Gorman, Kyle & Reiss, Charles. 2026. Get rich quick:
Why kids don’t need Occam’s Razor. *Glossa:
a journal of general linguistics* 11(1). pp. 1–31.
DOI: https://doi.org/10.16995/glossa.27793

Get rich quick: Why kids don’t need Occam’s Razor

Kyle Gorman, CUNY Graduate Center, New York, kylebgorman@gmail.com

Charles Reiss, Concordia University, Montréal, charles.reiss@concordia.ca

Most phonologists assume that phonological processes are encoded in terms of natural classes defined by sets of features. Thus, children acquiring a phonological grammar need some way to determine the natural classes identifying target and trigger segments. Exactly how children solve this *feature specification problem* (FSP) is unknown. Many phonologists endorse the view that the feature specifications should be *minimal* in the sense of containing the minimal number of feature specifications necessary to uniquely identify the class of segments, but no algorithm has been proposed to achieve this objective. We show that there is no guarantee that there is a unique minimal definition of every natural class. Then, drawing on recent work by Chen & Hulden (2018), we show that any algorithm which minimizes the number of features is intractable in the sense that it can only be solved by sifting through trillions of candidate natural classes. Recognizing that feature minimization is intractable, we propose an alternative objective, *maximization*, and a tractable algorithm that yields a unique result to implement it. Finally, we argue that the bias for minimal specification reflects a misapplication of Occam’s Razor to acquisition.

1 Introduction

Like many languages, Georgian has two laterals in complementary distribution: “plain” or light [l] occurs before the front vowels [i, e], and the “dark” or velarized [ƚ] occurs elsewhere.

(1) Georgian laterals (Robins & Waterson 1952):

leło	‘goal’
k’b̥iłs	‘tooth (dat.)’
tsʰoli	‘wife’
łɑmazad̥	‘prettily’
ɑłq’ɑ	‘siege’

The lightness or darkness of a Georgian lateral is completely predictable, so one can treat these two segments as allophones of a single underlying segment. Dark [ƚ] occurs before back vowels and all consonants, and thus one might choose to analyze light [l] as the allophone of dark /ƚ/ before front vowels. The generalization is stated informally below:

(2) ƚ → l / __ i, e

If one wishes to specify (2) more formally, in terms of natural classes, one has to decide exactly how general or specific the environment should be. Let us assume that Georgian vowels are featurally specified as follows:[1]

(3) Georgian vowel features (non-contrastive features omitted):

	/i/	/e/	/a/	/o/	/u/
Back	–	–	+	+	+
High	+	–	–	–	+
Low	–	–	+	–	–

Then, a few possibilities for the rule’s environment are:

(4) a. … / __ [–Back]
b. … / __ [–Back, –Low]
c. … / __ [–Low, –Round]
d. … / __ [–Back, –Low, –Round]
e. … / __ [–Back, –Low, –Round, –Nasal, –Creaky, +Voice, …]

Unfortunately, there is no obvious way to decide between these specifications on the basis of Georgian alone, because all Georgian –Back vowels are –Low, –Round, –Nasal, and so on,

---

[1] We discuss the status of [ɑ] in Georgian in §5.2. Note that while we use Back in this specification, we could just as well have employed Round.

so the specifications in (4) are all extensionally equivalent—that is, all of these specifications refer to /i, e/ in Georgian—and thus all of them are observationally adequate characterizations of the environment for this process. This remains true even if one assumes that features that are non-contrastive in Georgian (like Creaky) are not present at this stage of the phonological computation. If such features are absent, then one can ignore (4c–e), but it is still necessary to decide between (4a,b), which only refer to features, namely Back and Low, which are presumably contrastive in Georgian.

A similar point can be made about the natural class specification of the rule's target, although it only extensionally targets a single segment.[2] Under the analysis developed here, the only underlying lateral in Georgian is /ɫ/, and we are aware of no other processes introducing laterals in the language. Thus the following formulations are all also extensionally equivalent:

(5) a. [+Lateral] → ...
b. [+Lateral, +Sonorant] → ...
c. [+Lateral, +Sonorant, –Continuant] → ...
d. [+Lateral, +Sonorant, –Continuant, +Voice, ...] → ...

Discussing the Georgian laterals, Hale & Reiss (2008: 7) write that "[n]o language-internal evidence would bear on the matter of selecting the correct formulation of the rule". Let us call the problem of deciding how to represent the relevant natural classes governing phonological processes the *feature specification problem* (FSP).[3]

Thus far we assumed a rule-based formulation of the Georgian alternation, and will continue to use rules rather than constraints. We do so for ease of exposition, but it is done without any loss of generality because the FSP arises in any theory where phonological processes are represented via featurally-defined natural classes, and presumably, for all phonological processes in all spoken languages.[4] Indeed, the FSP is not even specific to phonology. In realizational approaches to morphology, including Distributed Morphology, there are rules of *referral* (e.g., *impoverishment*) which manipulate morphosyntactic features and rules of *realization* (or *exponence*) responsible for inserting morphs. Both types of rules, by hypothesis, are conditioned by natural classes defined by morphosyntactic feature specifications, and thus such theories instantiate a form of the FSP.

In our usage, an *objective* is a claim about what makes some solution to a problem (here, the FSP) preferable to others. For instance, in section 3 we show that much received wisdom

---

[2] We assume singleton targets and triggers still represent natural classes, an assumption we justify in §2.1.3.

[3] In rule-based frameworks, there is also the matter of determining how the structural change portion of the rule ought to be encoded. We do not discuss this issue here because we believe it is a much simpler task than the FSP. Given pairs of input/output segments like [ɫ] and [l], one can uniquely identify the change simply by computing the set differences between their featural specifications.

[4] Indeed, it is generally believed that the phonology of signed languages is also organized around natural classes (e.g., Sandler et al. 2011), so this problem apparently arises independent of modality.

assumes good solutions to the FSP are "minimal" in terms of the number of features specified. Minimization is a possible objective, but not the only one that can be imagined. An *algorithm*, on the other hand, is a computational procedure guaranteed to satisfy some objective. Below we also show that there exists no tractable algorithm for finding the minimal specification of a natural class—we also show there may not be a unique minimal specification at all—and this result leads us to reject minimization as an objective. Then in section 4, we propose an alternative objective and an efficient algorithm which implements it.[5] Finally, we use the term *heuristic* to refer to a computational procedure which, unlike an algorithm, cannot be guaranteed to satisfy any particular objective.

