological features have no substantive properties, instead, segments are assigned features by learning strategies set to the task of devising a computational system for a phonology that is consistent with the requirements of UG. I address two problems for such a substance-free model. The first is the Card-Grammar problem, which has been suggested to argue for universal substantive features, on the premise that, otherwise, language data cannot be stored in a fashion necessary to correct learning errors. The Card Grammar problem disappears, in a suitably modular theory of mind with learned interfaces, where the mind still can retain information not parsed in a particular grammar. The second problem is the need for a demonstration, not just an assertion, that a reasonable theory of grammar and learning which has no access to phonetic substance can yield a coherent system of feature assignments. This is accomplished by modeling the learning of features necessary for the phonology of Kerewe.

Keywords: phonology, features, substance-free grammar, learning

Résumé

Cet article soutient que les traits phonologiques n'ont pas de propriétés substantielles, mais que les segments se voient attribuer des traits par des stratégies d'apprentissage dont la tâche est de concevoir un système informatique pour une phonologie qui soit cohérente avec les exigences de l'UG. J'aborde deux problèmes soulevés par un tel modèle sans substance. Le premier est le

This article is a development of 'Phonology Ex Nihilo', a presentation and handout from said presentation given to the phonology group at University of Tromsø on 6 Dec 2006, and expanded on at the Phonology Theory Agora, Nice on 16 March 2019 and Workshop on Theoretical Phonology, Concordia University on 8 May 2020. I thank the audiences of these venues for their helpful comments, as well as the editors and reviewers of *Revue canadienne de linguistique/Canadian Journal of Linguistics*. I especially thank Sylvia Blaho, Alex Chabot, Mark Hale, Martin Krämer, Bruce Morén-Duolljá, Mary Paster, Charles Reiss, Curt Rice, Tobias Scheer, Christian Uffmann, Shanti Ulfsbjörninn and Bert Vaux for pertinent feedback.

problème de la ‘grammaire des cartes’, qui a été suggéré pour plaider en faveur de traits universels substantiels, en partant du principe qu’autrement, les données linguistiques ne peuvent pas être stockées de manière à permettre la correction des erreurs d’apprentissage. Le problème de la ‘grammaire des cartes’ disparaît, dans une théorie modulaire appropriée de l’esprit, avec des interfaces apprises, où l’esprit peut encore retenir des informations qui ne sont pas analysées dans une grammaire particulière. Le second problème est le besoin de démontrer, et pas seulement d’affirmer, qu’une théorie raisonnable de la grammaire et de l’apprentissage qui n’a pas accès à la substance phonétique peut produire un système cohérent d’assignations de traits. Ceci est accompli en modélisant l’apprentissage des traits nécessaires à la phonologie du *kerewe*.

Mots clés: phonologie, traits, grammaire sans substance, apprentissage

1. INTRODUCTION

One of the most influential ideas of phonological theory, originating with Jakobson (1949), is that speech sounds are defined by the conjunction of a set of autonomous features.¹ Jakobson proposes that “features” (as they are now known) are binary, and have acoustic manifestations. Rather than viewing the sounds standardly symbolized as *t*, *ð* and *s* as unanalyzable atoms, Distinctive Feature theory defines such segments, following the widely-used version proposed in Chomsky and Halle (1968), in terms of the specific universal feature specifications.²

(1)	[t]	[ð]	[s]
	$\left[\begin{array}{l} -\text{sonorant} \\ -\text{continuant} \\ +\text{coronal} \\ +\text{anterior} \\ -\text{voice} \\ -\text{strident} \end{array} \right]$	$\left[\begin{array}{l} -\text{sonorant} \\ +\text{continuant} \\ +\text{coronal} \\ +\text{anterior} \\ +\text{voice} \\ -\text{strident} \end{array} \right]$	$\left[\begin{array}{l} -\text{sonorant} \\ +\text{continuant} \\ +\text{coronal} \\ +\text{anterior} \\ -\text{voice} \\ +\text{strident} \end{array} \right]$

This system for segment definition allows classes of sounds to be defined via the features which the members of the class have in common. All three of the above segments have in common the features [–sonorant, +coronal, +anterior], the segments *t* and *s* have in common [–sonorant, +coronal, +anterior, –voice], *t* and *ð* have in common [–sonorant, +coronal, +anterior, –strident], and *ð* and *s* have in common [–sonorant, +coronal, +anterior, +continuant]. Phonological rules are (by hypothesis) stated in terms of sets of segments defined by the features.

There have been numerous changes to the theory of features, ranging from models such as Jakobson’s or Fant and Halle’s (1952), up to geometric models

¹Abbreviations used: FP: Formal Phonology; H&R: Hale and Reiss (2008); Mo: module; RSFP: Radical Substance-Free Phonology; SPE: *The Sound Pattern of English*; UFT: Unified Features Theory; UG: Universal Grammar; VOT: voice onset time. Features: Ap: Approximant; Co: Coronal; Fr: Front, L: Labial; N: Nasal, Pa: Palatal; Pl: Place; Rd: Round; St: Stop; Ve: Velar; Vl: Voiceless.

²This is only a partial list of applicable feature specifications.

such as Clements and Hume's (1995), especially regarding how features relate to one another, and to the question of feature privativity. Theories of features have, for the most part, tacitly accepted two basic premises. First, it has generally been assumed that features are defined in terms of phonetic properties – articulatory, acoustic, or both. Second, it has been assumed that all languages draw on the same pre-defined set of features, which are provided by Universal Grammar (UG). This article sets forth an alternative view, namely that features are not defined in terms of physical substance, and are not listed in UG. Instead, the features employed in the grammar of any language are arrived at inductively, strictly on the basis of the representational requirements of a grammar. There are two basic requirements of phonological grammars. First, every sound-type (segment) of a language is represented by a unique configuration of features and prosodic elements: if *b* and *p* are both sounds in the phonology of a language, they have different representations. Second, rules of a phonological grammar apply in configurations that are identified via representational differences. When a rule applies differently in the context of *b* versus *p*, the difference comes from an interaction between how the rule is stated and the representational differences between the segments. This theory of features is formal, in that the inferred features are a consequence of the form of phonological rules which refer to features. Features are not based on physical substance.

The precision with which features have been defined has varied over the course of generative phonology. The SPE (*Sound Pattern of English*) theory of features is one of the most phonetically fine-grained theories, having at least 26 defined features. Subsequent theories, especially within feature-geometric trends, generally posit fewer and more abstract features. For example, the distinction between front and back vowels in SPE theory is governed by the feature [back] where front vowels are [-back] and the distinction between alveolars and labials is governed by [Coronal], whereas in Unified Features Theory (UFT: Clements and Hume 1995), they are governed by the same feature. The phonetic definition of [Coronal] in UFT is less specific, in abstracting away from exact details of tongue-raising. Likewise, in Bradshaw (1999), voicing and L tone are represented with a single feature. The Parallel Structures Model (Morén 2003) abstracts away from phonetic definitions even further, so that the features [Open] and [Closed] may distinguish different tone registers, vowel height, laryngeal constriction, or fricative vs. stop, depending on what structure these features are predicated of. Element Theory (Kaye, Lowenstamm and Vergnaud 1985 *et seq.*) likewise abstracts away from the fine-grained details of articulation and acoustics so that the unary element “H” may be realized as High tone, aspiration, or frication. See section 1.2 of Blaho (2008) for detailed analysis of similarities between various substance-free approaches to phonology.

Odden (2006) and Blaho (2008) advance the claim that phonological theory does not require or allow features to have any substantive definition at all. That theory, Radical Substance Free Phonology (RSFP), holds that phonological primitives have no intrinsic phonetic content, and no aspect of phonological grammar refers to the phonology-external substance of the segments referred to by the

primitives. In RSFP, [Coronal] is a formal label which, in conjunction with other formal labels, gives every segment a unique identifier, and these identifiers can be exploited to refer to classes of segments in operations of grammars – the rules. A feature conventionally called [Coronal] makes no claim about the tongue or F_2 .

RSFP differs from similar accounts, as discussed by Blaho, and also differs from the approach of Dresher (2009), which (similar to most theories) assumes some degree of universality and phonetic definedness in the set of features: but similar to that work, RSFP holds that the basis for feature learning is at least partially the fact of being active in the phonology (Dresher 2009 attributes underspecification choices to the evidence of rule activity). RSFP has some similarity to the approach of Mielke (2008) as well as Hale and Reiss (2008) in not attributing to phonology any reference to phonology-external factors, but differs from Mielke in attributing decisions about features to the requirements of phonological computation, which are part of the innate language faculty, namely the abstract syntax of features and rules. As held by Hale and Reiss, not everything is learned. Mielke, in contrast, (apparently) attributes feature learning to guidance by phonetic substance.

This article outlines how RSFP “works” in a system of phonological computations. The discussion of the theory of phonological computations (section 2) is brief because the theory adds nothing to existing autosegmental theories of computation and aims to take away much. The main focus of the article is showing how feature learning is possible. Features can be learned because rules can be learned, but learning that a rule exists and what its proper statement is depends on there being a theory of rules and rule-learning – this is the point of section 2. Section 3 addresses the card-grammar parsing problem raised by Hale and Reiss (2008), which suggests a logical basis for requiring detailed phonetic definitions of features as part of phonological UG. It is shown that while UG must encode the formal concept “feature” and provide a syntax of features and computations, no specific features with phonetic definitions are required in phonology. The problem with the card-grammar argument is that it does not distinguish between the information available to the extra-phonological interface device that parses sounds into grammatical segments which are operated on by the grammar, and the information actually available within the entire grammar or the mind. Putative phonetic properties of features are at most an aspect of *phonetic* computations, and only under certain assumptions about what the grammar of phonetic implementation is. Attention is then directed to the central problem of distinguishing phonetic computations from phonological computations, and the foundational question of whether there exists a phonetic component to grammar which performs phonetic computations.

Section 4 focuses on the logic of feature learning, starting with a very simple toy language, then proceeding to an analysis of the Bantu language Kerewe. Feature learning is based on the requirements of grammatical computation, so this section focuses on giving the correct analysis to the rules of Kerewe, from which the feature analysis of segments can be inferred.

2. THE FORMAL THEORY OF FEATURES AND RULES

As indicated above, the formal theory of features in RSFP is minimal. RSFP is, at the level of grammatical theory, part of an answer to the question “What are phonological representations?”, answered in a way that is consistent with the principles of Formal Phonology (FP) (see Odden 2013). FP is a development of the perspective on phonology articulated in Hale and Reiss (2008) (henceforth H&R), though differing from that work in certain ways. The framework presupposes that which is argued for in H&R, that phonology is a specific cognitive faculty which performs computations on stored representations, and this faculty constitutes part of a broader theory of human behavior. It does not posit a collection of substantive axioms as to what the tools of phonology are; instead, it posits a logic of determining what those tools might be. Most important for the purpose of this article, and also to some extent for distinguishing FP from the H&R model, formal simplicity is the fundamental arbiter behind the choice of theories of grammar and specific languages. If analyses A and B both function correctly as accounts of a fact, but A is simpler than B, then A is to be chosen over B.³

In this article I only address segmental features, leaving analogous questions about suprasegmental representation to separate investigation. However, there is no compelling reason to limit the radical substance-free approach to just segments. A segment is, formally, a set of features and dominance relations.⁴ A segment is formally the root of a feature tree, and everything under it. This is, of course, a simplified description of various autosegmental theories of features, including the feature geometry of Clements (1985) or of Sageyan (1986), Unified Features Theory, and the Parallel Structures Model.⁵ It is simpler than Sageyan geometry in eliminating certain stipulations such as the claim that [Coronal], [Labial], [Dorsal] must be immediately dominated by [Place]. Any dominance relation possible in one of these representational accounts is, in lieu of compelling theoretical reasoning to the contrary, also possible in RSFP.

2.1 Privativity

An important formal question arises as to the nature of features in RSFP: are features privative, or are they value-attribute pairs, that is, binary (or more)? Or, for that matter, is this a fact that needs to be learned by the child? At this point, I apply the logic of FP to the question, concluding that features are privative (thereby also illustrating the simplicity-based logic of FP). First, the alternative whereby a child

³The complex question of how comparative complexity is determined is explored in depth in Odden (2021).

⁴Conventionally, there is at least a two-way distinction drawn between necessarily terminal features such as [Voice] and potentially pre-terminal nodes like [Place]. This distinction is superfluous in the present substance-free account and can be replaced with the more general term “node”, as discussed in Odden (2021). In this article I will simply talk of “Place” or “Laryngeal” as being features.

⁵Dependency theories of segment representation are sufficiently different from autosegmental models that I will not pursue the question of whether element-learning is likewise possible in that approach.

must also learn whether features in a language are binary versus privative will be dismissed, on grounds based on the card-grammar logic: there has to be some fixed fact that serves as the innate basis for learning. That fixed fact is the formal nature of phonological computations. A child does not learn from scratch what it means to be “a rule in the grammar”, the child already knows that. The nature of rules depends crucially on whether features are just nodes (privative), or are value-attribute pairs. A theory of rules which says that rules change values of features is incompatible with a theory that says that rules change relationships between nodes. What the child does need to learn is, what are the rules of a particular grammar?

Still, we could assume that features are all binary, or all privative – what fact of a substance-free formalist framework says that features are privative? Suppose a child is learning a language with the segmental inventory of American English, and a rule does something to a set of vowels before the consonants {t, d, tʃ, dʒ, θ, ð, s, ʃ, n, l, r} but not before any other consonants. This class of segments has some characteristic that makes them similar in some way, to the exclusion of any other segment in the language, and we label that characteristic [Coronal]. We could also use a juxtaposition of value, and attribute [+coronal]. The latter theory of the syntax of features implies that the complement of this class is automatically inferred: the members of the set {p, b, m, f, k, g, ŋ} are automatically [–coronal]. The defect of the theory of binary features is subtle: it makes a claim for which there is insufficient evidence. There is no evidence that the complement of {t, d, tʃ, dʒ, θ, ð, s, ʃ, n, l, r} has such a similarity in phonological behavior. In not labelling the set {p, b, m, f, k, g, ŋ}, the privative theory does not put “lack of evidence” on a par with “evidence”.


Privative theory is formally simpler, because it posits one concept – a feature (more generally, a node) – whereas binary feature theory requires three concepts – “value”, “attribute”, and “feature” (the conjunction of value and attribute), and binary feature theory requires unnecessary theoretical propositions to the effect that a value cannot exist independent of an attribute (there are no floating plusses), nor can an attribute exist independently of a value. As emphasized in Odden (2013), Occam’s Razor is an essential tool of theorizing in FP. If there were sufficient evidence for introducing all three concepts into phonological theory, it would be possible to do so, but in lieu of such evidence, the simplest system of causal concepts is to be adopted. That system is the theory of privative features.

A reasonable counter-argument against privative features is that it seems to imply – incorrectly – that voicelessness cannot spread. This potential prediction might follow if voicedness (and not voicelessness) were the universally-assumed specified value, where voicelessness is the result of not specifying a segment with [Voice]. It is generally assumed that a rule can refer to the fact that a segment has a particular feature, but cannot refer to the lack of a specification for a feature, from which it would follow that a rule deleting [Voice] before a segment lacking the specification [Voice] would be a formally impossible rule – thus, voicelessness cannot spread.⁶


⁶This is a statement about assumptions made in the field, not an acceptance of the validity of the assumption, which needs independent scrutiny and validation, and that is beyond the scope of the present article. However, I do in fact adopt that assumption here.

There are at least two reasons why behavioral asymmetry is not mandated by privative features, especially in RSFP. First, the presumption that voicing is universally implemented via the feature [Voice] is untenable. The premise that all languages employ the specification [Voice] is directly counter to the premises of RSFP, where it is possible for a language to specify the laryngeal distinction as [Voiceless] (names are arbitrary, features have no intrinsic interpretation). Indeed, nothing in RSFP precludes having [Voice] and [Voiceless] coexist in a language. When voiceless segments act as a class under a rule, that behavior motivates the existence of a feature [Voiceless]. Nothing prevents a language from having a fact pattern motivating a feature [Voice] as well as a fact pattern motivating a feature [Voiceless]. The second reason why privativity does not mandate the phonological inertness of the presumed “unmarked” member of an opposition is that alleged spreading of a supposedly unspecified terminal node can be accomplished by spreading of a dominating preterminal node.

The logic of this argument is made clear in Lombardi (1991), who shows the ill-suitability of the asymmetry claim. There may exist a dominating node (such as [Laryngeal]), and when voicelessness seems to spread, it is not the terminal voicing feature that spreads, it is the dominating [Laryngeal] node that spreads, even when there is no lower-level feature corresponding to voicelessness. Spreading of voicelessness is formally just as possible as spreading of voicedness.



Moreover, if a language can have both the features “voiced” and “voiceless”, and two features can be dominated by a node, then a representation like (3) is possible.