Given that many possible feature specifications, even for this simple pattern, are extensionally equivalent, it is important to note that a solution to the FSP—ultimately, in the form of an algorithm—is not a purely philosophical quest. We discern three groups of stakeholders.[6]

The first stakeholder is the child acquiring Georgian. Clearly, children do acquire Georgian, including the alternation described above, so under the assumption that phonological processes are encoded in terms of natural classes, the language acquisition device must include, among other things, an algorithm solving the FSP.

We acknowledge that different speakers of the socio-political construct "Georgian" may converge on subtly different I-language grammars. However, we put aside the possibility of interspeaker variation in such cases, because there is no reason to think these "Georgian" speakers acquire different representations of lateral allophony, nor is such a hypothesis falsifiable at present. To see why, it is helpful to compare the above example of the FSP to a superficially similar scenario. Han et al. (2007) claim there are two populations of Korean speakers: some whose grammars raise verbs to a higher tense projection, and others whose grammars do not. Han et al. suggest that individual variation in Korean reflects the fact that the critical examples distinguishing the two grammars—involving negation and quantifier scope—are vanishingly rare in the primary linguistic data. Yet they are able to establish interspeaker variation using a truth-value judgment task, and this task is applicable because there are Korean sentences which are interpreted differently by different populations. In contrast, there is at present no data from Georgian that distinguishes between the possible natural class specifications of the lateral alternation because the various specifications in (4–5) are extensionally equivalent. Putting aside the possibility of interspeaker variation, as we do here, is thus a justifiable scientific idealization. For similar reasons, we put aside the possibility of intraspeaker variation, i.e., the possibility that children acquiring Georgian converge on multiple extensionally equivalent rules or grammars

---

[5] Our notions of *objective* and *algorithm* are related to Marr's (1982) distinction between *functional* and *algorithmic* levels of cognitive description, respectively.

[6] Our discussion here is inspired by Hale & Reiss (2008: 18–22), who refer to the issues facing these three groups of stakeholders as the *human's problem*, the *linguist's problem*, and the *AI problem*, respectively.

differing in how they encode the natural classes involved in this alternation. Of course, it may be possible at some point in the future to establish the existence of interspeaker and/or intraspeaker variation with respect to natural class specification—e.g., with novel behavioral and/or neuroimaging techniques—and such results would naturally impact our understanding of the FSP.

The second stakeholder is the linguist studying Georgian. An informally specified rule like (2) is likely adequate as an instrumental description of this minor detail of Georgian phonology. Yet, many linguists share our assumption that phonological processes are encoded in terms of natural classes and are ultimately interested in how children acquire and encode these processes. Thus, they are interested in the objective and algorithm children use to solve the FSP.

The third stakeholder is the engineer who wishes to model the pronunciation of Georgian, e.g., for applications like text-to-speech synthesis systems. As for the descriptive linguist, an informal specification like (2) is adequate, but speech technologists have long recognized that natural classes—along with many other basic constructs from linguistic theory—provide a way to structure generalization and combat data sparsity, and the engineer who adopts the assumption of natural classes also requires a solution to the FSP. While this engineer may share some of the scientific concerns of the linguist, they need not worry whether or not the solution they adopt is actually used by children.

2 Formal preliminaries

Before we take up the FSP itself, we first motivate and define our notion of features, segments, natural classes, and rules. Then, with this background, we provide a formal definition of the FSP.

2.1 Features, segments, natural classes, and rules

The linguistic notion of a natural class has precursors as far back as Pāṇini (e.g., Kiparsky 1991) and Dionysius Thrax (e.g., Sen 2024), and plays a major role in generative phonology and its structuralist forebears. There are three closely-related intuitions underlying this notion. First, while phonological processes may target single segments, there may also be multiple target segments, sharing specific articulatory and/or acoustic properties. This motivates the use of featurally-defined natural classes to define targets of phonological processes. Second, processes which target multiple segments may apply articulatory and/or acoustic parallel changes to these targets. This motivates the use of phonological features to describe the changes triggered by processes, though the “change” portion of a phonological rule cannot be equated with a natural class, because it does not in any sense contain segments as members (e.g., Bale et al. 2020). Third, while processes may be triggered in the environment of individual segments, the same (or parallel) changes may occur when the target segments occur within the environment of various triggers, again sharing specific articulatory and/or phonetic properties. This motivates the use of featurally-defined natural classes to also define triggers of processes. These three intuitions are tightly intertwined. In particular, the notions of natural class targets and featurally-defined changes seem impossible to motivate independently, since one is needed to justify the other. Once again, we speak in terms of rules but similar intuitions can be stated in constraint-based theories insofar as these constraints are defined in featural terms.[7]

Our interpretation of natural classes and related notions below is mentalistic: we are concerned with them as they exist in the head of the child acquiring, or the adult who has acquired, Georgian. However, we note that these notions have some relevance even for instrumentalists, like the linguist interested in (merely) describing Georgian or the engineer interested in modeling Georgian pronunciation. These stakeholders could describe the full pattern by listing out pairs of substrings like /łe/ ⇝ [le], /ło/ ⇝ [lo], and so on, but feature-based natural class descriptions often result in far greater conciseness. However, we do not consider conciseness to be an objective in and of itself. In our view, rules in I-languages are concise only to the extent they are encoded by featurally-defined natural classes, just as lexical representations are concise—i.e., involve underspecification—only to the extent that such representations are selected by the protocols of the language acquisition device, whatever these are.[8] We see no reason to assume that the minimization of redundancy plays a role in such protocols, and we believe any putative reduction of redundancy must be preceded by a stage of maximally rich representations.

2.1.1 Features

In the form of substance-free phonology we espouse here (the “Concordia school” of e.g., Reiss 2017), features are innate, universal mental categories involved in the transduction to and from articulatory and acoustic phenomena. Sets of features define the targets, changes and environments (triggers) of rules. Feature specifications are abstract information structures of phonological computation, in contrast to theories in which they are imbued with articulatory or acoustic properties which are in some sense “visible” to the phonology. In line with the assumptions of Cognitive Phonetics (e.g., Volenec & Reiss 2018; 2019; 2025), the featural representations which are the output of phonological computation are in turn the input to a language-universal system of phonetic implementation. These assumptions have major consequences for the purview of phonology. They suggest, for instance, that many details of *coarticulation* are in fact computed by the Cognitive Phonetics module, as are effects other approaches relegate to a (poorly theorized) system of “language-specific phonetic implementation”, a notion we reject. For ease of exposition, we adopt common assumptions about

---

[7] In contrast, Flemming (2004) proposes to define natural classes in terms of the interaction of different markedness constraints. However, for Flemming, these markedness constraints are themselves defined by feature specifications interpreted along the lines we give below. Thus Flemming’s proposal is essentially (re)definitional and maintains a principled role for our more-traditional notion of natural class.

[8] For further thoughts on these matters see Inkelas 1995; Hale & Reiss 2003, and Gorman & Reiss 2025a.

the featural inventory, and will at various points draw on a table of IPA symbols and their mapping onto feature specifications provided by Hayes (2009).

2.1.2 Segments

We assume that segments are themselves sets (i.e., unordered bundles) of valued features, and make no further assumptions about internal structure of these bundles (*pace* feature geometry). In other work (Gorman & Reiss 2025c), we endorse the view that these bundles are in turn linked to higher-level structures (i.e., X's on a timing tier), but such structure plays no role for the FSP and will be ignored henceforth. Similarly, in other work, we argue for universal binary (or *equipollent*) specification of features. Valued features are written as, e.g., +F or –G. We assume, without loss of generality, that segments must be *consistent*; that is, a segment cannot be specified both +F and –F for any feature F. However, we do not assume that segments must be *complete*, so underspecification is permitted.

(6) Segments as sets of features:

$$/\text{i}/ = \left\{\begin{array}{c} +\textsc{Syllabic} \\ -\textsc{Back} \\ -\textsc{Round} \\ +\textsc{High} \\ -\textsc{Low} \\ +\text{ATR} \\ \ldots \end{array}\right\} \qquad /\text{e}/ = \left\{\begin{array}{c} +\textsc{Syllabic} \\ -\textsc{Back} \\ -\textsc{Round} \\ -\textsc{High} \\ -\textsc{Low} \\ +\text{ATR} \\ \ldots \end{array}\right\} \qquad /\text{I}/ = \left\{\begin{array}{c} +\textsc{Syllabic} \\ -\textsc{Back} \\ -\textsc{Round} \\ -\textsc{Low} \\ +\text{ATR} \\ \ldots \end{array}\right\}$$

2.1.3 Natural classes

Like segments, we assume that natural classes are also defined by a set of valued features, but interpret them as sets of segments (i.e., sets of sets of valued features). The members of a natural class are all the segments whose specification is a *superset* of the set used for the class specification (e.g., Bale & Reiss 2018; Bale et al. 2020). In other words, a natural class is characterized as a partial description of a set of segments, each of which is subsumed by the set in the class representation. For example, the class of –Back, –Low vowels is the set of vowels that are supersets of the set {–Back, –Low}, as in (7).

(7) Natural classes as sets of segments:

$$\begin{bmatrix} -\textsc{Back} \\ -\textsc{Low} \end{bmatrix} = \left\{x \mid x \supseteq \left\{\begin{array}{c} -\textsc{Back} \\ -\textsc{Low} \end{array}\right\}\right\} = \{\text{i}, \text{ɪ}, \text{y}, \text{e}, \text{I}, \ldots\}$$

Some members of the class may be underspecified, such as /I/. Also note that an expression like (7) gives an intensionally defined set of segments, but not all such segments may be present in a given language. In lieu of the unwieldy set-builder notation, we use square brackets to denote natural classes of segments, as opposed to the curly brackets which denote segments and other simple sets of features. This notation captures a type-theoretic distinction ignored in earlier phonological models. This definition of natural classes admits natural classes with singleton (i.e., single-segment) extensions, as shown for fully-specified /i/ below.[9]

(8) Natural class with a singleton extension:

$$\begin{bmatrix} +\textsc{Syllabic} \\ -\textsc{Back} \\ -\textsc{Round} \\ -\textsc{Low} \\ +\text{ATR} \\ \ldots \end{bmatrix} = \left\{ x \mid x \supseteq \begin{Bmatrix} +\textsc{Syllabic} \\ -\textsc{Back} \\ -\textsc{Round} \\ -\textsc{Low} \\ +\text{ATR} \\ \ldots \end{Bmatrix} \right\} = \{\text{i}\}$$

We take this to be a welcome result. Phonological patterns which, extensionally speaking, target or are triggered by single segments are widely attested. Extensionally speaking, Finnish assibilation (e.g., Kiparsky 1973; 1993), for example, targets only [t] and is triggered only by a following [i], as in /halut-i/ ⇝ [halusi] 's/he wanted'. In contrast, Harms (1968: 26) defines natural classes as a "class of segments that can be specified with fewer features than any individual member of the class", which would seem to rule out (8), and Davenport & Hannahs (2020: 96) more explicitly state that natural classes "may consist of any number of sounds, from two, as in [t, d], to many, as in the class of all vowels". Given the existence of many putative "rules" like Finnish assibilation, we see no reason to posit that targets and triggers of phonological rules may either be single segments or natural classes. Natural classes alone will do.

Finally, we note that not all sets of segments can be expressed as natural classes. For example, given the feature specification (3), there is no natural class that contains the Georgian vowels /i, o/, but no other vowels in the language; any natural class that contains /i, o/ also contains /u, e/.

2.1.4 Rules

For exposition, we adopt a relatively traditional form of *feature-changing* phonological rule, commonly used in early generative phonology, which is defined by the target, environment, and change.[10] The target is defined by a natural class. For the cases of interest to us, the environment

---

[9] This interpretation of natural classes entails that not all singleton sets can be expressed as natural classes. Suppose that f, g are segments such that $f \subset g$, so that f is underspecified w.r.t. g. Then, there exists no natural class containing f but excluding g (Gorman & Reiss 2025c).

[10] In other work in the Logical Phonology framework (e.g., Bale et al. 2014; Gorman & Reiss 2025b), we argue that apparent feature-changing processes should be decomposed into separate subtraction and unification steps. However, this proposal, relating as it does to the semantics of the change portion of a rule, does not bear directly on the FSP, and we put it aside here.

is defined by triggers (themselves defined by natural classes) and their relative position with respect to the target, denoted by the familiar __ symbol in the rule specification. Finally, the change effected by a rule is defined by a feature specification, not a natural class.[11] This rule format is highly restrictive, and is unable to express rules which remove specified features (as in the delinking rules of later autosegmental theories) or which insert or delete full segments, but this is once again without loss of generality, as these types of rules pose no new difficulties. The same is true for environmental conditioning by syllable structure or stress. These simplifications allow us to focus on the logic of the FSP.