This structure is analogous to the widely-accepted hypotheses that [Laryngeal] dominates [Constricted Glottis], [Spread Glottis] and [Voice], and that [Place] dominates [Labial], [Coronal], [Dorsal], [Radical]. Since the technical device of nodes dominating nodes can apparently accomplish everything that is accomplished by binary feature specifications, and nodes dominating nodes is a representational device with independent usage (at least in any autosegmental theory), it follows that without other evidence to support treating features as value-attribute pairs, features should be treated privatively.⁷

⁷Lombardi’s analysis also depends heavily on an invariant and phonetically-defined set of features and dominance relations, plus various stipulations regarding well-formed structures,

2.2 Order of features

A characteristic of many theories of feature geometry is that dominance relations are pre-specified by UG: [Place] dominates [Coronal] and not the other way around, [Laryngeal] dominates [Voice], but not [Nasal]. Because RSFP denies that specific features are in UG, it follows that dominance relations are not in UG. RSFP allows language-specific conditions, so if there is evidence for a node “Place” and for a node “Coronal”, there may also be evidence for a condition on representations saying that [Coronal] may (or must) be dominated by [Place]. Such rules are learned on the basis of pertinent evidence – the requirements for setting up a system of grammatical computations.

2.3 The theory of rules

The RSFP analysis of feature learning depends on having a theory of phonological computations combined with general principles of conceptual learning. That means: we have to have a theory of phonological computations. My argument is conducted in a minimalist FP-friendly rule-based model, but being rule-based is not a prerequisite for an RSFP analysis of feature learning. See Blaho (2008) for an OT-based instantiation of RSFP, though a computational model where constraints are learned. In Blaho’s model, UG does not contain specific constraints like Ident[Nasal], it contains a constraint form: Ident[___], where existence of the target feature must be learned. In that account, when the existence of a feature is learned – [Coronal] for example – the existence of a class of constraints Ident[Coronal], also OCP[Coronal], *[Coronal], and so on is thereby learned.

It is well beyond the scope of the article to articulate and defend a complete theory of phonological rules, nor is it necessary to do so. No novel assumptions about rule theory are required to facilitate feature learning. I follow the standard assumption that a rule combines two representations, the structural description which identifies the class of strings that undergo the rule, and a structural change which describes how the string is changed. A fundamental desideratum of autosegmental rule theory has been that the structural change should be reduced to a single operation. I assume, specifically, that a rule is limited to the insertion or deletion of a single dominance relation or node in the representation. That is, a rule can insert or delete an association line, or a segment, feature, or other phonological constituent. I do not assume abbreviatory schemata as encountered in the SPE theory of phonology, or rule-independent repairs or limitation constraints as encountered in parametric versions of autosegmental rule theory.⁸ There are a number of

which result in the prediction that if voicelessness spreads by some rule, glottalization and aspiration must also spread. Those premises are not valid in FP and RSFP. RSFP primarily posits the device of domination, allowing a language recourse to conditions on domination such as “[Laryngeal] dominates [Voice]” if there is evidence for such a condition.

⁸A reviewer wonders whether the analysis with minimal assumptions about features in UG is made possible by “quite a rich machinery elsewhere in the phonology”, that is to say, invoking certain representational and computational assumptions from Autosegmental Phonology. It is important to distinguish the wide range of grammatical devices posited in the history of

independent questions about rule theory that need to be answered, irrespective of the theory of features assumed. For example, is “structure preservation” a valid phonological concept? are rules subject to “blocking” conditions, and if so, what is the syntax of such expressions?

A salient unanswered question will remain unanswered for lack of relevant evidence: What is the theory (if any) of feature realization? Under the premise that *s* in some language, perhaps English, is [Coronal, Anterior, Continuant], a legitimate and interesting question is, by what mechanism is [Coronal, Anterior, Continuant] realized as *s* in the language? The general answer is, it happens in and after the phonetic component, and we need to at least understand the nature of phonetic computations (see section 3.3). We need a fully-articulated theory of phonetic implementation. There seem to be some differences in how *s* is pronounced in American English, Andalusian Spanish, Basque, Icelandic, Korean, and Modern Greek, and except in the case of Basque, these differences do not seem to reflect phonological patterns; instead, they are simply language-specific details about how a sound is pronounced. Does the phonetic grammar directly control the superior longitudinal, inferior longitudinal, transverse, and genioglossus muscles? Or, more likely, does the phonetic component produce a more abstract symbolic output which is the input to some other cognitive model? We simply don't know how phonological features are physically implemented, we just know that they are implemented, perhaps indirectly.

For the sake of concreteness, I assume that feature matrices like [Coronal, Anterior, Continuant] are somehow interpreted as “*s*, as pronounced in this language”, and I use transcriptional symbols to stand for a lower-level cognitive representation related to a grammatical feature matrix – perhaps an auditory image, a set of articulatory instructions, or maybe an autonomous linguistic phonetic object. A feature specification is thus a direct or indirect index to other cognitive entities, some of which are outside of grammar.

An important question cannot be resolved here, regarding the analysis of particular facts. The claim of this article is that a child induces a feature analysis of segments based on the requirements of phonological representation and computation – two sounds are different in the phonology, there is a phonological rule which identifies one sound to the exclusion of another. Therefore, we need to know whether a given rule is in fact part of phonology, or whether it is something else. This article does not claim to reveal how we are to distinguish the results of morphological computations or phonetic computations. Ultimately, we need coherent and mutually-informed theories of phonological, morphological, and phonetic computation. No amount of purely phonological argumentation will rule out an analysis where the surface [s], [z], and [ɪz] variants of the English plural suffix are computed in the morphology, nor will pure phonological reasoning inform us non-arbitrarily whether aspiration in English is computed by the phonology, or by the phonetics.

Autosegmental Phonology from the extremely sparse set of computational assumptions made here, and explored more extensively in Odden (2021). The alternative pre-autosegmental theory set forth in Chomsky and Halle (1968) is in fact much richer in computational apparatus that even the most all-encompassing version of Autosegmental Phonology.

3. THE CARD GRAMMAR ARGUMENT FOR PRE-DEFINED FEATURES

Hale and Reiss (2008, ch. 2) put the question of feature innateness into sharp focus, pointing to the necessity of some innate basis for learning. As they put it, “ya gotta start with something”. That is, not everything about language can be learned. This is a restatement of the Innateness of Primitives Principle (Pylyshyn, Fodor, Jackendoff), expressed in Jackendoff (1990: 40) as: “In any computational theory, ‘learning’ can consist only of creating novel combinations of primitives already innately available”⁹

In order for a child to learn that *s* has specification X and *θ* has specification Y, the child must be able to store the fact that the language has a segment *s* which is distinct from *θ*. This section scrutinizes (and rejects) the implication that this entails phonetically-defined innate features as part of Universal Grammar. The argument to that effect ultimately depends on a premise that is not self-evidently true, that phonological feature assignments arise deterministically from an unlearned direct interface between the auditory system and the phonological component. Given the alternative that features are assigned by a learned interface which relates pre-phonological cognitive representations to representations suited to phonological computation, there is no logical impediment to learning phonological features.

3.1 The Card Grammar argument

Hale and Reiss advance the “card grammar” argument for innateness of phonological features based on the supposed impossibility of learning contrastive feature assignments without the specific features being already available in UG. In their argument, cards correspond to sentences of natural languages, but the argument applies to any representation in grammar. Here I review the argument, laying bare the required assumptions. The argument holds if you make certain assumptions, and not otherwise.

The argument explores the consequences of different models of UG for what could be learned, using a stripped-down model of language, Card Grammar, where a grammar is a set of conditions on cards. Each model of card UG provides a set of primitive features and operators defined for those features. A card *c* is grammatical with respect to grammar *G* iff *c* satisfies the conditions imposed by *G*. In example UG₁, the primitives are NUMBERCARD (henceforth “#”), the suits ♣, ♦, ♥, and ♠, and the operator AND. Grammar *G*₁ is the rule [#], which means that only a card with the property [NUMBERCARD] is grammatical. A physical card [K♦]¹⁰ is ungrammatical, because it does not have the property [#]. Such an input is parsed

⁹I assume that the intended claim is that the thing-combinations which constitute “learning” are *reducible* to innate primitives, not that they *are* themselves innate primitives, but it is possible that the authors believe that learned concepts like “apple”, “fork”, “canid”, “mammal” and so on are themselves innate cognitive primitives.

¹⁰By “physical card”, I mean a mind-external actual card, notated with “|”, and not a mental representation of a card. This is analogous to H&R’s use of body brackets as in $\mathfrak{k}^{\mathfrak{h}}\mathfrak{a}\mathfrak{e}\mathfrak{t}$ to refer to a physical production of “cat”. I believe my use of these brackets clarifies their original intent, which is to say how a mind-external stimulus maps to a mental representation.

by UG_1 simply as $[\diamond]$ – the physical property identifiable as “K” is not assigned any representation by UG_1 .¹¹ The cards $|6\diamond|$, $|3\clubsuit|$ are parsed as $[\#\diamond]$, $[\#\clubsuit]$ (likewise $|3\diamond|$, $|6\clubsuit|$ are parsed as $[\#\diamond]$, $[\#\clubsuit]$). These cards are grammatical since they follow the rule requiring $[\#]$. G_3 has the rule $[\spadesuit]$, which means that any of $|2\spadesuit\dots A\spadesuit|$ are grammatical. The physical cards $|2\spadesuit|$ to $|10\spadesuit|$ are parsed by UG_1 as the same thing, $[\#\spadesuit]$, and $|J\spadesuit|$ through $|A\spadesuit|$ are parsed as $[\spadesuit]$. Following the rule of G_3 , $|2\spadesuit\dots 10\spadesuit|$ are all judged to be grammatical $[\#\spadesuit]$, and $|J\spadesuit|$ through $|A\spadesuit|$ are also judged to be grammatical, parsed as $[\spadesuit]$. The important point is that given the particular primitives provided by UG_1 , there are only eight possible mental card representations: $[\#\clubsuit]$, $[\#\diamond]$, $[\#\heartsuit]$, $[\#\spadesuit]$, $[\clubsuit]$, $[\diamond]$, $[\heartsuit]$, $[\spadesuit]$.¹²

UG_3 has a richer set of representational primitives, including [picturecard] = [P], individual numbers [2, 3, 4, 5, 6, 7, 8, 9, 10], and $[\pm\text{red}]$. This allows $|2\heartsuit|$ and $|3\heartsuit|$ to be parsed distinctly as $[2 +\text{red}]$ and $[3 +\text{red}]$, but does not allow $|2\heartsuit|$ to be distinguished from $|2\diamond|$, which are both represented as $[2 +\text{red}]$.¹³ On the other hand, UG_4 has the very impoverished inventory of features, containing only $[\diamond]$. Accordingly, physical cards are all parsed as either $[\diamond]$ or are not parsed at all. The upshot of this analysis is that unless UG has a very rich representational inventory, all physical inputs will be parsed as the same thing, or not parsed at all, and therefore there would be no basis for learning that $|2\heartsuit|$ is in fact distinct from $|2\diamond|$. Thus, the existence of particular features could not be learned. As they say:

It should now be obvious that we are heading toward the conclusion that children must “know” (i.e. have innate access to) the set of phonological features used in all of the languages of the world. This is how the IofPP will be extended in this chapter; but it is equally clear that the same conclusion holds for primitive operators like the AND and OR of card languages, or whatever are the operators of real grammars (in both phonology and syntax). (Hale and Reiss 2008: 38)

H&R do not actually claim that the set of innate phonological features has to be phonetically defined; instead, there has to be a set of features available in UG . The feature [Coronal] must by this argument be available in UG , but it need not have anything to do with the tongue. Other assumptions (see Hale et al. 2007) could have the consequence that features have their traditional phonetic consequences.


¹¹An alternative would be that J, Q, K, A are parsed as $[\#]$. A detailed study of human perception of parrot “speech” might shed light on how humans parse stimuli that are well outside the norms of human production, but clearly parsing sounds as speech is not very strict.

¹²In their footnote 3, H&R say that only one of these suit features can characterize any given card, thus $[\clubsuit\diamond]$ or $[\clubsuit\spadesuit]$ are not considered. This could either be because of the nature of the physical inputs, or because this is a property of all versions of card UG . For our purposes it does not matter what restricts inputs, but in language it obviously matters substantially.


¹³Note that in the printed version of this article, hearts and diamonds are ‘grey’ but the reader should consider them ‘red’, as they are in a real-life deck of cards.

3.2 What parses?

The concepts of representation and input relied on above require a cognitive domain; since there are no absolute inputs, there are inputs *to* something and representations *in* something. In saying that a sound is parsed, we mean that it is assigned a featural representation in phonology.

An intended speech sound enters the human body, is mechanically transduced to a pressure wave in the cochlea, causing neuronal excitation which is the neural basis for creating a first representation of a sound. This signal is parsed into other representations, proceeding to the auditory cortex, and perhaps ultimately, to the phonological component. In cognitively processing an instance of *s*, *B#*, or  (the sound of a car crash), different representations are created, but at early stages of processing, the same kind of representation is created – a raw acoustic image (the structure of the ear does not sort out whether a sound is linguistic). The physical sound $|s|$ may be interpreted as a language sound, but it can also be interpreted as escaping steam, or as a person or other thing making a snake noise. It is an empirical question exactly how this happens: what is clear is that UG does not directly convert the pattern of neuronal excitation from the cochlea into phonological features. The fundamental question for understanding how physical sound maps to phonological representations is: What is the sequence of representation-to-representation mappings that takes place prior to phonology?

Modular theories of grammar typically hold that each mental module has a distinct set of primitives. As a mental object passes through the various modules, it is subject to that many symbol-to-symbol translations: these translation devices are called interfaces. To clarify the consequences of interfaces for the Card Grammar argument, I set forth the Revised Card Grammar argument, which introduces multiple modules and interfaces in a theory of mind, MT_1 , demonstrating that phonological features can be learned. Assume the following physical inputs:



Perhaps because of the nature of the sensory apparatus, MT_1 does not process differences in card size at all, so in transducing physical inputs to a first mental form, the difference between $|E|$ and the remaining inputs is irretrievably lost. This is analogous to the loss of information in an acoustic signal above or below certain frequencies. This physical transduction provides a representation in module Mo_1 where $|A-D|$ are all distinct, and results in $in [_{Mo_1}J♣], [_{Mo_1}J♣], [_{Mo_1}♣J],$ and $[_{Mo_1}♣♣]$ respectively. $|E|$ is not represented distinctly in Mo_1 : we will say that only the size aspect of the signal is lost so it is parsed as $[_{Mo_1}J♣]$, but it is possible that an input is entirely rejected. After the computations of Mo_1 are performed, the output is

passed through interface I_{1-2} which converts and discards some information, in this case the typeface difference, with the result that Mo_2 receives $[_{Mo_2}J\clubsuit] \leftarrow [A]$, $[_{Mo_2}J\spadesuit] \leftarrow [B,D,E]$, and $[_{Mo_2}\clubsuit J] \leftarrow [C]$. With respect to language, this discarding of earlier information is similar to discarded spatial information in auditory processing. Although we can hear that “hat” uttered close to the right ear is not the same as “hat” uttered close to the left ear, that difference is thrown away on the path to grammatical computation.

Since our goal is to understand how features could be learned, we now consider the case where some aspect of interface I_{2-3} between Mo_2 and Mo_3 depends on experience: the mapping may be learned. Learning constitutes the postulation of a hypothesis as to the nature of the system which operates on the relevant data. As new data become available, hypotheses may be reinforced, or they may be subject to correction when the hypothesis is contradicted by known facts. To draw a phonological analog, a hypothesis could be initially advanced by a German-learning child, based on known facts, that the German word [bunt] ‘federal’ is underlying /bunt/. That hypothesis will eventually be overridden in light of complicated facts regarding the German word [bunt] meaning “colourful”, and other forms of the words ‘federal’ and ‘colourful’. Eventually the child corrects the system of rules and representations such that ‘federal’ is /bund/ and there is a devoicing rule.

Based on initial data, a child might postulate that $[_{Mo_2}J\clubsuit] \rightarrow [_{Mo_3}J\clubsuit]$ and $[_{Mo_2}\clubsuit J] \rightarrow [_{Mo_3}\clubsuit J]$, and also that $[_{Mo_2}J\spadesuit] \rightarrow [_{Mo_3}J\spadesuit]$, that is, the distinction between $[_{Mo_2}J\clubsuit]$ and $[_{Mo_2}J\spadesuit]$ may be eliminated. At this point, Mo_3 contains only $[_{Mo_3}J\clubsuit]$ and $[_{Mo_3}\clubsuit J]$: but the basis for retrieving the third distinction still exists in Mo_2 . Subsequent experience can establish the incorrectness of the interface mapping; for example it might be discovered that computations in Mo_3 treat supposed instances of $[_{Mo_3}J\clubsuit]$ differently, depending on whether they derive from $[_{Mo_2}J\clubsuit]$ or $[_{Mo_2}J\spadesuit]$. The mind still stores a distinct representation of [A], namely $[_{Mo_2}J\clubsuit]$, and the interface mapping can easily be corrected so that $[_{Mo_2}J\clubsuit] \rightarrow [_{Mo_3}J\clubsuit]$. Although computation within Mo_3 is limited to the primitives of Mo_3 , learning about the interface from Mo_2 to Mo_3 is not an operation in Mo_3 . The mind which is still learning the rules still has access to information that was erroneously discarded. The key to solving the Card Grammar problem is that the phonological component of a grammar does not learn; it is that the mind learns about the phonological component.

H&R use the Card Grammar argument to claim that the traditional Subset Principle of learning theory is wrong. Given the Card Grammar argument, it would be impossible to correct the hypothesis that a language only has the vowels {i, a, u} to the hypothesis that the language has {i, i, ε, a, ɔ, u, u}. If parsing of inputs is absolutely limited to just {i, a, u}, then a child could not gain awareness that {i, ε, ɔ, u} also exist. As shown above, a child can learn that the inventory {i, a, u} is an error, since phonological feature assignment is not performed by the cochlear nucleus.

3.3 What part of the mind would have universal features?

The card-grammar argument is not specifically about phonological features; it is presumptively an argument about the broad language faculty. Even though the language

faculty must have some primitive properties that form the basis for learning, we cannot conclude that the relevant primitives are specifically a list of features, as standardly envisioned in phonological theory. Still, H&R propose that universal features are an innate part of UG, which raises two important questions: what is UG, and where in UG would these claimed universal features exist?

3.3.1 *What is UG?*

Chomsky (1965) advances the concept of UG as an architectural mechanism regulating the operation of specific grammars:

The grammar of a particular language, then, is to be supplemented by a universal grammar that accommodates the creative aspect of language use and expresses the deep-seated regularities which, being universal, are omitted from the grammar itself. Therefore it is quite proper for a grammar to discuss only exceptions and irregularities in any detail. It is only when supplemented by a universal grammar that the grammar of a language provides a full account of the speaker-hearer's competence. (Chomsky 1965: 6)

The historical source of the concept UG is rationalist philosophy of the Renaissance period, which Chomsky is attempting to put on a cognitive footing. Chomsky and Halle (1968) also characterize UG as follows:

A universal grammar is a system of conditions that characterize any human language, a theory of essential properties of human language. It is reasonable to suppose that the principle of the transformational cycle and the principles of organization of grammar that we have formulated in terms of certain notational conventions are, if correct, a part of universal grammar rather than of the particular grammar of English. (Chomsky and Halle 1968: 43)

Roberts (2016) summarizes the theory of UG as “the scientific theory of the genetic component of the language faculty”, that it is “the theory of that feature of the genetically given human cognitive capacity which makes language possible, and at the same time defines a possible human language”. This architectural view is the one which I assume.

Hale and Reiss (2008: 2) state an alternative view: “Once we accept the existence of a language faculty, it is also uncontroversial that this faculty has an initial state, before any experience with a particular language. Under this view Universal Grammar, the theory of this initial state, is a topic of study, not a hypothesis”.

This follows the view of Chomsky (1980: 7), seeing UG to be the initial state of the child or the language faculty: “In a highly idealized picture of language acquisition, UG is taken to be a characterization of the child's pre-linguistic initial state”. And, later: “These and many other questions must be considered in the development of a comprehensive theory of UG, as a characterization of the initial state of the language faculty” (Chomsky 1980: 138).

Under this view, it is hard to see how UG has potency for regulating computation of linguistic forms, or how UG could have an effect on an already quadrilingual child learning a fifth language.¹⁴ Obviously, these radically different views of the nature of

¹⁴Experience with language begins *in utero*, starting at around 30 weeks. It remains completely unclear at what stage the brain has developed to the point that the language organ is

UG substantially affect arguments about whether UG must contain pre-specified features. Especially relevant to the debate is the extent to which a special language-specific theory of learning is necessary, or can the universal properties of language be explained by the interaction between the fixed architecture of grammar and general mechanisms of automatic learning such as employed in learning eating, walking or visual tracking. As proposed here, the UG theory of grammar specifies what are possible representations and computations in the phonological component. Language acquisition results from such general learning strategies, which are set to the task of inducing a set of UG-consistent computations and representations given the primary linguistic data. The contribution of UG to language acquisition thus resides in prescribing the general nature of the mental object that is to be acquired, and not what the set of possible mental objects are. Even if there exist specific learning strategies for language, it does not follow that the learning device needs to be pre-coded with the specific features of a language. What needs to be known is simply that sounds are represented and operated on in phonology using features. The features resulting from learning must integrate with the grammar that is also being learned.

3.3.2 *Where, if anywhere, are the universal features?*

Even if universal substantive sound properties are encoded somewhere in UG, it is an open question what those properties are, and how they relate to distinctive features as they exist in phonology. It is possible that the linguistic phonetic component contains a collection of genetically predetermined primitives, which influence what objects are presented to the phonological component, thus indirectly determining how physical sounds map to feature matrices. For example, language sounds in the linguistic phonetic component could be single symbols analogous to IPA letters, and the theory of linguistic phonetic sounds might reduce to being whatever expressions can be constructed in the IPA – for example, $[\text{phonetic } i \text{ } i \text{ } \bar{i} \text{ } \dot{i}]$. If this is the alphabet with which phonetic forms are represented and how phonetic computations are carried out, we would have a basis for postulating defeasible interface hypotheses about what phonology receives. A child could, based on evidence, learn that $[\text{phonetic } \bar{i} \text{ } \dot{i}]$ do behave phonologically like distinct objects, and could therefore recover from the error of assuming that $[\text{phonetic } \bar{i} \text{ } \dot{i}]$ both map to $[\text{phonological } i]$. A child could then learn a phonological contrast between tense and lax vowels, even if that hypothesis had been previously rejected.

In order to argue about the nature of the phonological component based on the phonetic component, we need a theory of that component as a computational device, and obviously we need to determine whether there even exists such a component. See Samuels et al. (2022) for discussion of language-specific phonetics, especially the fact that the Gomera dialect of Spanish is a single language with just one phonological grammar but two modes of production, spoken and whistled, which require

fully developed: arbitrarily, I assume that the language faculty is fully formed prior to 30 weeks.

two distinct systems of transduction to body outputs. See Hamann (2011) for an overview of issues regarding the relation of phonetics to phonology. As observed by Hamann (2011: 203), “the nature of the interface cannot be unearthed by experimental studies alone. It depends to a considerable part on the theoretical assumptions we make, and on the aim we have in mind with our phonological and phonetic descriptions”.

If devoiced and underlyingly voiceless obstruents in German and Dutch are physically or perceptually distinct, this fact has no implications for the theory of phonological computation in the event that final devoicing is a phonetic rule, but has substantial implications for phonology if this is a phonological rule.

The answer to the question of what linguistic primitives are required for phonetic computations depends on having a theory of phonetic computation (not a taxonomy of phonetic versus phonological phenomena), and that theory in turn depends on determining whether a phonetic component exists, a matter to be resolved based on how the assumption solves what would otherwise be severe problems for the theory of phonological computation and representation. Chomsky and Halle (1968) deny that there is a component of phonetic computation, holding that phonetic and phonological features and computations are indistinguishable. The SPE approach to phonetics reduces all language-specific differences in phonetic properties to either non-contrastive features (for instance [suction], invoked for phonetic differences in the production of labiovelars between Guang languages such as Late versus those in Yoruba and Ibibio (see Ladefoged 1964)), or to rules assigning integer values to a particular feature, where in SPE $[1F_i]$ is the maximum degree of property F_i and $[nF_i]$ for some value n substantially greater than 1 is the smallest degree of the property, total lack. In that framework, the slight difference between “phonetics” and phonology is (by hypothesis) that features do not have whole number values in the lexicon, they only have the values $\{u, m, +, -, 0\}$, and phonological outputs (which are directly interpreted by some motor control device) always have numeric feature coefficients assigned to features.¹⁵ The consequences of this approach for the theory of phonological computation are not trivial to assess, and the literature on phonetic interpretation *qua* phonology within the “no-phonetics” tradition is not extensive.¹⁶ It is clear that within this approach, the theory of phonological computations would have to be expanded to include arithmetic operations and integer or fractional variables, which have no phonological justification.

Ascertaining whether there is a component of linguistic phonetics is necessary, so that we can know what constitutes the input to the transducer yielding phonological representations as outputs. Indeed, we need a theory of transducers into grammatical components. One possibility, very similar to the SPE program, is suggested in Hale et al. (2007: 647) where there are two transducers between phonology and

¹⁵Postal (1968: 61–62) holds that representations automatically map to fully specified numeric forms by virtue of universal interpretive principles, thus some number is always associated with “+” and “–”.

¹⁶The primary works within that tradition are King (1969), Fromkin (1972), Peters (1973), Anderson (1974), Johnson (1975) and Clifton (1976).

lower-level body functions, namely a “transduction of features (the input) to some gestural score (...) transduction of a percept (the input) to features (the output)”, where “these two transducers are innate and invariant – they are identical in all humans (barring some specific neurological impairment) and do not change over time or experience (i.e., they do not ‘learn’)”. They propose that “we can consider the transduction process, too, as invariant in that the *relationship* or *mapping* between a particular feature bundle and a particular gestural score is a deterministic (and thus consistent) conversion process and, similarly, that the relationship or mapping of a particular perceptual input to a feature bundle is deterministic”.

Additionally, they propose that transductions can be context-sensitive. Clearly, there is a substantial difference of opinion regarding the nature of phonological transductions. In the perspective advanced here, transductions between grammatical components are learned, and are also probably context-free symbol-replacement lists.

We cannot expect to resolve all questions about the theory of transduction into/from grammar or within grammar here. We can, however, ask what kinds of facts could be relevant to answering these questions. The main question of interest is whether there is a linguistic component of phonetics. The potential evidence for a linguistic phonetic component falls into three categories. First, phonological features refer to ranges of physical facts, not precise measurements of such facts. The range of facts related to a feature may depend on which language the feature is instantiated in. For example, languages with the feature [Spread Glottis] on consonants differ in how a segment is realized when specified or not specified with that feature, via language-specific target voice onset time (VOT) values. Second, the realization of a feature may vary in a language-specific way that depends on surrounding context. For example, F_0 values of a H or L tone may be adjusted upwards or downwards from a target value, depending on surrounding tonal context (e.g., H before L may be subject to a language-specific process of pitch-raising). Such cross-linguistic variation is extremely informative, since it provides a basis for constructing a theory of phonetic computations, which allows us to evaluate the adequacy of competing theories of phonetic grammar. Finally, the time course of the realization of a feature may be language-specific, thus lip protrusion associated with a rounded consonant might, on a language-specific basis, be timed close to the release of a consonant, or it might be timed earlier and be easily detectable on a preceding segment. Below we consider specific facts that lend *prima facie* credibility to the claim that there is a language-specific component of phonetics. Needless to say, the argument cannot be properly evaluated without a theory of phonological computations, since the alternative hypothesis (as set forth in SPE) is that by definition, all linguistic variation comes from phonological rules, so the theory of phonology must include whatever devices are necessary to enable such computations.

As an example of language-specific degree in the realization of a feature, consider how the feature [Spread Glottis] (the feature underlying the phenomenon of aspiration) is realized as differences in voice onset time, since this is a well-known example of variation between languages. Cho and Ladefoged (1999: 216–217) document an example of degree-variation between languages in the voice

onset time lag which implements an aspiration contrast, in languages with a phonemic contrast :

(4)	[k]	[k ^h]	N speakers
	<i>msc</i>	<i>msc</i>	
Jalapa Mazatec	23	80	6
Gaelic	28	73	11
Apache	31	80	8
Khonama Angami	20	91	6
Tlingit	28	128	4
Hupa	44	84	3
Navaho	45	154	7

As Cho and Ladefoged (1999: 214) note, crosslinguistic comparison of VOT requires “a body of data from a number of widely different languages, all of which have been collected and analyzed in the same way”. There is no information on token variability within and across these languages, so it is impossible to know whether the measured mean VOT of 23 msc in Jalapa Mazatec [k] is statistically different from the 28 msc of Tlingit. Based on the magnitude of differences, it is reasonable to conclude that [k] falls into at least two different subtypes, with Khonama Angami having the shortest VOT and Navaho having the longest, and [k^h] falls into at least three subtypes, with Gaelic having the shortest VOT, Navaho having the longest, and Tlingit being in the middle – though Khonama Angami probably represents a fourth type. Within the SPE model of phonetic implementation, this suggests language-specific integer targets along the following lines, where, for example, Navaho [k^h] is assigned the value [1spr.gl.] and Gaelic [k^h] is [4spr.gl.].¹⁷

(5)	[k]	[k ^h]
Jalapa Mazatec	7	4
Gaelic	7	4
Apache	7	4
Khonama Angami	7	3
Tlingit	7	2
Hupa	6	4
Navaho	6	1

These phonetic differences between languages are fodder for a theory of phonetic computation. Similar language differences in vowel realization are well-known; see for example Disner (1983), who finds that [i] in German and Norwegian differ in that [i] is higher in German (has lower F₁) – this result holds for bilingual speakers, and depends on which language the individual is speaking. See Vaux and Samuels (2015) for further discussion of language-specific segment target differences and the problem which they pose for dispersion theory.

Contextual differences in phonetic realization of segments can also be language-specific, though this issue has not been as well studied. Lebanese Arabic, Logoori,

¹⁷The lack of [5spr.gl.] corresponds to the lack of a language, in the above subset, having VOT in the neighborhood of 50-70 msc VOT for velars. Cho and Ladefoged indicate a mean VOT of 56 msc for Yapese, which has no aspiration contrast.

Finnish and North Saami all have a four-way phonological contrast in segment length between VCV, VC:V, V:CV and V:C:V structures. All four languages differ in how such length in these prosodic subclasses is realized in terms of segment duration. Data from North Saami and Logoori come from my own research. Lebanese Arabic data are presented in Khattab and Al-Tamimi (2014), who give mean durations in the four relevant contexts, along with information on the statistical significance of different durational means. Data for Finnish derives from Dunn (1993), who, in her Appendix 2, presents numeric data separated according to speaker, an additional extraneous factor of V_2 length, and distinguishing /p/ from /m/. The Finnish duration values in example (6) are the mean of reported V and C durations within the four-way prosodic subclasses studied here, averaging across speakers, segmental differences in C, and disregarding vowel length in σ_2 .

For example, phonological vowel length in Finnish is realized with much more prolongation, compared to what is found in Logoori. In addition, the languages differ in how durations are determined in V:C:V. In Logoori, there is no interaction between consonant and vowel duration, whereas Finnish and North Saami have a durational trade-off where neither vowels nor consonants in V:C:V are prolonged as much as expected.

(6) <i>Language</i>	<i>Prosodic context</i>	<i>V duration</i>	<i>C duration</i>
Lebanese Arabic	VC	78	84
	VC:	77	182
	V:C	166	99
	V:C:	149	181
Logoori	VC	76	99
	VC:	81	201
	V:C	146	97
	V:C:	140	217
Finnish	VC	109	95
	VC:	121	227
	V:C	240	97
	V:C:	215	178
North Saami	VC	79	73
	VC:	76	235
	V:C	169	90
	V:C:	148	181

Taking the context before __CV to best reveal the duration target for long versus short vowels, and the context V__V to best reveal the duration target for long versus short consonants, we arrive at the following differences between languages in the long-to-short ratio.

(7)	V:/V	C:/C
Lebanese Arabic	2.13	2.17
Logoori	1.92	2.03
Finnish	2.20	2.39
North Saami	2.14	3.22

Arabic and North Saami have essentially the same degree of prolongation associated with vowel length. The increase in Finnish long vowels is more than that of

Arabic and North Saami, and the increase in Logoori vowels is less than in the other languages. Similarly, North Saami prolongation of long consonants is most substantial, and Logoori is least substantial. These differences can be expressed as a cross-linguistically variable ratio determining how much longer long vowels or consonants are, compared to short vowels and consonants.

Beyond differences in target duration associated with contrastive length for vowels or consonants, there are also language-specific contextual differences in how consonant and vowel duration is computed in connection with phonological length. One pattern found in these data (and in comparable data from other studies) is that long vowels have greater duration before a single consonant, and shorter duration before a geminate. A second effect is that long consonants have greater duration after V than they do after V:. The degree of shortening of long vowels and consonants next to long consonants and vowels depends on the language. Example (8) gives the degree of shortening of long vowels and consonants associated with adjacent long segments. The first column is the duration of V: before C: divided by the duration of V: before C in the language, and the second column is the duration of C: after V: divided by the duration of C: after V in that language (in nanoseconds).