2.2 Innateness and universality

A full-throated defense of our assumption that the phonological feature inventory is innate and universal would take us too far afield—though see Chomsky & Halle 1965 and Reiss & Volenec 2022 for this—but we will defend it as a useful scientific idealization for isolating the FSP.

For sake of argument, suppose instead that distinctive features are induced, on a language-specific basis, from the primary linguistic data; this appears to be the position of Mielke (2008), Archangeli & Pulleyblank (2015), and Chabot (this issue), among others.

Consider a child acquiring Georgian who hears substrings like the […le…, …ɫs…, …li…, …ɫɑ…, …ɫq'…] found in the words listed in (1). First, on what basis is this child to first infer that different renditions of [l] are instances of the same segment, different renditions of [ɫ] are instances of a distinct segment, and similarly for the flanking segments [e, i, …]? We take these inferences to reflect innate categorical perception of segments, as documented by Janet Werker and colleagues (e.g., Werker et al. 1981; Werker & Tees 1984; Werker & Lalonde 1988; Polka & Werker 1994). Categorical perception, in turn, reflects the fact that the percepts themselves, the segments, are featurally specified. Second, supposing that the child has perceived renditions of the various segments in the primary linguistic data, on what basis are they to determine that [l] and [ɫ] are instances of the same higher-level category (i.e., surface realizations of /ɫ/)? Third, supposing that the child has identified the allophonic relationship between [l] and [ɫ], on what basis is the child to determine that the distribution is framed in terms of shared properties of [i, e], rather than separate rules for [i] and [e]? We take the latter two inferences to reflect an innate imperative to cast sound patterns in terms of features and natural classes whenever possible, an instance of what Fodor (1980: 33f.) calls *epistemic boundedness*: children can perceive natural class generalizations, and are incapable of not perceiving them. This, it seems to us, requires that features be epistemically prior to language-specific natural classes, since the latter

---

[11] As noted above, this is a departure from the notation of early generative phonology: we use curly braces rather than square brackets, because the change is not interpreted as a natural class and does not denote a set of segments. See Bale & Reiss 2018: §11.1 and Bale et al. 2020 for discussion.

are defined in terms of the former. More plainly, a learner can’t detect a pattern involving features without having the featural categories defining that pattern. We see no way to circumvent the chicken-and-egg dilemma of jointly (i.e., simultaneously) learning the featural representation of the relevant segments and the change and environment of the rule. Mielke, Archangeli & Pulleyblank, and other proponents of “emergent” features and natural classes, have little to say about these crucial matters. However, even if one is not convinced by our brief defense of feature nativism, it should be clear that it is more manageable to study the (yet unsolved) feature-specification problem without the additional complexities that emergentist accounts of features introduce.

2.3 The feature specification problem

We are at last in a position to provide a formal definition of the FSP. This definition is intended to be sufficiently generic to apply to any theory in which rules are defined in part by targets and triggers, which are in turn defined by natural classes (§2.1.4), natural classes are defined by sets of feature specifications interpreted as sets of segments (§2.1.3), and segments consist of sets of valued features (§2.1.2).

The FSP is a subproblem in the acquisition of phonological grammars encountered during the induction of individual rules. We suppose that, via some alignment procedure to be determined, the child has isolated lists of segments—represented as sets of valued features—behaving as targets of the candidate rule. Similarly, any observed trigger segments for this candidate rule have also been isolated into similar lists, one for each non-target position in the environment. We provide an example specification of this form for Georgian /ɫ/-fronting, using IPA characters as stand-ins for the full featural specifications.

(9) Alignment specification for Georgian /ɫ/-fronting:
Targets: ɫ
Triggers at position $t + 1$: i, e

A *feature specification function* is then a function from sets of feature specifications, such as the one that defines the targets above, to natural class specifications for those targets. This function is applied separately to the list of targets, and to each of the trigger positions in the specification, and the resulting natural classes are then combined with a featural representation of the change to build the phonological rule.

We are not in a position to provide a full algorithm by which children might induce a full phonological grammar. This problem is much bigger than the FSP itself, but it demands a solution to the FSP under the assumption that rule targets and triggers are defined by natural classes. We anticipate that the rule induction system will need to inspect candidate natural classes for various properties in the course of induction. First is the case where the input feature specifications share no common features, so the maximal natural class is simply [ ]. As Bale et al. (2020) note, this “universal” natural class contains all segments, because (by a theorem of set theory) every set is a superset of the empty set. Again, this is not specific to the objective or algorithm: the universal natural class is the only natural class that contains both input segments in this case. The rule induction system will have to determine whether the generalization from some to all segments is valid or invalid, perhaps by considering segments not observed in, or known not to be, members of the relevant class in the primary linguistic data.[12] In other cases, an algorithm in some cases returns natural classes whose members include elements other than the input, but not quite as broad as the universal [ ]. For example, in the Hayes (2009) feature table, which largely follows traditional assumptions regarding place or manner features, any natural class which contains coronal plosives like /t/ and dorsal fricatives like /x/ will necessarily also contain coronal fricatives like /s/ and dorsal plosives like /k/. Again, this follows not from the specific objective or algorithm used, but from the structure of the feature system and our insistence that rules are defined by natural classes, and any rule induction system will have to determine what to do with the resulting natural class. For instance, if the generalization from /t, x/ to /s, t, x, k/ is known to be inconsistent with the primary linguistic data, the rule induction system presumably must conclude that the putative “process” targeting, or triggered by, /t, x/ is actually decomposed into two or more rules. Thus our insistence that rules are defined by natural classes provides a principled mechanism for an induction algorithm to determine when a process can or cannot be handled by a single rule.

Potential solutions to the FSP, then, are algorithmic definitions of feature specification functions. In the next section, we analyze feature specification functions implementing the received wisdom that an optimal feature specification is “minimal”, and show that such functions necessarily have undesirable properties. Then, in section 4, we argue for an alternative solution, one producing a feature specification that is in some sense maximal.