(8)	V: c: /_C	C: v: / v_
Lebanese Arabic	0.90	0.99
Logoori	0.96	1.08
Finnish	0.90	0.78
North Saami	0.88	0.77

The Logoori pattern is the simplest: segment length is realized as simple doubling of a segment's duration target, and there is no significant interaction between V: or C: duration as a function of following C-length or preceding V-length, respectively. The Arabic pattern is similar to that of Logoori, but a long vowel is somewhat shorter before C: than it is before C, and that difference is statistically significant.¹⁸ Contextual differences for both V-length and C-length in Finnish and North Saami are statistically significant. Most obviously, duration of long consonants in those two languages is strongly influenced by the length of the preceding vowel – a long consonant has only about 3/4 of its expected duration, based on the duration after a short vowel.

3.4 What segments does phonology receive, and from where?

The present theory of feature learning assumes that children mentally categorize tokens of sounds in some manner, that this categorization is outside of phonology, and from grammar-external auditory representation, a system of computations and a feature analysis of those sounds is possible. In other words, the model takes segments to be logically fundamental in phonological acquisition, and perhaps are

¹⁸In Logoori, the duration differences between long consonants associated with long-V versus short-V context are not statistically significant; likewise the differences between long vowels associated with long-C versus short-C context are not significant. Khattab and Al-Tamimi (2014) report that the difference degree of vowel shortening associated with following C: is statistically significant.

perceptual primitives, but does not treat the phonological analysis of those segments as a pre-phonological fundamental. The model in section 4 is based on the logic that if segments $\{s_1 \dots s_i\}$ computationally group together and $\{s_{i+1} \dots s_j\}$ are excluded, a feature expression selects $\{s_1 \dots s_i\}$ but not $\{s_{i+1} \dots s_j\}$. The child must therefore know what the segments of the language are, and especially which ones require a distinct representation in the phonology. In hypothesizing a feature assignment based on a phonological rule where b, d, g become p, t, k before an obstruent, it is highly relevant to know if any of d^z, z, b^h, p^h also exist in the language. Feature learning presupposes prior analysis of the stream of speech – in German, a child must have knowledge of inflectionally-related words like [bunt], [bundə] to know that there is final devoicing. It is insufficient to experience acoustic waveforms: there must already be analysis of continuous speech into segments.

Segmentation of continuous physical sound has two aspects: reduction to a sequence of discrete mental units, and analysis of units into types. Discretization and categorization are not specifically linguistic functions. To understand the pre-grammatical prerequisites for phonological analysis, we must know how a child learns that physical signals {I,II,III} are discretized into segmental sequences $\{(a_1, a_2, a_3), (b_1, b_2), (c_1, c_2, c_3), \dots\}$. The ability to mentally isolate an aspect of continuous reality and treat it as “a thing” is pervasive in human cognition, and is not specific to language, so is not a particular aspect of UG. The fact that continuous signals are reduced to sequences of units is necessitated by the fact that grammar is the manipulation of discrete units; but discretization itself is necessary for other forms of auditory perception and for visual perception. What may be uniquely linguistic is the specific unit that the mind parses – especially given the plausible choices of syllable versus segment as “most basic percept”.

A continuous signal must obviously be physically converted from a form of external energy to something in the mind, whose nature is determined by the sensory apparatus (e.g., the inability of human eyes to detect x-rays, the inaudibility of a 5 Hz sound at normal, safe amplitudes). Physical mechanisms within the head (including the outer ear) cause inner hair cells of the cochlea to transduce physical sound to a pattern of electrical impulses, resulting in a tonotopic map which is interpreted in the primary auditory cortex. Much cognitive processing takes place, the end point of interest for us being that the continuous signal is converted into a series of things and relationships between things.

We also need to know how a child determines that discrete tokens $\{a_1, b_1, c_1 \dots\}$ are subsumed under one category X, but tokens $\{a_2, b_2, c_2 \dots\}$ are subsumed under a separate category Y. For language sounds, at the acquisitional state prior to devising a treatment in a specific grammatical component, that means at the coarsest level learning whether d and \underline{d} are both part of the language. Categorization of things means discerning that parsed tokens $\{a, b\}$ are similar in a perceivable way, and are distinguishable from $\{c\}$ in that same way. A sine wave at 100 Hz is similar to a square wave at 100 Hz, and distinct from a sine wave at 105 Hz. A sine wave at 100 Hz is, likewise, similar to a sine wave at 105 Hz, and distinct in that respect from a square or triangular wave at 100 Hz. A sine wave at 10,000 Hz might not, on the other hand, be perceptibly distinguishable from a sine wave at 10,005 Hz. Speech-

sound categorization is carried out at a much higher level than is involved in detecting differences and similarities between constructed signals differing in F_0 and amplitude distributions; indeed, spectral properties rarely relate to categories (conventional groupings with an attached label such as “brassy”), and frequency properties relate to labels not available to most people (“middle C”, “high C”, “C”, “B sharp”). Though we don’t know *how* it happens, we know that it does happen: that children learn that a certain range of physical sounds are “the same thing” in their language.


The aspect of sound-to-symbol conversion most relevant to phonology is that which is specifically about language. While it might be interesting to know how people can learn to identify different musical instruments or notes based on sound, this is not the same task as learning speech segments. Starting from the assumption (apparently made by Hale et al.) that humans can perceive very subtle distinctions in speech sounds, learning to correctly categorize speech sounds involves learning which differences are inconsequential, and which ones are important. The range of variation in the phonetic properties of segments observed in a collection of tokens of the word ‘cat’ is an example of inconsequential difference. So too, probably, is the difference observed in a collection of tokens from a number of speakers, at least to the extent that they are speaking the same language. Such variation has in common that it is not observed to correlate with anything linguistically relevant. A case where variation in sound realization might correlate with a property of the grammar is when it systematically correlates with other sounds – the appearance of perceptible [r³] before [q] and [i] anywhere else could signal a rule-governed distinction. It could also signal a nonlinguistic physical necessity imposed by the nature of the articulatory gestures required to utter [i] followed by [q]. This raises the question of what a child’s genetically-dictated “knowledge” of anatomy is and what are the acoustic consequences of that anatomy. Do we know, in advance of experience, that it takes N milliseconds to move the tongue from point A to B, for all possible points? Alternatively, do we learn this in the course of babbling, by observing that attempts to produce perfect [iq] always result in something like [r³q]? Or, do we learn that it is possible to produce both [r³q] and [iq], but in the language of the environment, one never encounters [iq] even though one encounters [ik]?

If a child has learned that a fact about a language sound is linguistically relevant (is not an unavoidable physical requirement of doing something else), this does not say whether it is expressed in the phonetic grammar, as opposed to being expressed in the phonological grammar. A child is in a privileged position to answer the question, because it is not prejudiced by the presumptions of a transcription, which linguists must overcome. In reporting that a word is pronounced [r³q], linguists build a phonological analysis into the lowest level of data reporting, rather than using a non-linguistic system of tongue-movement notation. A classical example of building phonological analysis into data reporting is the case of Marshallese vowels. Marshallese has been said to have 12 pure vowels and 24 diphthongs at the phonetic level. Research leading up to Bender (1968) revealed a rich system of vowels and consonants, including previously non-obvious consonant qualities (rounding and palatality) and correlations between vowel and consonant quality. Bender (1968) sets forth a phonological analysis which reduces the set of phonological vowels

from 12 to 4 – a set of central vowels differing in height, which are neutral for frontness/backness and roundness. The richer set of apparent surface vowels such as [o] can be treated as arising from a phonetic process interpreting vowel phonemes specified only for height. Bender (1968) explicitly recognized the problem of earlier impressionistic transcriptions like /jok/, /koj/, stating that relistening to an item which had been transcribed as /jok/ revealed the actual quality to move from front to back with increasing rounding, all at mid height: [t^{yɛə}ok^w].¹⁹ Similarly, /koj/ came to be perceived as [k^{wə}ɛt^y]. And as the phonetic facts of other mixed environments were reexamined, each proved to be capable of similar interpretation as resulting from competing consonantal influences on a less fully specified vowel. A benefit of attributing some surface vowel qualities to rules involving surrounding consonants is a simplification in rules for computing the surface forms of affixes, as noted by Bender.

Choi (1992) investigates this question acoustically, arguing that there is a continuous interpolation from one consonantal articulatory state to another in cases like [j^{ɛə}ʌt^y] ‘well-sifted’, [p^{yə}ʌo^kw] ‘wet’. Apparent short monophthongs like [o, e] only appear between consonants of the same secondary articulation (all consonants are analyzed as palatalized, rounded, or velarized) (e.g., [r^wor^w] ‘bark’, [t^ɛp^j] ‘cheek’, [p^yʌp^y] ‘triggerfish’). Choi posits that vowels are unspecified for the features [Palatal] and [Velar] (and presumably [Round]), and remain unspecified going into phonetic interpretation (it is a common assumption in current feature theories that there does not exist a set of universal default rules assigning values for all unspecified features). The surface variation in realization of phonological /ə/ results from interpolation between consonantal targets. In the phonetic component, vowels are assigned a target for F₁, which instantiates the phonological height specification, but they have no F₂ target. F₂ instead derives by a phonetic interpolation function between C₁ and C₂, which are the bearers of F₂ targets. When the flanking consonants are of the same vocalic type, the interpolation function returns a constant F₂ value for all times between C₁ and C₂, but when the consonants differ, the function returns continuously varying values of F₂. This continuously varying F₂ path is often discretized in linguistic transcriptions as a sequence of micro-vowels, in forms like [j^{ɛə}ʌt^y] = /l^jət^y/.

The relevant acquisitional question is, what are possible forms for a child to contemplate as the output of phonology, given the classes of physical things to be modeled? There is, arguably, only one reasonable possibility for “j^{ɛə}ʌt^y”: the phonology produces [l^jət^y]. The alternative of /l^jət^y/ → [j^{ɛə}ʌt^y] is factually arbitrary in a manner not supported by any consideration. As a symbolization of phonetic fact, it is arbitrarily imprecise, skipping many intermediate steps in the vocalic continuum which suffer no pre-theoretical disadvantage. A more accurate symbolization of

¹⁹I use the symbols employed by the data sources (Choi 1992, else Bender 1968). Constructed forms like [j^{ɛə}ʌt^y] follow standard IPA interpretations. The interpretation of letters like [i] depends on the theory of features – my interpretation is that it is a high vowel which is neither front nor back, and is not specified as round. Hale and Reiss (2008) assume binary features, and so employ the symbol  for this vowel.

pronunciation would be something like $[j^{iic}_{\emptyset}{}^{aiu}t^y]$. Either alternative faces serious problems in receiving a coherent phonological representation – what system of computation maps $/l^j_{\emptyset}t^y/$ to $[j^{iic}_{\emptyset}{}^{aiu}t^y]$? Letter-strings like $[j^{\emptyset}{}^{\wedge}t^y]$ or $[j^{iic}_{\emptyset}{}^{aiu}t^y]$ are phonologically incoherent, without a theory of the “micro-segments” represented by raised letters or the even finer-subdivided “nano-segments” in $[j^{iic}_{\emptyset}{}^{aiu}t^y]$. As emphasized above, we need a model of phonetic grammar to evaluate such claims. Given a phonetic model such as set forth by Choi, there *is* a theory of phonetic interpretation and a model of Marshallese phonetic grammar that generates correct physical outputs, operating on a phonological output $[l^j_{\emptyset}t^y]$.

In the case of $[r^w{}^{or^w}]$, when combined with the independently necessary place-interpolation rule of phonetics, either $[r^w{}^{er^w}]$ or $[r^w{}^{or^w}]$ *qua* phonological output can be credibly related to the physical facts (the possibility that rounding of the vowel derives from interpolation between consonants does not preclude the possibility that rounding is also present in the input). When no phonetic fact dictates which of two (or more) phonetic forms are the output of the phonology, the resolution of the question must fall to inspection of the resulting systems of phonetics and phonology. Whichever system is the simplest, that is the system learned, under the premise of FP. Obviously, we need much more information about Marshallese phonology, in order to advance any argument about whether it is simpler to assume $[r^w{}^{er^w}]$ or $[r^w{}^{or^w}]$.²⁰

There are also two *prima facie* plausible accounts of certain patterns of English consonant allophony: perhaps they result from late phonological rules, or perhaps they are from processes of phonetic implementation. Even assuming that underlying representations do not contain all of $/p, b, p^h/$ (a claim that cannot be stipulated arbitrarily), their systematic presence in pre-physical mental representations is not in serious doubt. Is there a phonological rule assigning aspiration, or is it part of phonetic implementation? What kind of process derives $[ʔ]$ ($[hɪʔ]$ ‘hit’, $[kaʔŋ]$ ‘cotton’, in some dialects) or $[r]$ (*write* ~ *writing* vs. *ride* ~ *riding*, $[r^w{}^{aiŋ}]$ in both cases). If these processes are not phonological, what does that entail about the nature of phonetic computations: Can phonetic processes perceptually neutralize a distinction between words? Apparently, phonetic implementation must be able to refer to syllable or foot structure in a phonetic account of consonant allophony, which is not a strikingly controversial claim. Is there some logical principle regarding the nature of grammars that dictates that aspiration should *not* be treated phonologically, or phonetically? This is the kind of knowledge required to decide what segments are present in phonological output. Different knowledge – phonological knowledge – is required to figure out whether certain segments are missing from underlying forms.

To summarize this section, the argument that phonological features are logically unlearnable and must be built into UG depends on a number of assumptions that have not been established. It relies on the claim that phonological features are provided by an invariant direct interface with non-linguistic articulatory and

²⁰An even less phonological account is suggested by Hale and Reiss (2008), to the effect that grammar does not operate on $/l^j_{\emptyset}t^y/$ and $/r^w{}^{er^w}/$ at all. It is beyond the scope of this article to explore what theory of phonetic implementation might underly their proposal, since Hale et al. (2007) also do not have a phonetic component.

perceptual systems. This precludes the existence of a linguistic component of phonetic interpretation. The consequences of that move for the theory of phonological computation and representation are quite significant, since this amounts to expanding the scope of phonology in a manner that renders phonology less coherent, not more coherent.

4. HOW ARE FEATURES LEARNED?

It does not suffice to just assert that the features of a language are learned. The purpose of this section is to show how feature-learning proceeds. In the first subsection, I sketch the basic logic of feature acquisition using a simple constructed phonology. The second subsection gives an account of feature learning in the Bantu language Kerewe.

4.1 Introduction

The greatest challenge for showing that features need not be innate and can be learned is that there is little by way of explicit logical framework for discussing the acquisition of phonological grammars (see Hale and Reiss 2008) for extensive discussion of this point.

Any language has a set of segments $A = \{s_x \dots s_z\}$, where each segment is instantiated in the grammar with a distinct feature matrix M_i which is the structured conjunction of features $\{F_i, F_j \dots F_n\}$. The fact that a segment $[s_x]$ exists in a language is sufficient reason for it to be assigned some set of features which renders it distinct from any other type of segment in the language. The specific features posited for a segment are motivated by how the segment functions in phonological computations. I assume a theory of learning where assignment of features to segments is made in such a way that the set of feature specifications in the rule system is the simplest possible, and the set of features invoked for the language is minimal. More generally, I assume the simplicity-driven model of phonological theorizing: Formal Phonology (Odden 2013). The theory of UG predetermines what is a possible computational system for generating the data, and the theory of learning tells us how (and to what extent) the set of possible grammars is further narrowed down – the simplest system is the one that is learned.

Take a simple hypothetical language, with the segments $p, t, k, m, n, \eta, i, a, u$. This language has the following phonological fact, which motivates a rule that is to be learned.²¹


$$(9) \{p, t, k\}_n \rightarrow \{m, n, \eta\}_n / _ \{m, n, \eta\}$$

²¹In shorthand notations for factual generalizations that a rule expresses, $\{a, b, c\}_i$ is always paired with at least one other instance of $\{e, f, g\}_i$, and means “ a in the context $e, b \dots f, c \dots g$, respectively”, but $\{a, b, c\}$ means “any of a, b , or c ”. In example (9), p, t, k become respectively m, n, η as notated with subscripts, but the trigger is any of m, n, η , notated by the lack of subscript.

The child has awareness of the facts represented as (9), and automatic learning creates a rule which has the effect that when the sounds *p, t, k* stand before any of *m, n, η*, the former become respectively *m, n, η*. Knowledge such as (9) is outside of grammar (grammars do not contain the data that they generate): that factual knowledge is the basis on which a grammatical rule like (10) can be posited. The learning algorithm may therefore evaluate the hypothesis that the grammar has the following rule:

$$(10) [F_i] \rightarrow [F_j] / _ [F_j]$$

If (10) is correct, it follows that *p, t, k* are $[F_i]$ and *m, n, η* are $[F_j]$. Expression (10) is a possible rule only given a certain view of “rule” as provided by UG – the SPE theory of rules. It is not difficult to recast (10) as a feature-spreading rule in the privative autosegmental theory assumed here. The same fact pattern would motivate (11), which inserts an association relation between two segments in case the second has the property $[F_i]$ and the first has the property $[F_j]$




There are other UG-consistent statements of the rule that could describe the data. In the SPE model, one might posit the following:

$$(12) \text{ a. } [F_i] \rightarrow [F_j] / _ \left[\begin{smallmatrix} F_i \\ F_k \end{smallmatrix} \right]$$

$$\text{ b. } [F_i] \rightarrow [F_j] / _ [F_k]$$

The rule in (12a) is rejected by the child learning the language, because it is more complex, compared to (10) – it posits an extra feature F_k , and because it employs three features in the rule rather than two.²² Analysis (12b) is rejected compared to (10) simply because it posits an extra feature F_k .