3 Minimization

We first consider the hypothesis that natural classes are specified *minimally*, by which we mean that their intensional characterization, in terms of feature specifications, contains the minimal number of specifications needed to uniquely identify the extension (i.e., set of segments). Because of the way we propose to interpret natural classes (§2.1.3), this intensional minimality corresponds to a natural class which is extensionally maximal—for example, the natural class characterized with one valued feature as [+NASAL] contains more segments than the natural class characterized with two valued features as [+NASAL, +LABIAL]. Of course, some of the possible

---

[12] Volenec & Reiss (2020: 28f.) and Chomsky (2021) discuss one type of inconsistency of this sort: rules which have a broader intension than their extension because they are bled by an earlier rule.

segments denoted by a natural class description may not be present in the relevant language, or may not exist at the relevant stage of the derivation.

This minimization objective is hinted at, assumed, or argued for, by virtually all prior work of the FSP. However, we show that this objective does not correspond to a tractable, well-behaved cognitive algorithm.

3.1 Why kids don't need Occam's Razor

One belief that seems to be shared by virtually all generative phonologists is a preference for minimally-specified natural classes, all else held equal, reflecting a more-general preference for conciseness in the grammars of individual languages. We discern three related motivations for conciseness of this sort before presenting our heretical alternative.

The first motivation for conciseness, which is often made explicit in the writings of early generative phonology, and is also echoed in current work, is the post-war fascination with information theory, henceforth IT. When first announced (Shannon 1948), IT was seen as something of a "universal acid", applicable to virtually any scientific problem. IT also drew some of its cachet from Shannon's strategically important wartime work on cryptanalysis, encryption, and fire control, among other problems.[13] Indeed, the prestige of IT was so great that Shannon himself warned his theory had become something of a "scientific bandwagon".

> …workers in other fields should realize that the basic results of the subject [IT] are aimed at a very specific direction, a direction that is not necessarily relevant to such fields as psychology, economics, and other social sciences. (Shannon 1956)

Writing of his research during this era, Halle confesses that he and colleagues had made the mistake jumping on the bandwagon:

> In the 1950's I spent considerable time and energy on attempts to apply concepts of information theory to phonology. In retrospect, these efforts appear to me to have come to naught. For instance, my elaborate computations of the information content in bits of the different phonemes of Russian (Cherry et al. 1953) have been, as far as I know, of absolutely no use to anyone working on problems in linguistics. And today the same negative conclusion appears to be to be warranted about all my other efforts to make use of information theory in linguistics. (Halle 1975: 532)

---

[13] Given IT's military origins, the influence of the generous defense funding for scientific research that characterized the post-war period (see Chomsky 1993) should also not be discounted. For example, in a paper entitled "On the role of simplicity in linguistic descriptions", Halle (1961) acknowledges funding from the U.S. Army's Signal Corps, the Office of Naval Research, and the U.S. Air Force's Office of Scientific Research, Air Research and Development Command, as well as the National Science Foundation.

IT's influence on early generative linguistics extends beyond explicit computations of the bits of information, though, towards a general preference for grammatical conciseness. For instance, Chomsky states a preference for intensionally minimal and extensionally maximal classes:

> In particular, we prefer rules that apply to large classes of elements and that have a simple and brief specification of relevant context; and we prefer a set of rules in which the same classes of elements figure many times. These and other requirements are met if we define the complexity of the morphophonemic component in terms of the number of features mentioned in the rules… (Chomsky 1966: 542–543)

This theme is developed in *The Sound Pattern of English* (*SPE*; Chomsky & Halle 1968:§8.1), where this preference is expressed using a feature-counting evaluation metric which favors the most concise (in terms of features) empirically adequate grammar. Kenstowicz & Kisseberth, for example, provide a concise formulation of the *SPE* evaluation metric:

(10) Conciseness Condition: If there is more than one possible grammar that can be constructed for a given body of data, choose the grammar that is most concise in terms of the number of feature specifications. (Kenstowicz & Kisseberth 1979: 336)

The second motivation for grammatical conciseness is the notion that the child's task during language acquisition can be likened to a search for concise descriptions of the primary linguistic data. This idea is a manifestation of a more general idea that the best theory in any domain is the one that is most compact: as Chaitin (2006: 77) writes, "[a] useful theory is a compression of the data; compression is comprehension." This notion is commonly expressed in terms of the notion of *minimal description length* (MDL), which appears in many domains including computational cognitive models of language acquisition. For example, Rasin et al. (2021), in a simulation of phonological grammar learning, use MDL to motivate natural class induction. In contrast, our assumptions (see §2.1.3) require no further motivation for natural classes, because they are the only mechanism UG provides for encoding rule targets and triggers.

The third and final motivation for conciseness appears to be an appeal to Occam's Razor as applied to grammatical description. As Halle writes below:

> An important principle governing scientific accounts of all kinds is that they must be concise. Since the accounts proposed here are formulated in terms of feature complexes, conciseness in the present context implies that the accounts must use as few features as possible. (Halle 2019: 4)

Halle is not alone in appealing to Occam's Razor. The following quotations are all drawn from phonology textbooks:

> Between two grammars that generate the same set of descriptions of sound-meaning pairs, the one containing the smallest number of rules, and the rules that have the most general scope, is chosen; in other words, the simplest grammar. This is an absolutely fundamental point. All linguists proceed in this manner, whatever school they belong to… (Dell 1980: 138)

> Linguists, like other scientists, like to provide the most general statement of a rule or a principle. (Spencer 1996: 136)

> There are good reasons to include only just as many features in a rule as are needed. (Hayes 2009: 92)

> …one should use the minimum number of features required to specify all and only the sounds in the class. (Zsiga 2012: 282)

> The usual principle adopted in phonology is that simpler rules, which use fewer features, are preferable to rules using more features. […] rules are stated in terms of the simplest, most general classes of phonetically defined segments… (Odden 2013: 61, 66–67)

Odden continues, providing an explicit example in favor of intensional minimality and extensional maximality:

> In this example we only have direct evidence for the change after *m*, so it would be possible to restrict our rule to the more specific context "after *m*." But this would run counter to basic assumptions of science, that we seek the most general explanations possible, not the most restricted. (*ibid.*: 89)

In a slight variation on this theme, Halle & Clements present conciseness less as a scientific principle than a cross-linguistic tendency:

> …the languages of the world prefer to deal with sets of sounds that require few specified features for their identification rather than sets that require many. (Halle & Clements 1983: 9)

The only departure from this orthodoxy we are aware of comes from Zimmer, whose critique is directed at feature-counting metrics:

> The fairly widespread assumption that feature counting will automatically lead us to choose the preferable description from two or more competing ones, as long as they use the same features and the same conventions for writing rules, has never, to my knowledge, really been supported by detailed and convincing arguments… (Zimmer 1970: 97–98, cited in Hyman 1975: 113)

We acknowledge Occam's Razor as an unavoidable principle of scientific practice, but suggest that at least some of the quotations above apply it at the wrong level of description. In particular, we see no motivation for applying Occam's Razor to the grammars of individual languages. *Pace* Halle (2019)—who takes the Razor to apply to "scientific accounts of all kinds"—the natural classes and conditions that define the phonological rules of languages are ontologically different objects than operations like Merge. A linguist might posit Merge as a fundamental part of the language faculty, and thus (potentially) part of every grammar. If this is correct, then no child posits Merge or modifies it in the construction of a particular I-language; rather, Merge is given. In contrast, a linguist might posit a phonological process like devoicing of obstruents in the coda of syllables as a component of a particular language, and the linguist knows that the process is composed of primitives such as valued features, reference to precedence, structural relations among segments, and so on. The linguist builds the rule out of a hypothesized set of primitives, and—if they are interested in how children acquire and encode rules—they hope that the child has also constructed the corresponding I-language rule out of the same fundamental toolkit. The difference is that the linguist cannot be sure what the toolkit contains, whereas the child's hypothesis space is defined by the toolkit that UG provides, whatever that happens to be. This is why the linguist needs Occam's Razor and the child does not.[14] For the phonologist trying to decide if it is necessary to posit both ATR and RTR, or both BACK and FRONT as part of the universal feature set, Occam's Razor is relevant. A child, however, encodes representations using all the innately-given features, whatever they happen to be, and the child already knows what those features are.

These three bits of received wisdom cannot be fully disentangled. IT can be thought of as a method for operationalizing Occam's Razor, and MDL acquisition accounts are fundamentally also applications of IT. The MDL approach is also motivated in part by theories which draw strong parallels between cognitive development and the day-to-day work of scientists, such as Gopnik's (2003) "theory theory" proposed as an alternative to linguistic nativism. In the context of language acquisition, this theory likens the practices of linguists—including, apparently, the obsession with grammatical conciseness we document above—to that of "little linguists": children acquiring language. As we suggested above, such analogies are flawed, because there is a fundamental difference in the two domains under consideration. The scientist starts with no iron-clad domain-specific contentful concepts, only tentatively hypothesized categories and the general concepts of the science forming faculty, but has the ability to make and test hypotheses, and to reason from negative evidence. Because the scientists are starting from a blank slate of domain-specific concepts, they ought to only enrich the model when forced to.[15] In contrast, the child already has a well-defined space of concepts, but has no way of predicting which are redundant, and no reason to do so.

---

[14] Even if we grant Halle's contention for the sake of argument, we share Zimmer's skepticism that feature counting alone suffices to measure the relevant dimensions of complexity.

[15] Arguably the scientist starts with something worse than a blank slate because common-sense notions often have to be rejected to make progress.

Having established that there is a received wisdom holding that feature specifications are intensionally minimal, we first consider formal properties of minimization as an objective for the FSP. Then, building on work by Chen & Hulden (2018), we show that no tractable algorithm exists to compute this objective.

3.2 Uniqueness

According to Robins & Waterson, Georgian /ɫ/-fronting is triggered by a following /i, e/. Using the vowel feature matrix in (3), the natural class [–Back] contains /i, e/ and no other segments, and appears to be the unique, intensionally minimal characterization of that set of trigger segments which does not include any other Georgian phonemes. The minimization objective potentially provides a mechanism to decide between extensionally equivalent characterizations of the trigger natural class listed in (4). However this is not generally the case because there may be multiple ways to minimally specify a given natural class. We provide three concrete examples of non-uniqueness below, though we conjecture that such cases are quite common.

3.2.1 Maranungku

Maranungku has five vowel phonemes: /i, æ, ə, ɑ, ʊ/ (Tryon 1970). These phonemes can be uniquely identified by the three-feature system below:

(11) Maranungku vowel features (after Archangeli 1988):

	/i/	/æ/	/ə/	/ɑ/	/ʊ/
Back	–	–	–	+	+
High	+	–	–	–	+
Low	–	+	–	+	–

Despite the apparent simplicity of this system, there are at least two ways to minimally characterize certain singleton natural classes:

(12) a. /ɑ/: [+Back, –High] or [+Back, +Low]
b. /ʊ/: [+High, +Back] or [+Back, –Low]

For both /ɑ/ and /ʊ/, the two natural classes are equally minimal. For example, /ɑ/ is the only +Back vowel which is –High and which is +Low, but it is not the only +Back vowel, so its singleton natural class has characterizations using two features. Similarly, a minimal natural class identify characterizing /ʊ/ must mention that it is either +High or –Low to distinguish it from /ɑ/. Though Tryon's phonological description is quite brief, there is some evidence that the grammar of Maranungku needs to target /ʊ/ to the exclusion of the other vowels. For example, /ʊ/—and no other vowel—lowers to [o] (a "a higher mid back rounded vocoid") before [j], as in /pʊjtʲ/ ⇝ [pojtʲ] 'not to know'. The issue of uniqueness only becomes more acute if one imagines that non-contrastive features are also considered by the FSP. For instance, if we bring the feature Tense into consideration, then /ʊ/ can also be minimally characterized as either [+High, –Tense] or [+Back, –Tense].

3.2.2 Balearic Catalan

According to Wheeler (2005), the dialect of Catalan spoken on the Balearic islands has an eight-phoneme vowel inventory: /i, e, ɛ, a, ə, ɔ, o, u/. Wheeler proposes an unusual feature specification for this inventory which mixes privative and equipollent features, but here we give a more traditional specification based on the Hayes (2009) feature system:

(13) Balearic Catalan vowel features:

	/i/	/e/	/ɛ/	/a/	/ə/	/ɔ/	/o/	/u/
Front	+	+	+	–	–	–	–	–
High	+	–	–	–	–	–	–	+
Low	–	–	–	+	–	–	–	–
Round	–	–	–	–	–	+	+	+
Tense	+	+	–	–	–	–	+	+

Given an eight-vowel system like Balearic, there are 255 (= 2^8–1) non-empty subsets. Of these, 32 are extensionally distinct natural classes according to (13), and of these, seven have two equally-minimal but non-unique intensional specifications:

(14) a. /i, e/: [+Front, +Tense] or [–Round, +Tense]
b. /o, u/: [–Front, +Tense] or [+Round, +Tense]
c. /i/: [–Front, +High] or [+High, +Round]
d. /e/: [+Front, –High, +Tense] or [–High, –Round, +Tense]
e. /ə/: [–Front, –Round, –Tense] or [–Front, –Low, –Round]
f. /o/: [–Front, –High, +Tense] or [–High, +Round, +Tense]
g. /u/: [–Front, +High] or [+High, +Round]

Indeed, of the eight vowel phonemes, only three have a singleton class with a unique minimal characterization. (14a,b) shows that the issue with uniqueness is not confined to singleton classes, and cannot be eliminated by treating singleton target or trigger segments as somehow distinct from “true” natural classes (cf. the discussion in §2.1.3).