These SPE-style rules can also be recast as autosegmental rules, for proper comparison with (11). Rule (12a) spreads $[F_j]$ to a preceding segment, provided that the segment on the right is also $[F_k]$, and (12b) inserts $[F_j]$ into the first segment when the following segment is also $[F_k]$.




²²More specifically, that hypothesis would not be entertained, unless there is a specific reason to consider the possibility of such a rule.

Again, the alternative analyses are more complex, either specifying more things in the rule, or positing additional features that are not necessary. From rule (10) or the autosegmental analog (11), we learn the featural analogy $p:m::t:n::k:\eta$. Stops are $[F_i]$ and nasals are $[F_j]$. We also know that vowels are not $[F_i]$, since vowel segments do not change before a nasal.

Another phonological fact of the language is the following.


$$(14) i \rightarrow u / \{p, k, m, \eta\}_$$

This implies a rule²³ of the form:



The rule in (15) allows us to identify $[F_a]$ and $[F_b]$ as features of i and u , respectively, and it also identifies $[F_b]$ as a feature common to labials and velars. We now know that $[F_b, F_j]$ identifies the class $\{m, \eta\}$. Again, the class $\{p, k, m, \eta\}$ could be analyzed as $[F_c]$ and the rule could be stated as inserting a node and association relationship, but that invokes an additional unnecessary feature and a more complex rule.

Finally, in this language, $i \rightarrow a / _\{k, \eta\}$. Thus:



Based on segment behaviour in phonological rules, we have sufficiently learned the features of this language to the point that all segments are distinctively represented.

(17)

p	t	k	m	n	η	i	a	u
	i	i					a	
b		b	b		b			b
			j	j	j			
		c		c		c		

The point of this brief artificial sketch is to illustrate the reasoning of RSFP feature learning, and not to solve all problems in the theory of rules or representations. Put simply, when a class of segments functions together, that is because they have a shared property in the grammar. The task of feature acquisition is finding the simplest system of properties that accounts for those cases of grammatical functioning-together that can be observed in the primary linguistic data.

It should be obvious that RSFP depends on having a well-defined and simple theory of rule formalism. The broad structure of rules (11) and (16) is the same. Why, then, not posit one rule that does both things?

²³Henceforth I dispense with SPE-theoretic statements of rules.

$$(18) \left\{ \begin{array}{c} p \ t \ k \\ i \end{array} \right\}_i \rightarrow \left\{ \begin{array}{c} m \ n \ \eta \\ a \end{array} \right\}_{i/} - \left\{ \begin{array}{c} m \ n \ \eta \\ k \ \eta \end{array} \right\}_i$$

This might be a possible rule using SPE notations such as braces and angled brackets, since the SPE theory of rules has relatively little to say about what constitutes a formally-possible rule.²⁴ Embedded in a theory with minimal rule machinery such as the present autosegmental framework, such an expression has no corresponding formal rule. Therefore, a conjectured generalization along the lines of (18) is never entertained as a hypothesis about the rule system, and the possibility of merging all rules where the focus precedes the determinant into a single rule therefore does not lead to a rule which affects the computation of the interface rules mapping segments to feature matrices.

One methodological point remains, regarding the above model of feature learning. In (17), *t* is specified [F_i] and *n* is specified [F_j]. But some instances of [n] derive from /t/ by (10), which adds the specification [F_j], so [n] from /t/ would be represented as [F_i,F_j], but underlying /n/ would be represented as [F_j]. One solution to this problem (derivational history being maintained in the representation) would be to presume that all [F_j] segments are made to be [F_i] (via language-specific rule). An alternative solution is simply that [F_i,F_j] and [F_j] need not be pronounced differently: both are produced as [n]. There is an analogous problem in vowels, where in the proposed system *i* is [F_a], *u* is [F_b], and *a* is [F_c]. Since [a] can derive from /i/, some instances of [a] are expected to be [F_a,F_c]; and since /i/ also becomes [u], some instances of [u] are expected to be [F_a,F_b]. In other words, all vowels would be [F_a]. Now the problem is that [F_a] no longer identifies *i*; what identifies *i* is the fact of being [F_a] with no other feature specification. As noted in section 2, there has been an assumption that a rule cannot refer to the lack of a feature.

The underlying issue is one that transcends privativity vs. binarity, and features by universality vs. by learning, namely the plausible and widely-adopted but unproven assumption that languages have rules pertaining to licensed combinations of primitives – structure preservation. If a rule would create a structure that is illicit in a language, it has been conjectured that the illicit structure is brought into conformity with the rules defining the object in question (segment or syllable). Whatever the theory of primitives and feature values is, it is at least credible to contend that some mechanism in the grammar of Arabic (most dialects) indicates that *p* and *v* are not segments of the language, even though free combination of independently necessary primitives might allow their existence. Or, perhaps this is an accidental gap: there simply happen to be no lexical items or derivational results containing such a segment, but the segments are not grammatically illicit.

²⁴SPE does define the formal properties of a simple rule, ZXAYW → ZXBYW, but virtually never employs simple rules, instead proffering rule schemata which are expressions for evaluating infinite sets of simple rules in the grammar. RSFP, in contrast, does not allow a grammar to have an infinite set of rules, and does not employ schemata such as braces, angled brackets, and English-language conditions.

The concept of structure preservation has historically been fraught with problems, such as the premise that it is defined in part by the questionable notion “contrast” (so-called allophonic rules are not limited by a structure-preservation requirement). The RSFP view of segments and features denies the significance of the taxonomic phoneme, holding that if p^h , t^h , k^h , r are segments in the phonology of English, then p^h , t^h , k^h , r are segments that have to be assigned features.²⁵ RSFP does not take an a priori stance as to what the segments of a language are: this is an empirical question (where competing answers contribute to language change).

One analytic trend in coping with structure preservation has been to make gaps fall out from the featural analysis of segments – a perfectly valid approach to the problem. If a language has a voicing contrast in obstruents but not in sonorants, and since feature dominance relations are not predetermined by UG, it may be that the feature for voicing is a structural dependent of the feature that distinguishes obstruents from sonorants. If a language has i , e , a , o , u and not ϕ , u , the language may simply not employ the feature [round], and lip protrusion is an articulatory fact about back vowels, having no phonological significance. I take the matter of structure preservation to be a real issue that ultimately needs to be addressed, but also, persistent feature-cooccurrence relations are not a theoretical artifact of RSFP’s making. Whether there is any real problem pertaining to asymmetries and gaps in feature specification remains to be seen.

4.2 Acquisition of Kerewe features

This section demonstrates the logic of feature acquisition to segments in Kerewe, a Bantu language spoken in Tanzania. The goal is not only to show that it is possible to arrive at a feature specification of the segments of the language based solely on phonological behaviour, but also to exemplify the dependence of this analysis on a logic of acquisition and a theory of rule formulation. Since the point of the discussion is to demonstrate how features are learned, I forego extensive empirical discussion and draw on a deeper analysis of the data being developed elsewhere, only providing basic illustrative examples.

As a starting point, the surface segments of Kerewe are as follows.

- (19) p t tʃ k f s b d dʒ g β v z l h m n ŋ ŋ j w
i u e o a

Besides these obvious segmental distinctions, Kerewe is a tone language, but tonal properties are not analyzed here. There is also a robust lexical contrast in the length of vowels (*ekisiβo* ‘tether’, *ekisiβo* ‘fasting’, *emboga* ‘vegetable’, *embooga* ‘infected eyes’, *kuhata* ‘to dislike’, *kuhaata* ‘to peel’, etc.), and a limited but productive nonlexical surface contrast between single and geminate nasals in verbs in utterance-initial position, deriving from /n+C/ sequences, for instance, *n-aahúla* ~ *nn-ahúla* ‘choose me!’, cf. *ku-jáhúla* ‘to choose’. There is good evidence for the classical

²⁵To state the point somewhat differently, “contrast” in RSFP means “is a segment in the language”.

autosegmental treatment of length as involving a many-to-one mapping between segments and higher prosodic units – moras or skeletal positions. It has also been classically assumed that syllabicity and length are represented with suprasegmental non-featural prosodic objects. RSFP does not have a principled commitment to including or excluding higher prosody from the set of learned representational objects, though obviously it would be advantageous if prosody could also be entirely learned. Because so many extraneous issues would arise, I ignore prosody and presuppose a standard moraic phonology account of syllabicity and length. The main consequence for the analysis of segmental features is that the distinction between *j, w* and *i, u* might be just prosodic, or it might also be featural (but there is no evidence in Kerewe for a featural difference). When *j, w* act differently from *i, u*, that could be because their prosody is distinct. In the course of ignoring prosody, it may be useful to know the pattern of segment sequencing which is usually handled by a set of rules of syllable structure. As in most Bantu languages, syllables are of the form (C³V)*, between one and three consonants in the onset followed by a vowel, which defines a syllable.²⁶ The optional left margin of the syllable is always a nasal (homorganic with the following consonant) and the optional right margin of the onset is a glide *j* or *w*.

Some consonants have very limited distribution in Kerewe. Two of them are so limited that it is unclear whether they actually are segments of the language. Those segments are *v* and *dʒ*, which each exist in under a half-dozen recent loan words. The fact of being very low frequency does not per se justify ignoring them. What is not clear is whether they actually exist as the result of normal language acquisition. They may be analogous to *ø, y, x, ll* in English, which are sounds that educated adult speakers can make in pronouncing *milieu, Übermensch, Bach, Xhosa*, but which are not acquired or represented in the same way that *p, t, θ, ɪ* are. It is unknown what the acquisitional facts are surrounding foreign phonemes in Kerewe. It is possible that *v* is more generally nativized as *β* but is pronounced by educated speaker in certain words such as “driver” as *v*, using extragrammatical information. It is known that *dʒ* is an originally Jita phoneme, and most or all Kerewe speakers are bilingual in Jita. There is a widespread strategy of nativizing Jita *dʒ* as *zj* which may be resisted in certain words (especially those coming from Swahili). Because *v, dʒ* do not appear in a context where they clearly undergo or are excluded from a phonological rule, there is little evidence for what their feature assignment is.

The consonant *ŋ* appears robustly before *k, g, ŋ*, but never before any other consonant, and exists in very few words in root-initial position. Although the distribution of *ŋ*, like that of *v*, is highly limited, the evidence suffices to provide a feature assignment to *ŋ*. In like fashion, *b* (distinct from *β*) is distributionally limited. After a nasal, *β* always becomes *b*, and this rule is the main source of instances of *b*. But there are a few instances of *b* which are not so derived.²⁷

²⁶In utterance-initial position, a short V syllable, one without an onset, is also possible.

²⁷It is also unclear how stable underived [b] is. Examples seem to come from loanwords, and sometimes /b/ is pronounced [β], for example in *kubómola ~ kuβómola* ‘to tear down’, influenced by Swahili /bomoa/.

The rules are discussed in a heuristic order, where rules that can be reasonably understood alone are considered before rules that crucially depend on prior knowledge. There is no implication of ordering in the process of learning: the child has to create an integrated analysis that accounts for all of the available facts.

4.2.1 Consonant features

The first rule to be considered is nasal place assimilation. Certain prefixes (1SG subject or object prefix; class 9–10 noun prefix) have /n/ which assimilates in place to a following consonant. This is illustrated in (20) with the 1SG subject prefix /n/. A process of hardening is also attested in some examples.

(20)	m-bazílé	‘I counted’	a-βazílé	‘class 1 counted’
	m-boomílé	‘I was dull-witted’	a-boomílé	‘class 1 was dull-witted’
	m-mazílé	‘I finished’	a-mazílé	‘class 1 finished’
	m-pweezílé	‘I helped’	a-hweezílé	‘class 1 helped’
	m-pekesílé	‘I made fire’	a-pekesílé	‘class 1 made fire’
	m-fuzílé	‘I cleaned’	a-fuzílé	‘class 1 cleaned’
	n-naaβílé	‘I bathed’	a-naaβílé	‘class 1 bathed’
	n-daβílé	‘I passed’	a-laβílé	‘class 1 passed’
	n-tegílé	‘I fished’	a-tegílé	‘class 1 fished’
	n-ziikílé	‘I planted’	a-ziikílé	‘class 1 planted’
	n-sonílé	‘I sewed’	a-sonílé	‘class 1 sewed’
	ɲ-ɲamwíílé	‘I ruined’	a-ɲamwíílé	‘class 1 ruined’
	ɲ-tʃumisílé	‘I stabbed’	a-tʃumisílé	‘class 1 stabbed’
	ɲ-ganílé	‘I told a story’	a-ganílé	‘class 1 told a story’
	ɲ-kwaasílé	‘I touched’	a-kwaasílé	‘class 1 touched’
	ɲ-ɲoloosílé	‘I groaned’	a-ɲoloosílé	‘class 1 groaned’

There are no clear examples of initial /d/ in verb roots – *d* is historically the post-nasal allophone of /l/ – but there is a synchronic initial contrast in noun roots, as shown by forms with a vowel-final noun prefix such as *aka-lezu* ‘little beard’, *aka-dege* ‘little airplane’. Compare *en-dege* ‘airplane’, *en-dezu* ‘beard’.

The underlying form /n/ of the 1SG subject prefix is motivated in (21) before the past prefix *-a-*.

(21)	<i>Simple past</i>		<i>Perfective (hesternal)</i>	
	n-a-gáβá	‘I divided’	ɲ-gaβílé	‘I divided’
	w-a-gáβá	‘you divided’	u-gaβílé	‘you divided’
	tw-aa-gáβá	‘we divided’	tu-gaβílé	‘we divided’

In general, any nasal can appear before any vowel, but a nasal before a consonant is always homorganic with the consonant.²⁸ The pattern that the child will be aware of is as follows:

²⁸There is a non-phonological process of phonetic implementation whereby the vowel *u* may have reduced duration when preceded by *m* and followed by a consonant, so that the vowel in *omuβígi* ‘trap fisherman’ is noticeably shorter compared to *oβuβígi* ‘act of trap-fishing’. In my data, the vowel is present in words like *omuβígi* most of the time, but it is

- (22) $n \rightarrow m / \text{ ____ } p, b, f, m \text{ (v) (w) } <\beta \text{ h}>$
 $n \text{ ____ } tʃ, ɲ \text{ (d}_3\text{)}$
 $\eta \text{ ____ } k, g, \eta$
 $n \text{ ____ } t, d, s, z, n \text{ } <l> \text{ } <j>$

The significance of parenthesized (v, d₃) is that there is an attestation gap for these segments. The consonant *v* never appears after a nasal. Root-initial *d*₃ is also unattested, so we cannot unambiguously observe alternating prefixes with final /n/ appearing before *d*₃. However, *d*₃ does appear root-internally after a nasal in a very few words – *kuβúúɲdʒa* ‘to peddle’ and the noun *embúúɲdʒá* ‘jigger’, as well as *emooɲdʒo* ‘toy top’. It is far from certain that the nasal in these words is the result of applying an assimilation rule: the underlying roots may well simply be /βúúɲdʒ, mooɲdʒo/. Still, this constitutes weak evidence pointing to a possible place of articulation specification for *d*₃. To avoid building a theory of features on a weakly supported conclusion, I set aside the question of place feature for *d*₃, until we come to the point of obligatorily distinguishing *d*₃ from other sounds.

The glide *w* also never appears in the relevant context (root-initial position), underlyingly or in derived forms, where it could be preceded by an alternating nasal. Angled brackets in (22) indicate segments which have been independently removed from rule inputs by another rule. The glide *j* does underlyingly appear in this context, but because of a prior deletion rule, it is not present when assimilation takes place. Likewise /β, l, h/ are present in underlying forms, but have become *b, d, p*, respectively.

There is some freedom in the statement of the rule in question. All and only consonants are triggering segments – no consonants must be explicitly excluded from the rule, but all vowels are excluded. However, all consonants are within the onset of the syllable, so provided that the domain of the rule is stated as being the onset, all segments in that domain trigger the rule. Because the language has no sequences of obstruents at any stage of the derivation, no consonants have to be excluded from the input as non-alternating. As far as trigger segments are concerned, because of the restricted combinatorics of complex onsets in Kerewe, the trigger consonant is always the unrestricted obligatory consonant, C₂ in the (C₁)C₂(C₃) onset template.