3.2.3 Barcelona Catalan

The vowel inventory of the Barcelona dialect of Catalan (Mascaró 1976) is much like the that of the Balearic dialects described by Wheeler, except that [ə] is not a phoneme in its own right. Mascaró (1976) proposes a novel equipollent feature specification for this inventory; it uses CP (“constricted pharnyx”) instead of Low, Back instead of Front, Labial instead of Round, and ATR instead of Tense:

(15) Barcelona Catalan vowel features (after Mascaró 1976: 40f.):

	/i/	/e/	/ɛ/	/a/	[ə]	/ɔ/	/o/	/u/
ATR	+	+	–	+	+	–	+	+
Back	–	–	–	+	+	+	+	+
CP	–	–	–	+	–	–	–	–
High	+	–	–	–	–	–	–	+
Labial	–	–	–	–	–	+	+	+

This new inventory and feature specification also has a number of minimal but non-unique feature specifications:

(16) a. /i/: [–Back, +High] or [+High, –Labial]
b. /ɛ/: [–ATR, –Back] or [–ATR, –Labial]
c. /ɔ/: [–ATR, +Back] or [–ATR, +Labial]
d. /u/: [+Back, +High] or [+High, +Labial]

Thus which classes have non-unique minimal specifications is a function of the choice of feature specification for a given inventory, but the problem of uniqueness does not necessarily go away under alternative specifications.

3.3 Tractability

Keeping in mind that natural classes may have multiple minimal specifications, what algorithm might allow us to find all of the minimal specifications? First consider a truly naïve algorithm, which iteratively considers all possible natural class specifications. If $\mathscr{F}$ contains n features, there are 3^n intensionally distinct natural classes to consider,[16] giving a worst-case complexity of $O(3^n)$. To give a sense of scale, there are almost 23 trillion natural classes to consider with Hayes’s 30-feature inventory, a number that is “essentially infinite” for our purposes. Even this feature system is insufficient, since it does not provide a complete set of features needed to distinguish less-common types of phonation (e.g., creaky voice) or non-pulmonic airflow mechanisms (e.g., clicks, ejectives, or implosives) which are contrastive in many languages.

Two tricks are potentially applicable to the naïve algorithm. First, it makes sense to consider natural classes with k valued features before considering those with $k + 1$ features, since our goal is to identify intensionally minimal classes. Second, it makes sense to limit the search to natural

---

[16] The 3 base here reflects the fact that for some feature F, a natural class as defined above may contain +F, –F, or be underspecified w.r.t. F. An alternative definition of natural classes using privative specification would have 2^n natural classes to consider, though such systems often use a larger inventory of features than equipollent systems.

classes whose specifications are consistent with the specifications of the segments in the extension of the natural class. Anticipating results presented in the following section, this limitation can be enforced by computing the maximal natural class—which can itself be computed efficiently—and then limiting the search to subsets of that class. However, these tricks do not affect the worst-case computational complexity of this naïve algorithm; the former may permit one to halt earlier on average, whereas the latter results in a smaller effective value of n.

In recent work Chen & Hulden (2018), henceforth C&H, prove that no algorithm can meaningfully improve on the exponential worst-case complexity of the naïve algorithm outlined above. While we essentially assume their results as presented, and refer the interested reader to their paper, we provide a high-level review of their proof, but we first consider its relevance to the FSP.

Tractability is the constraint on cognitive capacities which holds they must be implemented by biological-cognitive systems (such as human brains) subject to reasonable assumptions about these systems' resource limitations. One widely-debated hypothesis regarding tractability is the *P-cognition thesis* (e.g., Frixione 2001). An algorithm is said to run in polynomial time if it is guaranteed to terminate in time proportional to some polynomial function of the size of the input, and the P-cognition thesis holds that cognitive capacities are limited to functions which can be computed in polynomial time. While the P-cognition thesis has been disputed from many angles (see, e.g., van Rooij 2008), it seems to have been implicitly adopted by many in the computational phonology community. For example, Eisner (1997), Idsardi (2006), and Heinz et al. (2009) debate whether it is possible to find the optimal candidate in Optimality Theory in polynomial time. Similarly, Heinz (2010) and Chandlee et al. (2014) prove their algorithms for acquiring phonotactic generalizations and phonological mappings, respectively, are decidable in polynomial time. These debates and demonstrations draw their force from their authors' implicit adoption of the P-cognition thesis.

C&H's method involves relating the minimization objective to the well-studied *set cover* problem. In this problem, one is presented with a set U (the "universe") and a set of subsets of U denoted $\mathcal{S}$, and an integer k. One must then decide whether there exist k or fewer member sets of $\mathcal{S}$ whose union equals U. A worked example, using integers as set elements, is given below.

(17) Sample set cover problem:

Problem:

$$\begin{aligned} U &= \{1,2,3,4,5\} \\ \mathcal{S} &= \{\{1,2,3\},\{2,4\},\{3,4\},\{4,5\}\} \\ k &= 2 \end{aligned}$$

Solution: Yes; e.g., $\{\{1,2,3\} \cup \{4,5\}\} = U$.

Set cover is a canonical NP-complete problem (Karp 1972): while it is easy to verify that a proposed solution is correct, finding such a solution requires one to search through exponentially many such sets in the worst case. To prove that minimization is NP-complete, C&H perform a *reduction from* set cover. A reduction from problem A to problem B is a mapping which converts any instance of problem A into an instance of problem B, such that if the answer to A is yes if and only if the answer to B is also yes. If this mapping can be performed efficiently (i.e., in polynomial time), and if A is known to be hard, then B must be at least as hard as A. Thus, under the standard conjecture that $P \neq NP$, C&H's reduction proves that minimization inherits the NP-complete intractibility of set cover. C&H show this result holds for both equipollent and privative feature systems.