- | | | | | |
|------|-------------------|----------------|-------------------|-----------------------|
| (23) | (N ₁) | C ₂ | (C ₃) | |
| | m | b | w | á ‘dog!’ |
| | | n | w | á ‘drink!’ |
| | ɲ | k | | é ‘few’ ²⁹ |

Given the option of referring the identification of target and trigger segments to positions within the syllable, picking out target and trigger may not directly provide direct evidence for features ([Nasal] and something similar to [Consonantal]). But a rule whereby the N₁ position is only filled with a nasal provides evidence for a feature that we can term [Nasal] (the label used to refer to features is completely arbitrary in RSFP and carries no implications about phonetic substance), and the rule whereby the


plausible to expect that under Swahili influence (where there is categorial vowel deletion) the process is becoming phonologized among younger speakers.

²⁹Numeric subscripts refer to noun classes, thus ‘few’, with class 9 agreement.

C₃ position is only filled with a glide provides evidence for something that identifies glides ([V-Place] – another arbitrary label – see section 4.2.2). In fact, there is good evidence in a number of contexts that surface post-consonantal glides derive from high vowels via the rule of glide formation, thus [nwá] is /nu-á/.

The nature of the change is more informative regarding features. The structural change of the rule is conditional, in that the exact output segment depends on which subset of segments appears to the right. It is obvious that the place specification of the input nasal is replaced by the place specification of the following consonant. In an autosegmental representation, a node labeled Place spreads to a preceding segment in the onset. The underlying formal premise is that every phonological rule operates on one phonological unit: one element of the input segment is changed.³⁰ Since there are four distinct outputs depending on the subset of segments which follows, there must be four distinct feature configurations dominated by the node [Place]. We observe that the set {p, f, b, m} functions as one class (termed [Labial]) which conditions *m*, {tʃ, ɲ} and possibly *dʒ* function as another group ([Palatal]) which conditions *ɲ*, and {k, g, ŋ} as a third ([Velar]) conditioning *ŋ*. Possibly, {t, s, z, d, n} are unified with a feature ([Coronal]); or perhaps they are not specified with [Place], and instead underlying /n/ is unchanged in that context, just as it is unchanged before a vowel. Since there is other evidence for a positive specification of [Coronal], we will at least initially adopt a four-feature account of consonantal place. The place assimilation rule is a standard autosegmental spreading rule.³¹ ‘X’ refers to the segmental root node.

(24)



On the basis of these alternations, we also gain evidence that some feature unifies the set {m, n, ɲ, ŋ}, namely [Nasal], which is a property of the input *n*. That feature is inherited from *n* in the case of derived {m, ɲ, ŋ}: the reason why *n* becomes *m* before labials, and not some other consonant, is that the only difference between *n* and *m* is that the former is [Labial], and {n, m} have in common the property [Nasal]. At this point, we have learned the existence of the feature [Place] (the property which spreads), four nodes dominated by [Place] namely [Labial], [Palatal], [Velar], and [Coronal], as well as [Nasal] (inherited from /n/, and unifying the set {m, n, ɲ, ŋ}).

³⁰There is a competing scenario, proposed in Reiss (2003), whereby a rule may contain expressions quantifying over subsets of the features, for instance “all of the features of the set [coronal, anterior, back]”. I will not pursue that approach here, but assuming such a rule statement, one would simply gain information about the individual features and their set-theoretic unification *qua* the “place” set.

³¹It is beyond the scope of this work to satisfactorily address the highly relevant question of how to interpret autosegmental notation in RSFP-FP. The rule in (24) should not be taken to assert that X immediately dominated [Place]. At most, the rule asserts that a segment which is C receives [Place] from a following segment that is C.

The computational theory allows there to be a series of separate rules which individually spread [Labial], [Palatal] and [Velar], and since the alternating prefix is /n/, only three rules would be necessitated. Each of those rules would have the form of (24), applying to a specific [Place] feature. Although this would (for the moment) eliminate the argument for [Place], the simplification in features is counterbalanced by a significant complication in rule system. Therefore, the three-rule solution is rejected by the language-learning child.

The second rule to consider is post-nasal Fortition, whose effect, descriptively, is:

$$(25) \{\beta, l, h\}_i \rightarrow \{b, d, p\}_i / \{m, n\} \langle \eta, \eta \rangle \text{ ---}$$

Examples of this process are seen in (26), and we can add examples of the present tense where the 1SG subject stands right before the verb root.


(26)	n-dímá	‘I cultivate’	a-límá	‘class 1 cultivates’
	m-béézá	‘I carve’	a-βéézá	‘class 1 carves’
	m-panaantúká	‘I descend’	a-hanaantúká	‘class 1 descends’

Not all consonants can appear post-nasally in underlying forms. There is no clear example of rare underlying /b/ appearing here, but there are no problematic cases either.³² The consonants *f, s, z, p, t, tf, k, b, d, dʒ, g* do appear post-nasally as previous examples have shown, and they are not changed, so these segments must either be explicitly excluded, or the effect as realized on these consonants must be vacuous. No vowels trigger fortition, and no nasals have to be excluded as non-triggers. Again, the only morphemes which demonstrably trigger this process have the underlying form /n/. There are no input cases of *n, η, m* which precede /β l h/, so those segments could be excluded or included as necessary in order to achieve a more economical grammar (no savings results from excluding any nasals). As in the case of place assimilation, these facts can be expressed by limiting the rule to affecting C in the onset.

From the fact-pattern in (25), we conclude that *β, l, and h* have a feature in common that distinguishes them from the set {*f s z t tf k dʒ g*}. The pairs *β* vs. *b*, *l* vs. *d*, *h* vs. *p* are themselves distinguished by a feature: the members of the pairs {*β, b*}, {*l, d*}, {*h, p*} are the same except for some feature, which is the feature that is changed by this rule. The simplest solution is that the target-identifying feature and the changing feature are the same. We can mnemonically call this feature [Approximant], again without any implication that this class is defined by a phonetic property. The rule turns Approximants into non-Approximants. Thus *β/b* are the same (the features of *b* are inherited from *β*), save for the feature [Approximant], likewise *l/d* and *h/p*. The change performed by the rule can be formalized as deletion of [Approximant] when preceded by an onset segment, though the specific context might be stated in terms of the preceding segment being a nasal.

³²It is empirically unanswerable whether the underlying form of roots such as *gaamba* ‘say’ is /gaamb/ or /gaamb/. In the latter case, no rule is required to derive the surface form, and in the former case, the independently motivated rule Hardening will derive [gaamb]. No motivated mechanism of the grammar would turn /Nb/ into anything different from [mb].

(27)



By specifying only β , l , and h as [Approximant], we guarantee that only β , l and h are affected. Based on these two rules, we have a partial assignment of features to segments.

(28)	p	t	tʃ	k	f	s	β	z	b	d	g	l	h	m	n	ɲ	ŋ	
	Pl	?	Pl	Pl	Pl	?	Pl	?	Pl	?	Pl	?	Pl	Pl	?	Pl	Pl	
	L	Cor	Pa	Ve	L	Cor	L	Cor	L	Cor	Ve	Cor	L	L	Cor	Pa	Ve	
							Ap					Ap	Ap					
															N	N	N	N

A third rule, Spirantization, changes t , $tʃ$, d , and l to s , s , z , and z respectively before three of six morphemes which underlyingly begin with i .

(29) $\{t, tʃ, d, l\}_i \rightarrow \{s, s, z, z\}_i / __ i$

Hesternal perfective examples with the triggering suffix *-ile* illustrate this change.

(30)	a-has-ílé	‘class 1 disliked’	ku-hat-a	‘to dislike’
	a-fuz-ílé	‘class 1 washed’	ku-ful-a	‘to wash’
	a-laanz-ílé	‘class 1 interlaced’	ku-laand-a	‘to interlace’
	a-pekes-ílé ³³	‘class 1 started a fire’	ku-pékétʃ-a	‘to start a fire’

The following are examples of other consonants which do not change.

(31)	<i>Isg hesternal perfective</i>		<i>infinitive</i>
	n-dom-ílé	‘I made a speech’	ku-lom-a
	n-doβ-ílé	‘I fished with hooks’	ku-lóβ-a
	ŋ-kop-ílé	‘I borrowed’	ku-kóp-a
	ŋ-gon-ílé	‘I snored’	ku-gon-a
	m-bis-ílé	‘I concealed a fact’	ku-βís-a
	m-boj-ílé	‘I fought’	ku-βój-a
	ŋ-kop-ílé	‘I cooked improperly’	ku-kój-a
	n-dog-ílé	‘I bewitched’	ku-log-a
	m-bik-ílé	‘I announced a death’	ku-βík-a
	m-boh-ílé	‘I tied’	ku-βóh-a

This rule applies before the suffixes for nominalization *-i*-, causative *-i-*-, and perfective *-ile*-, but not stative *-ik-*-, applied *-il-*-, or causative *-isj-*.³⁴ The reason for the behavioural divergence in instances of *i* is that the trigger morphemes in Proto-Bantu had the vowel *i* but the non-trigger vowels had *ɪ*. Finding the proper

³³Most words which appear to have root-final *tʃ* actually have underlying /kʃ/ which surfaces as [tʃ], and they have a substantially different form in the perfective, cf [ku-βwáátʃ-a] ‘to greet’, [m-bwaak-ííʒjé] ‘I greeted’. The root /pekétʃ/ has true underlying /tʃ/.

³⁴The causative suffix appears as [j] since it always undergoes glide formation, but there is sufficient evidence that it is underlyingly /i/.

mechanism for identifying these specific morphemes as triggers is not essential to our present goal (a solution is discussed at the end of 4.3).

Rule (29) provides evidence that the members of the set {*t*, *tʃ*, *d*, *l*} have something in common. Any consonant (aside from marginal *v*, *b*, *dʒ*) can appear in the triggering context, and notably *p*, *β*, *k*, *g*, *m*, *n*, and *ɲ* do not change. There is no need to explicitly exclude *s*, *z* from the input class, since vacuous application of the rule to these segments also yields the correct output. The segments targeted by the rule are definable by the feature [Coronal], and the simplest analysis is that [Coronal] is unaffected by the rule (hence *s*, *z* are also [Coronal], the output segment having inherited that specification from its input).³⁵ We also observe that *t* and *d* map to distinct outputs – some distinguishing feature (a form of voicing) is preserved from the input. One option is that *t* has a voicing specification and *d*, *l* do not: or the converse.³⁶ The nature of the shared change also establishes that the relationship of *t* to *s* is the same as the relationship of *d*, *l* to *z*, that is, the rule adds (or subtracts) something in *t*, *tʃ*, *d*, and *l* to give *s*, *z*. We will eventually see evidence for a feature [Voiceless], necessitated by the fact that *k* is targeted by palatalization while *g* is not.

As for the change performed by the rule, it could be that the input segments gain the specification [Fricative], or else that an existing feature [Stop] is deleted. The overall strategy is to equate the sets {*t*, *tʃ*, *d*, *l*} versus {*s*, *z*, *n*} via feature specification, so that none of {*t*, *tʃ*, *d*, *l*} would change to *n*. Assume first that the members of the set {*t*, *tʃ*, *d*, *l*} have a feature that is deleted, [Stop]:

$$(32) \quad \begin{array}{c} X \\ \swarrow \quad \searrow \\ \text{Stop} \rightarrow \emptyset \\ \text{Coronal} \end{array} \quad [i]^{37}$$


If *n* does not have the specification [Stop], it will not undergo rule (32). If *n* is specified [Stop], and assuming another feature which positively identifies nasals (e.g., [Nasal]), either (32) requires an additional condition whereby the target cannot have [Nasal] (a complication which is theoretically problematic – requiring a rule to refer to the lack of a property), or an additional rule is required to reinsert [Stop] on [Nasal] segments. Another option is that the grammar generates stop and non-stop nasals, which are phonologically and phonetically indistinguishable.

Alternatively, a feature [Fricative] could be assigned by rule to otherwise unspecified *t*, *d*, *l* in this context.

³⁵The previous identification of [Palatal] as independent of [Coronal] now gives way to an analysis of [Palatal] as being a combination of [Coronal] and [Velar]. This decomposition of [Palatal] is partially motivated by *kj*-fusion to be discussed, and the present fact that *tʃ* acts as part of a class that includes [t,d,l].

³⁶It is also possible for all segments to be characterized by two such features, but such a hypothesis would not be considered without compelling evidence for its necessity.


³⁷Here, “[i]” refers to whatever properties identify the triggering vowels.



Again, the matter of how *n* is treated is a prominent problem: what unifies *t*, *d*, and *l* excluding *n*? Must the rule be formalized to exclude *n*, or can the rule apply vacuously? Assuming (33), if *n* is underlyingly specified as [Fricative], applying (33) to *n* does not change *n*. In comparing (33) which inserts both a node and an association with (32) which only deletes a node, [Stop]-deletion is the simpler rule.

We are not done with structure-preservation type problems with Spirantization, since both *d* and *l* become *z*. From Fortition (27) we know that *l* is Approximant and *d* is not. *z* is also not affected by Fortition, therefore it lacks [Approximant], which identifies targets. In the case of *d* → *z*, we could equally treat the change as [Fricative]-insertion or [Stop]-deletion. In the case of *l* → *z*, we also need a mechanism for removing [Approximant] (since *z* is not specified [Approximant]). When the spirantization change takes place, [Approximant] is also removed. This indicates that there is a more general feature-deletion process – removal of [Stop], and [Approximant] if present.

This dependency relationship has a straightforward solution, using representational machinery well-motivated in phonology which has been universally assumed in autosegmental theories of representation. If [Approximant] is a dependent of (dominated by) [Stop], [Approximant] will always delete when [Stop] deletes. The proposed representation of relevant segments is given in (34).



Under this analysis, it follows that *h* and *β*, the other two [Approximant] segments, are also [Stop] – not an expected outcome based on the substantive connotations suggested by feature names, but RSFP features have no substantive associations, and names are arbitrary. There is nothing peculiar about saying that *h* and *β* are [Stop], because “Stop” implies nothing about phonetics. This analysis flies in the face of the view that structural simplicity should reflect intuitions about what are “most basic” segments. See Odden (2013: 252): “the observation that a certain fact is ‘rare’ or ‘marked’ is irrelevant” – the fact that fricatives have less structure than stops is not a valid argument against this analysis.

A detail regarding spirantization needs attention: both *t* and *tʃ* become *s*. The difference between *t* and *tʃ* is place of articulation – some feature distinguished these stops (initially identified as [Palatal] but now analyzed as the combination of [Coronal] and [Velar]). Since *t* and *tʃ* have different place features, the fricative versions ought to inherit that difference. Yet both stops become the phonetically

identical fricative, *s*. There are three obvious responses to this problem; the problem is not overwhelming, but it is also not obvious which solution is correct. First, it could be that the unmodified outputs – *s* which is the fricative coming from *t*, and what we could symbolize as *f* for the fricative coming from *tf* – are simply not pronounced differently, even though they are featurally different. We cannot evaluate a phonetic-interpretation account without a substantially-supported theory of linguistic phonetic computation that addresses this option, which we presently lack. Second, we could appeal to the phonological notion of structure-preservation, meaning that there are well-formedness rules governing allowed feature configurations on segments (probably the same mechanism generates syllable well-formedness conditions). In that approach, the additional feature for palatals cannot appear on an oral non-fricative – the language does not have *f* or *ʒ*, and any rule which creates such a configuration automatically repairs the ill-formed structure. Saying how the repair is automatically brought about is complex (why not restore [Stop] or make the segment [Nasal]?). This brings us to the third option, that there simply is a rule in the language which deletes [Velar] from an oral non-stop (there is also no *x* or *ɣ* in the language). The problem with this approach is that “non-stop” is an appeal to lack of specification. We will not attempt to resolve this problem here.

One last alternation provides phonological evidence for the analysis of Kerewe consonants, namely a palatalization rule where the sequence *kj* becomes *tf*. While the marginal segment *ŋ* does not appear in a context where we can test whether it undergoes an analogous change (e.g., root-finally in a verb), *g* does, and *gj* is not changed. This process can be motivated with examples of the short causative *-j-*, which stands between the root-final C and the final suffix, *-a* in the following examples.³⁸


(35) <i>Infinitive</i>	<i>Causative infinitive</i>	
kuβuuka	kuβuutfa	‘comb’
kuhika	huhitfa	‘arrive’
kuseleka	kuseletfa	‘hide’
kuhiingga	kuhiingja	‘exchange’
kuloga	kulogja	‘bewitch’
kulima	kulimja	‘cultivate’
kuβííha	kuβííhja	‘be bad’
kulóβa	kulóβja	‘fish with hooks’
kukópa	kukópja	‘borrow’
kuβóna	kuβónja	‘find’
kuβísa	kuβísja	‘conceal’

Thus there is a feature, [Voiceless], which *k* has and *g* lacks, and the Palatalization rule refers to this feature in selecting the target. The rule also identifies *k* owing to it being specified [Velar], which excludes other places of articulation. By feature inheritance it follows that *tf* is also [Voiceless]. The simplest account of the rule where *kj* → *tf* is that the features of the input segments are merged into one segment: meaning that *tf* is [Velar] plus whatever place feature characterizes *j*. That

³⁸This suffix is also one of the triggers of spirantization, so examples of final *l,d,t* are changed for an independent reason.

feature is [Coronal], a fact which we knew from the outcome of Spirantization applied to *tf*.³⁹ Thus the hypothesized feature “Palatal” proves to be superfluous, and instead is to be analyzed as the conjunction of two directly-motivated features, [Coronal] and [Velar]. The evidence for this treatment of ostensive [Palatal] derives from the specific case *kj* → *tf*, therefore we lack direct evidence of this type for the treatment of *ɲ* (or *dʒ*, whose status is unclear) therefore we might maintain two analyses of palatals: those with the feature [Palatal], and those with the combination [Coronal]+[Velar]. That alternative is eliminated from consideration, since the resulting grammar is more complex (positing an extra feature). The desideratum of simplicity thus provides information that is not directly available in the speech signal – all phonological palatals should be treated the same, unless there is direct evidence to the contrary.

Finally, because *ki* does not change but *kj* does, as shown by *eki-tóóke* ‘banana’, cf. *etf-áála* ‘finger’ from /eki-áala/, it is essential that the input sequence be limited in domain to the Onset. Thus the Palatalization rule is as follows.



It is now appropriate to take stock of our feature analysis. We have identified [Place] as the item operated on by [Place]-assimilation, and the features [Labial], [Velar] and [Coronal] as the dependents of [Place] which are carried along in this assimilation. Previously-assumed [Palatal] has been supplanted with the combination of [Coronal] and [Velar] within a segment, on the evidence of the Palatalization rule. [Coronal] is also motivated in the target-selection aspect of the Spirantization rule, and [Velar] is likewise motivated via the Palatalization rule. The consonants *v*, *w*, and *dʒ* have been removed from the inventory below, because so far nothing in the phonology tells us anything about the make-up of these consonants. Question mark indicates uncertainties, situations where we cannot so far tell whether the segment has the feature in question.

(37)

p	t	tf	k	f	s	β	z	b	d	g	l	h	j	m	n	ɲ	ŋ
Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl
L	Co	Pa	Ve	L	Co	L	Co	L	Co	Ve	Co	L	Co	L	Co	Pa	Ve
?	St	?	?	?		St		?	St	?	St	St	?	?		?	?
						Ap					Ap	Ap	?				
?	VI	VI	VI	?	VI	?		?		?	?	?	?	?	?	?	?
														N	N	N	N

Not all consonants are fully differentiated at this point, but few consonants remain undifferentiated. In particular, the segments {p, f, b} are not yet distinguished from each other (they are all “labials”), and nor are {h, β} (both are “labial approximants”),

³⁹A side effect of this merger is that [tʃ] should not exist in the language, which is indeed the case.

and {v, d₃} are not yet distinguished from any other consonant. As the question marks indicate, it is possible that *p* is [Stop], or [Voiceless] – we have not seen positive evidence in the form of phonological rule behaviour that would justify an assignment. We can call on the fact that *h* becomes *p* after a nasal by deletion of [Approximant], to fill in missing features. Noting that *h* bears [Stop], the simplest analysis is that the output *p* bears it as well.⁴⁰ The assignment of [Stop] to *b* also follows from the analysis that fortition deletes [Approximant] *qua* dependent of [Stop]. In other words, *p* and *b* are both [Stop], and *f* is not. We cannot derive direct evidence from rule behaviour as to whether *k* or *g* are [Stop] or not. But because *kj* becomes *tf* by merger of features and *tf* is [Stop], we know that at least one of *k*, *j* is [Stop].

We have not provided any direct evidence for [Labial] – no rule specifically refers to labials. Now observe the place representation of the attested consonantal places.

(38)	k	t	tʃ	p
	Pl	Pl	Pl	Pl
	Ve	Co	Ve,Co	La

The four-way distinction in types of [Place] can, in fact, be derived without the invocation of a separate feature [Labial]. Instead, labials may be segments with a [Place] node, but no other feature under [Place], i.e.

(39)	k	t	tʃ	p
	Pl	Pl	Pl	Pl
	Ve	Co	Ve,Co	

There being no specific reason to posit [Labial] in addition to [Place], [Labial] can be eliminated from the feature inventory, and the desideratum of simplicity dictates that it is eliminated.

4.2.2 Vowel features

Now we turn to the analysis of vowels. As a class, vowels have a non-featural property in common, that they are dominated by one or two moras, and the difference between *j*, *w* versus *i*, *u* is at least a difference in moraic status. Since we are only investigating the acquisition of putatively phonetically-based features and not all representational entities, we will freely exploit the existence of contrastive prosodic properties as a means of avoiding postulation of features, making it harder to justify the invocation of a feature when a prosodic distinction is available.

For example, vowel segments are always syllable-final (there are no codas or diphthongs) and never syllable initial (except in short utterance-initial syllables). The grammar does not express these facts in terms of the concept “vowel”, it does so via rules regulating moras. One relevant rule is Glide Formation, which turns vowels other than *a* into corresponding glides.⁴¹ These examples also illustrate an

⁴⁰This follows from the fact that we have direct evidence that [h] is [Approximant], a feature dominated by [Stop].


⁴¹The mid vowels [e,o] do not generally appear in prefixes, except in absolute word-initial position. It is possible, but not guaranteed, that the class 9 and 2sg prefixes are underlyingly

optional rule of glide-deletion which deletes root-initial glides which are not word-initial (only *j* appears in the relevant context).

(40)	tu-janíká	twaaníká	‘we spread’
	o-játǰá	wáátǰá	‘you SG. light’
	mu-jégá	mwéégá	‘you PL. rejoice’
	gu-jolóǰá	gwoolóǰá	‘it ₃ is soft’
	βi-jologozíβwá	βjoologozíβwá	‘they ₈ are cleaned’
	li-jóká	ljóóká	‘it ₅ burns’
	e-jóká	jóóká	‘it ₉ burns’

The Glide Formation rule is potentially storable without reference to any features at all, given that non-application of the rule to *a* is due to an independently motivated rule which eliminates *a* + *V* sequences. The question of how to properly express rules of desyllabification with compensatory lengthening is a matter of longstanding controversy which we will not enter into: (41) is stated to explicitly perform all of the relevant prosody-to-segment relations, and it is a separate question how such rules are properly formalized.

(41)



The featurally-relevant thing that we learn from this alternation is that *j* and *i* are prosodic variants of each other, likewise *w* and *u*: that is, they are featurally the same (unless we find that there is an additional feature which correlates with the prosodic difference). Given that *e* and *o* also become *j* and *w*, we would extend the featural analogies to the sets *e, i, j* and *o, u, w*. Independently, we know the behavioural analogies unifying these sets of vowels from a rule of Vowel Fusion, to which we now turn.

Word-internally, *a* fuses with a following vowel, so that *ai* and *ae* become *ee*, *au* and *ao* become *oo*, and *aa* becomes *aa*.⁴² We can draw on the optionality of root-initial *j*-deletion to generate suitable vowel sequences which undergo Vowel Fusion, looking at examples of the present tense.

(42)	βa-jéémbá	βéémbá	‘they sing’
	βa-jíβá	βééβá	‘they steal’
	βa-joléká	βooléká	‘they point’
	βa-jahúlá	βaahúlá	‘they choose’

The sequence *a* + *u* is hard to come by, but the example [am-óóla] ‘shavings’ from /ama-úla/, cf. [elj-úúla] ‘shaving’, shows that the rule applies to *au* as well.

/i,u/ and are subject to word-initial lowering. I assume a more concrete underlying representation, /e,o/ which are the actually-observed vowels, but I do not crucially rely on it.

⁴²I follow the standard convention of notating long vowels as double vowels, meaning that they are a single segment with two moras.

There are at least two plausible approaches to the relationship between Glide Formation and Vowel Fusion, in terms of specifying which vowels undergo which process. One is that Glide Formation is explicitly restricted so that it does not apply to *a*, and Vowel Fusion subsequently applies to any vowel sequences that remain (*a* + *V*). The other is that Vowel Fusion explicitly applies only to *a* + *V*, and Glide Formation subsequently applies to any remaining vowel sequences. In the former analysis, the formalization of Glide Formation requires that target vowels be specified with a feature [A] which is lacking in *a*, and in the latter analysis Vowel Fusion requires the first vowel in the sequence to have a feature [B] which is found in *a* but not *i* or *u*. Either of (43a) or (43b) would seem to be possible partial specifications at this point (lacking height features).

- (43) a. a e o i u
 A A A A
 X Y X Y
- b. a e o i u
 B
 X Y X Y

Equally relevant is the fact that *a*, *e*, *o* have a feature in common, which sets these vowels apart from *i* and *u*. The merger of *ai* and *au* into *ee*, *oo* indicates that *a* intrinsically bears the feature that distinguishes mid vowels from high vowels. We can identify that feature as [Mid], whose existence is necessary regardless of how the targets of Vowel Fusion and Glide Formation are identified. Thus we select between the following feature assignments, where [X] and [Y] are whatever distinguishes front from back/round vowels.

- (44) a. a e o i u
 A A A A
 X Y X Y
 Mid Mid Mid
- b. a e o i u
 B
 X Y X Y
 Mid Mid Mid

The choice between these analyses would be arbitrary, unless some independent evidence exists for the [A] grouping or the [B] grouping.

We turn now to evidence which supports the [A] grouping, coming from vowel harmony. There is theoretical evidence for two such rules, one of which turns *i* into *e* after *e* and *o* (skipping over any consonants), with a second rule (likewise skipping over any consonants) turning *u* into *o* after *o* only. The theoretical premise behind the conclusion that there are two harmony rules is that rule formalism does not contain expressions encoding dependencies like “applies to X only if the trigger is also Y”, as could be expressed using SPE angled brackets notation. In the case of a suffixal front vowel, a [Mid] vowel causes *i* to become *e*. The last three examples show that no vowel can stand between the target and trigger vowels.

(45) <i>plain Verb</i>	<i>V + applied</i>	
kuβúga	kuβúgíla	‘to paddle’
kugaβa	kugaβíla	‘to divide’
kuβója	kuβójéla	‘to fight’
kuhoonga	kuhoongela	‘to bribe’
kumweemweesa	kumweemweesela	‘to smile’
kugeenda	kugeendela	‘to go for’
kugelula	kugelulila	‘to deduct’
kuβónékana	kuβónékanila	‘to appear’
kulémála	kulémánila	‘to be lame’


The triggering segments can be identified as [Mid], and the structural change is that the target becomes [Mid]. The vowel *a* does not trigger this rule: the analysis of *a* and the final version of this harmony rule are discussed below. Because, as we will see, the vowel *u* does not become *o* after *e* (it only lowers after *o* by a more specific rule), we must also restrict the target to a vowel that is front. We can tentatively state the rule as follows.⁴³



The reversive /ul/, often doubled, serves to motivate a second rule lowering *u* only after *o*.

(47) kusiβ-ul-a	‘to untie an animal’	kusiβ-ik-a	‘to be tethered’
kuhan-ul-a	‘to take in laundry’	kuhan-ik-a	‘to be hung up’
kusemb-ulul-a	‘to unwrap’	kusemb-a	‘to wrap up’
kuβóh-ólol-a	‘to untie’	kuβóh-a	‘to tie’

To identify the more restricted trigger *o*, this rule requires specification of a feature present in *o* and lacking in *e – o* is [Rd].



The obvious question that arises from these rules is, why does *a* not trigger application of (46), bearing in mind that *i* and *u* also do not cause lowering of *i*? As contemplated in (44), it is possible that *e*, *o*, *i* and *u* have a shared feature, or else *a* has a feature that is unique to it. The fact that *a* does not condition either vowel harmony rule motivates the decision that *e*, *o*, *i*, and *u* have something in common, a feature which is lacking from *a*, even though it has the feature which spreads, and harmony refers to that property. Now assigning mnemonic labels to the features that we have identified, the feature unifying vowels other than *a* is [V-lace], which


⁴³The rule spreads the feature [Mid] from vowel to vowel, where “vowel” is captured non-featuresally via reference to moraicity. This formalization should not be interpreted as saying that [Mid] is immediately dominated by μ.

correlates with presence of either Front or Round.⁴⁴ Although only *o* triggers (48), *o* is just the conjunction of a place feature ([Rd]) and a height feature ([Mid]), and we know from Fusion and Glide Formation that *o* and *u* are the same except for the distinguishing feature [Mid], therefore we know that *u* is [Rd].


(49)

a	e	o	i	u
	Vp	Vp	Vp	Vp
	Fr	Rd	Fr	Rd
Mid	Mid	Mid		

Rule (46) is therefore revised as follows, to include the restriction that only [Mid] vowels bearing [V-Place] trigger lowering of *i*.



With vowel feature specifications now resolved, we can now adopt the rule contemplated in (41), with the provision that the target bears V-Place. Vowel Fusion simply merges the content of the sole remaining prevocalic vowel *a*, which has just a [Mid] specification, with the feature specifications of the following vowel. The formulation below explicitly states all of the rearrangements involved, and it is left to separate theorizing within the theory of rule formulation to determine how this rule should be expressed. The underlying theoretical concept is that the output is the union of the features in the input.



At this point, the vowels have been fully distinguished from each other and from the glides.

(52)

i	u	e	o	a	j	w
μ	μ	μ	μ	μ		
Vp	Vp	Vp	Vp		Vp	Vp
Fr	Rd	Fr	Rd		Fr	Rd
		M	M	M		

⁴⁴As pointed out previously, feature names are arbitrary, and we could call this feature “Back”. As will be shown below, this is a preliminary hypothesis, to be supplanted with a superior analysis which uses exactly the same features for consonants and vowels, thus in fact [Coronal] and [Velar].

These feature assignments must be integrated with the feature assignments for consonants, and simplifications are possible. The first, for which there is direct evidence, is that the conjectured feature invoked solely for vowels is in fact identical to [Coronal] invoked for consonants. It was previously found that [Coronal], which unifies the set {t, d, s, z, l, n, tʃ} can come from the merger of *k* and *j*, justifying the conclusion that *j* is a non-moraic [Coronal]. We also know from Glide Formation that front vowels become the glide *j*, therefore front vowels have the feature [Coronal]. The alternation *ki-jóká* ~ *tʃóóká* ‘it (class 7) burns’ shows the combined effect of Glide Formation and Palatalization which supports the equation of [Coronal] with [Fr(ont)] as applied to vowels. The desideratum of simplicity motivates exploiting an existing feature [Velar], necessitated for consonant phonology, as the place feature underlying *o, u*, and *w*, rather than positing a unique feature for vowels.

The feature labeled [V-Place] is posited as the node dominating [Front = Coronal] and [Round = Velar] based on the need to exclude *a* as a trigger of vowel harmony: but is there a reason to treat [V-Place] as different from [Place], a feature motivated for consonants? There is not: simplicity dictates that since [V-Place] is superfluous in light of [Place], [V-Place] should not be added to the analysis. While vowels and consonants differ phonologically, the distinction moraic / non-moraic suffices to express that difference.

We should consider whether it is necessary to posit [Mid] as a distinct feature, or if there is already some feature in the consonantal inventory, which could be exploited to take its place. The feature [Approximant] is one possibility, as are [Nasal] and [Voiceless]. These are features motivated for consonants, which have found no other place in vowel phonology, and are therefore available for exploitation. Let us consider the possibility that “Mid” is simply the feature [Approximant] on moraic segments. That results in the following feature specifications for vowels.

(53)	i	u	e	o	a	j	w	
	μ	μ	μ	μ	μ			
	Pl	Pl	Pl	Pl		Pl	Pl	
	Co	Ve	Co	Ve		Co	Ve	
		Ap	Ap	Ap				

This analysis has two consequences which have to be empirically evaluated, but which cannot be resolved in this article.

The first consequence relates to the dominance account of [Approximant] and the fact that when stops lose [Stop], they also lose the feature [Approximant] – because [Stop] dominates [Approximant]. We have not found evidence for assigning any equivalent of [Stop] to vowels. Under the theory that “Mid” is really [Approximant], we might conclude that all vowels have the feature [Stop], and mid vowels additionally have [Approximant] thereunder. An alternative is that the dominance relationship between [Stop] and [Approximant] is rule-governed (as indeed it must be, by the logic of RSFP), and the question is then: What is the rule? It might be “[Approximant] must be dominated by [Stop]”, but it might also be “If a segment has [Stop] and [Approximant], [Stop] must dominate [Approximant]”. The formal theory of structural rules needs deeper investigation,

(55)	p	t	tʃ	k	f	s	β	z	b	d	g	l	h	m	n	ɲ	ŋ
	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl	Pl
		Co	Co,	Ve	Co		Co		Co	Ve	Co				Co	Co,	Ve
			Ve													Ve	
	St	St	?	?			St		St	St	?	St	St	?		?	?
							Ap					Ap	Ap				
	α	Vl	Vl	β	?	Vl	-α		-α		-β		α	?	?	?	?
														N	N	N	N
	i	u	e	o	a	j	w										
	μ	μ	μ	μ	μ												
	Pl	Pl	Pl	Pl		Pl	Pl										
	Co	Ve	Co	Ve		Co	Ve										
			?	?	?												
			Ap	Ap	Ap												

The grammatical facts have not yet given us a unique assignment of features to segments, even though they do determine the vast majority of features.

We must entertain the possibility that undecidable features are assigned at random from the grammatical perspective. This is especially obvious in the case of *v* and *dʒ* – if we assume that these segments are within the language’s segmental inventory. The only phonological information that we can glean from these segments is that they are not moraic since they appear after vowels (*edʒaházi* ‘ship’, *omundeléwva* ‘driver’). Barring an accidental-gap analysis, they are not [Nasal] since they do not appear preconsonantly. We have noted that in the rare instances where a non-alternating nasal appears before *dʒ*, it is palatal. Positing that *dʒ* is palatal ([Coronal] + [Velar]) is consistent with this fact, which could be sufficient evidence to assign that place specification to *dʒ* rather than another specification.

We have yet to (clearly) distinguish *f* and *v* from other consonants. Let us compare what features are detectable among the phonological labials, plus phonologically undetermined *v*. We see that, within that set, every possible combination of features is exploited (although we are not certain whether {p, h} are phonologically [Voiceless] or not). With respect to possible combinations of [Place] not dominating [Coronal] or [Velar], and also lacking the specification [Stop], there are only four remaining possibilities: [Voiceless] or [non-Voiceless], [Nasal] or [non-Nasal].

(56)	p	β	b	h	m	f	v
	Pl	Pl	Pl	Pl	Pl	Pl	Pl
	St	St	St	St	?		
		Ap		Ap			
	α	-α	-α	α	?	?	
					N		

Considerations of syllable structure suggest that neither *f* nor *v* are [Nasal]. In fact, owing to the above reasoning based on feature inheritance, a complete specification of *f* is available: it has a bare place specification, and an uncertain voicing specification. Obviously, *f* and *v* could be exactly the same except for their voicing specification. Phonological reasoning does not relate the voicing of *f* and *v* to that

of *p*, *b*, *h*, and *β*, so we would logically assign *f* and *v* the voicing values δ and $-\delta$, that is to say, distinct from each other but not relatable to the voicing of *p*, *b*, *h*, and *β*. It is also possible that *v* is assigned [Velar] place (but not [Coronal], since the combinatorics of [Stop] and [Approximant] for coronals has been exhausted). Thus the role for phonologically random feature assignment in Kerewe is very small.

A final issue regarding Kerewe features is the problem of the trigger of Spirantization: half of the suffixes which have initial /i/ trigger the rule. The question is, how do we distinguish those instances of /i/ which do not trigger the rule from those which do? Those which do not trigger the rule (/ik-/ ‘stative’, /il-/ ‘applied’, /isi-/ ‘causative’) are derivational suffixes with the shape VC(G), and those which do trigger the rule (/i-/ ‘causative’, /-i/ ‘nominalization’, /ile/ ‘perfective’) are derivational or inflectional suffixes that do not have that shape. There is no obvious phonological generalization that makes this distinction, nor is there a morphosyntactic generalization, so unless the rule simply enumerates the specific triggering suffixes, some arbitrary representational property is required to either trigger the rule or block it.

In substance-based phonological theories, this kind of problem is either resolved by invoking a diacritic feature such as [+D] which has no phonetic interpretation and only serves to distinguish those segments that trigger the rule, or else by invoking an abstract phonetic-featural distinction where, for instance, non-trigger *i* is underlyingly /ɪ/ and trigger *i* is /i/ (or vice-versa). In RSFP, phonological features are *all* abstract, and we only require phonologically distinctive behaviour to justify positing a feature. We have such behaviour here – then what feature distinguishes the /I/ which triggers spirantization from the /I/ which does not? There are plenty of gaps in feature combinatorics which allow two kinds of *i* to be distinguished, for example [Nasal], [Stop], or [Velar]. Nothing in the grammar favors one feature over the other, so speakers may be assumed to assign some feature at random.

4.4 Segments without rules

The theory of feature learning outlined here depends entirely on the behaviour of segments w.r.t. computations. We have to consider the possibility that the rules do not always uniquely identify all segments (as is the case for Kerewe), perhaps to the point that there are no phonological computations at all, so no basis for deciding between competing representations of any segments. This brings us into the contentious area of claims about what rules a grammar must contain. For example, it is not clear whether there are any segmental alternations in Vietnamese which would motivate phonological rules, but there are around two dozen consonants and nine vowels each of which need a distinct phonological representation.⁴⁶ Some system of features is needed to represent the following words.

⁴⁶Encoding historical data in synchronic phonology is not a valid basis for positing a phonological rule. It is perhaps less obvious that segment-minimalization in underlying forms is also not a valid basis for positing phonological rules. It is an open question whether there are “allophonic” rules in Vietnamese such as /ɔŋ/ → [ǎwŋ̃m], and why a child would learn this rule, rather than storing the invariant surface form.

(57) ti 'bureau' tu 'fourth' tu 'to drink'

In the worst case, the distinctions can be represented with a system of six features assigned arbitrarily to individual segments. We must also consider whether it is a fact of Vietnamese grammar that no words begin with two consonants, or end with two consonants.⁴⁷ This might justify rules constructing syllables, which may refer to a distinction “vowel” versus “consonant”. Similar gap-filling considerations could lead to discovering a rule excluding *b*, *d*, and *γ* and other consonants from syllable codas. Ultimately, there seems to be no fact of Vietnamese phonology determining what features distinguish *b*, *d*, *γ*. Analogously, Hawaiian appears to have no phonological alternations and cooccurrence restrictions seem to only suggest identifying “consonant” in order to say that there are no consonant clusters or final consonants. Since the segments of Hawaiian do not have “phonological patterns”, it does not matter what minimal assignment of features is employed to represent these segments.⁴⁸

It has been widely assumed, since SPE, that phonological grammars must contain computational statements that “account for” observable distributional tendencies – Morpheme Structure Conditions in SPE, perhaps surface Well-Formedness Conditions in theories such as OT. The underlying premise seems to be that given an inventory of *n* segments, all permutations of unbounded length constructed from those *n* segments should be equally-well attested. If there are significant attestation gaps (given a suitable theory of “significance”), we might presume that there are rules that govern distributional limits. Whether or not it is valid to hold the theory of computation responsible for such patterns is a completely separate question from the question of whether UG contains substantive information about particular features. In this article, I have not relied on static distributional properties to justify putative rules.⁴⁹ If it can be established that distributional gaps entail synchronic phonological rules, then such phonological rules would provide evidence for features. The important point is that distribution is not evidence for feature analysis; rules are evidence.

It is possible that the final featural analysis is influenced by learning artifacts. For example, *p* as distinct from *b* is marginal in Vietnamese, so at an early stage of acquisition, both sounds may be mapped to the same feature representation – they are not yet understood to be phonologically distinct segments. Further exposure to data may

⁴⁷It is clear that this is a fact about the language; the question is whether this is part of the grammar of Vietnamese. As has been repeatedly pointed out in the substance-free literature, there has been an excess of assumptions made in phonology to the effect that grammar contains reflexes of all forms of human sound-related behavior. Grammars do not encode all stateable observations about their languages.

⁴⁸UG provides analytic guidance. The alternative that the language has no features and instead has 13–18 atomic segments is precluded, in that the assumed theory of UG does not have a concept of “atomic segment”, it has conjunctions of features; and it is ruled out acquisitionally, in that decomposition into a smaller set of orthogonal features is simpler than positing a significantly larger set of segment-sized atoms.

⁴⁹I have appealed to syllabification in Kerewe, a construct which is phonologically motivated by various facts pertaining to vowel length, tone and reduplication.

correct the analysis, whereby some instances of assumed /P/ are distinguished from others: thus *p* and *b* could have the same features save for one distinction, because they were initially treated as *being* the same segment. Longitudinal evidence regarding development of segment perception in infants is sparse, so it is not possible to be more specific than to point out that if a phonological distinction is recovered – one assumed segment turns out to be two – remnants of earlier learning patterns may remain in later grammar. Ontogeny may partially recapitulate phylogeny. Ultimately, though, the theory does not mandate that the data must lead to a unique phonological analysis, to be accomplished by introducing various unmotivated auxiliary hypotheses.

5. CONCLUSIONS

It has been shown here that there is no logical requirement for UG to contain primitives expressing physical facts of speech – features need not be defined in terms of articulation or acoustics. There is no evidence that features are “defined” in grammar, nor that they are automatically assigned by any aspect of the language faculty. Instead, features are *used* distinctly to refer to how one segment is the same as or different from another segment in the grammar. There is no evidence for a UG limit on the number of (undefined) features, and any limit observed in a language emerges from the limited need for more features. However, UG clearly must contain formally-defined representational potentials. “Feature” is a concept of UG with a fixed formal nature. Likewise, grammatical computations have a fixed formal nature set by UG. These two aspects of UG, and the non-phonological ability to identify segments of a language, give rise to language-specific feature assignments, via the learned symbolic interface between the phonetic and phonological components. A key to a completely substance-free theory of phonology is recognizing that such computations are performed by a specific, highly-symbolic module in the mind. Processing and retaining physical inputs are performed by separate mental modules, which are not exclusively linguistic. The interfaces between those aspects of the mind and the phonological component are not part of UG and, I argue, are learned based on the formal requirements of creating a grammatical system that is accountable for perceptible facts which are outside of grammar.

We have seen that phonological features for Kerewe can be learned simply by reference to two considerations. First, when a set of sounds is identified by a rule, those sounds have a feature in common – if another set of sounds are excluded by a rule, those sounds lack the feature. Second, even in lieu of class behaviour in rules, the fact that a phonology contains distinct objects *p* and *b* means that some arrangement of features distinguishes those objects. Vietnamese seems to have no synchronic phonological alternations of the type $e \rightarrow \varepsilon / _ X$, $e \rightarrow r / Y _$, though there may be rules governing segment combinatorics whereby *uən*, *uəm* are possible syllable-final sequences but *ion*, *uxm* are not. Even if there are no rules in Vietnamese which treat *v* and *z* differently, they are independent sounds of the language, so must be represented with different features in the grammar. Since

phonological class behaviour is the primary driving force behind feature assignment, when there is no class behaviour but a distinction in sounds is still made, an interesting question arises. Is there any discernible pattern to feature assignment when the grammar is silent? Are available gaps in combinatorics exploited randomly (as presumed here), or are phonetic properties called on as a fallback method to reaching a uniform analysis, given a particular fact pattern which constitutes the primary linguistic data? The obvious theoretical question to address is: given a corpus of data constituting the basis for acquisition of a given language, must the theories of grammar and learning be expanded so that it is guaranteed that there is only one analysis of that data?

The theory of Formal Phonology does not impose such an a priori requirement on either grammatical theory or the theory of learning, and it certainly does not achieve uniformity of analysis by stipulating lists of substantive default assumptions. RSFP as a theory of features only maintains that the responsibility of the Language Acquisition Device is to construct the most economical possible grammar. To the extent that extragrammatical facts of speech such as acoustic similarity could be known to a child, such facts might influence the outcome of the formally random coin toss performed by learning theory in acquiring the feature-assigning interface.

REFERENCES

- Anderson, Stephen. 1974. *The organization of phonology*. New York: Academic Press.
- Bender, Byron. 1968. Marshallese phonology. *Oceanic Linguistics* 7(2): 16–35.
- Blaho, Sylvia. 2008. *The syntax of phonology: A radically substance-free approach*. Doctoral dissertation, University of Tromsø.
- Bradshaw, Mary. 1999. *A crosslinguistic study of consonant-tone interaction*. Doctoral dissertation, Ohio State University.
- Cho, Taehong, and Peter Ladefoged. 1999. Variation and universals in VOT: Evidence from 18 languages. *Journal of Phonetics* 27(2): 207–229.
- Choi, John D. 1992. Phonetic underspecification and target interpolation: An acoustic study of marshallese vowel allophony. *UCLA Working Papers in Phonetics* 82. Los Angeles: UCLA.
- Chomsky, Noam. 1965. *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 1980. Rules and representations. *Behavioral and Brain Sciences* 3(1): 1–61.
- Chomsky, Noam, and Morris Halle. 1968. *The sound pattern of English*. New York: Harper and Row.
- Clements, G. Nick. 1985. The geometry of phonological features. *Phonology Yearbook* 2: 225–52.
- Clements, G. Nick, and Elizabeth Hume. 1995. The internal organization of speech sounds. In *Handbook of Phonological Theory*, ed. John Goldsmith, 245–306. Oxford: Basil Blackwell.
- Clifton, John. 1976. Downdrift and rule ordering. *Studies in African Linguistics* 7(2): 175–194.
- Disner, Sandra. 1983. Vowel quality: the relation between universal and language-specific factors. *UCLA Working Papers in Phonetics* 58. Los Angeles: UCLA.
- Dresher, B. Elan. 2009. *The contrastive hierarchy in phonology*. Cambridge: Cambridge University Press.
- Dunn, Margaret H. 1993. *The phonetics and phonology of geminate consonants: A production study*. Doctoral dissertation, Yale University.
- Fromkin, Victoria. 1972. Tone features and tone rules. *UCLA Working Papers in Phonetics* 21: 50–75. Los Angeles: UCLA.

- Hale, Mark, Madelyn Kissock, and Charles Reiss. 2007. Microvariation, variation, and the features of Universal Grammar. *Lingua* 117(4): 645–665.
- Hale, Mark, and Charles Reiss. 2008. *The phonological enterprise*. Oxford: Oxford University Press, Oxford.
- Hamann, Silke. 2011. The phonetics-phonology interface. In *Continuum Companion to Phonology*, ed. Nancy Kula, Bert Botma and Kuniya Nasukawa, 202–224. London: Continuum.
- Jackendoff, Ray. 1990. *Semantic structures*. Cambridge, MA: MIT Press.
- Jakobson, Roman. 1949. On the identification of phonemic entities. *Travaux du Cercle Linguistique de Copenhague* V: 205–213.
- Jakobson, Roman, C. Gunnar M. Fant, and Morris Halle. 1952. *Preliminaries to Speech Analysis: the Distinctive Features and their Correlates*. Technical Report 13. Massachusetts: Acoustics laboratory, MIT.
- Johnson, Robert. 1975. *The role of phonetic detail in Coeur d'Alene phonology*. Doctoral dissertation, Washington State University.
- Kaye, Jonathan, Jean Lowenstamm, and Jean-Roger Vergnaud. 1985. The internal structure of phonological elements: A theory of charm and government. *Phonology* 2(1): 305–328.
- Khattab, Ghada, and Jalal Al-Tamimi. 2014. Geminate timing in Lebanese Arabic: The relationship between phonetic timing and phonological structure. *Laboratory Phonology* 5 (2): 231–269.
- King, Robert. 1969. *Historical linguistics and generative grammar*. Englewood Cliffs, NJ: Prentice-Hall.
- Ladefoged, Peter. 1964. *A phonetic study of West African languages – An auditory-instrumental survey*. Cambridge: Cambridge University Press.
- Lombardi, Linda. 1991. *Laryngeal features and laryngeal neutralization*. Doctoral dissertation, University of Massachusetts.
- Mielke, Jeff. 2008. *The emergence of distinctive features*. Oxford: Oxford University Press.
- Morén, Bruce. 2003. The Parallel Structures Model of feature geometry. *Working Papers of the Cornell Phonetics Laboratory* 15: 194–270. Los Angeles: UCLA.
- Odden, David. 2006. Phonology ex nihilo. Ms. <<https://doi.org/10.5281/zenodo.5512563>>
- Odden, David. 2013. Formal phonology. In *A Festschrift on the occasion of X Years of CASTL phonology and Curt Rice's Lth birthday*, ed. Sylvia Blaho, Martin Krämer and Bruce Morén-Duolljá. *Nordlyd* 40.1: 249–273.
- Odden, David. 2021. Phonological ontology. Ms. <<https://doi.org/10.5281/zenodo.5514165>>
- Peters, Ann. 1973. A new formalization of downdrift. *Studies in African Linguistics* 4(2): 139–153.
- Postal, Paul. 1968. *Aspects of phonological theory*. New York: Harper and Row.
- Reiss, Charles. 2003. Quantification in structural descriptions: attested and unattested patterns. *Linguistic Review* 20(2–4): 305–38.
- Roberts, Ian. 2016. Introduction. In *The Oxford handbook of Universal Grammar*, ed. Ian Roberts. Oxford: Oxford University Press.
- Sagey, Elizabeth. 1986. *The representation of features and relations in non-linear phonology*. Doctoral dissertation, Massachusetts Institute of Technology.
- Samuels, Bridget, Samuel Andersson, Ollie Sayeed, and Bert Vaux. 2022. Getting ready for primetime: Paths to acquiring substance-free phonology. *Journal of Canadian Linguistics/La Revue canadienne de linguistique* 67(4): 552–580.
- Vaux, Bert, and Bridget Samuels. 2015. Explaining vowel systems: Dispersion theory vs. natural selection. *The Linguistic Review* 32(3): 573–599.