We assume C&H's results in full. To reiterate, the force of C&H's reduction is this: no polynomial-time algorithm is guaranteed to find the minimal natural class. Thus, under the P-cognition thesis, feature minimization is inherently intractable; it is not a candidate objective for the FSP. Some other objective is needed, and it is to this we now turn.

4 Maximization

Having rejected minimization as an objective, we propose a counterintuitive alternative: perhaps natural classes are specified maximally, with as rich a feature specification set as possible. This is the suggestion of Hale & Reiss (2003), though we will offer new arguments for it here.

4.1 Uniqueness

Consider again the case of Georgian /ɫ/-fronting, triggered by a following /i, e/ and no other segments. The maximal natural class is then the class which contains /i, e/ and which is defined by as many valued features as possible. This class is unique, as can be informally demonstrated via contradiction. Suppose, contra our claim that this class must be unique, there are actually two natural classes containing /i, e/, C_1 and C_2, which are distinct but both "maximal". Because they are distinct, there must be some specifications in C_1 but not C_2 (and vice versa). Then, we should be able to construct an "actual" maximal natural class C_3 which contains all the valued features in C_1 (resp. C_2) as well these additional valued features. But this contradicts our earlier assumption that C_1 and C_2 are maximal, showing our initial assumption—the existence of a maximal but distinct C_1 and C_2—was ill-founded.

As alluded to above, the maximal natural class is a natural starting point for minimization algorithms, like exhaustive branch-and-bound strategies or the greedy heuristics considered by Chen & Hulden. Such solutions must be subsets of the maximal natural class. For example, a minimal natural class containing /i, e/ must not contain any valued features not also in both /i, e/ or the class would not contain both /i/ and /e/. As such, maximization can be thought of as a *satisficing* strategy in the sense of Simon (1956), in contrast to the optimizing strategy represented by minimization.

We now provide a simple algorithm for computing the maximal natural class.

4.2 The algorithm

Let S and T be feature specifications for two segments. Then, the maximal natural class containing $\{S, T\}$ is defined by their intersection, $S \cap T$.

(18) **Algorithm:** valued feature intersection $R = S \cap T$:

```
$R \leftarrow \emptyset$
for $c\mathrm{F} \in S$ do
    if $c\mathrm{F} \in T$ then
        $R \leftarrow R \cup \{c\mathrm{F}\}$
    end if
end for
```

Since this algorithm loops over S, and a consistent segment has no more than n specifications, this algorithm runs in linear time, specifically $O(n)$. By any account, this means this algorithm is highly tractable.

Applying this algorithm to /i, e/ using the 24 vowel features provided by Hayes (2009: §4.10), one obtains a natural class specified by 23 features: every one except HIGH, the one feature for which /i, e/ have inconsistent specifications:

(19)
$$\begin{bmatrix} -\textsc{Back} \\ -\textsc{Low} \\ -\textsc{Round} \\ -\textsc{Nasal} \\ \ldots \end{bmatrix} = \left\{ \varsigma \mid \varsigma \supseteq \begin{Bmatrix} -\textsc{Back} \\ -\textsc{Low} \\ -\textsc{Round} \\ -\textsc{Nasal} \\ \ldots \end{Bmatrix} \right\}$$

While intersection is defined in terms of pairs of segments, intersection is associative (and commutative), so it easily generalizes to more than two segments. That is, given specifications S, T, U, it is the case that $S \cap T \cap U = (S \cap T) \cap U = S \cap (T \cap U)$ and so on. One special case is when the extension of the class is a single segment; here, one need only return the feature specification of that segment in the form of a natural class.

As we stated in the introduction, there are a number of different extensionally equivalent feature specifications for the class /i, e/: minimal, maximal, or various points between. However, there is one corner case where the maximal natural class obtained via intersection may have a different extension than a minimal class. Consider the maximal and minimal natural classes for the segments /i, e/, repeated from (4a,e):

(20) a. [–BACK]
b. [–BACK, –LOW, –ROUND, –NASAL, –CREAKY, –VOICE, …]

Both of these classes contain Georgian /i, e/ and no other segments in this language. Let us suppose, however, it were possible to introduce a new phoneme, say /æ/, into the mind of a Georgian speaker. This vowel is, presumably, a member of the minimally specified class (20a). But it is presumably not a member of (20b) since it is +LOW. Thus the two approaches make different predictions about whether the novel /æ/ would or would not trigger /ɫ/-fronting; minimal specification predicts that it would, whereas maximal specification predicts it would not. While there is some precedent for using informal production experiments to test similar hypotheses (e.g., Halle 1978), it is not immediately clear to us whether it is at present possible to truly introduce a novel phoneme into the mind of an adult speaker.

5 Discussion

It may come as a surprise to find that as counterintuitive an objective as maximization is empirically adequate and unique, and further that as simple an algorithm as intersection is a tractable way to realize it. Below, we consider some alternative objectives and algorithms, then consider the negative impact that an implicit bias in favor of minimal description has had on phonological analyses.

5.1 Alternative objectives and algorithms

It is possible to imagine other objectives beyond minimization or maximization, and indeed, some of these may correspond to tractable algorithms. At present, though, it is not clear what other objectives are worthy of consideration. As an illustration of this point, we take the unusual step of quoting from comments made by an anonymous reviewer of an earlier version of this work.

The reviewer first observes that an implementation of a branch-and-bound algorithm considered by Chen & Hulden does, in their experiments, find minimal specifications, "albeit with a fairly large search space". It would be more precise to say, though, that branch-and-bound has the same "essentially-infinite" search space as the infeasible brute-force algorithm and thus the same worst-case complexity. In other words, it is not an alternative to the minimization objective under the P-cognition thesis.

The reviewer then discusses C&H's experiments with a brute-force "greedy algorithm", a heuristic in our terminology. Of these experiments, the reviewer writes that "but even then the failure is one of positing 4 features in a case where 3 would have sufficed. That is actually encouraging!" However, one ought not to reason solely from apparent best case performance. First, C&H report experiments with just six natural classes using an English-only feature system which defines almost 23 trillion intensionally distinct natural classes, hardly a sufficient basis on which to determine the overall difficulty of the minimization problems children might be expected to encounter. Rather, one ought to refer to known *approximation ratios* for the greedy algorithm, i.e., upper bounds on how much worse a set cover solution found via the greedy algorithm is than exhaustive search (e.g., Slavík 1997); whether these admittedly abstract results ought to be be "encouraging" for proponents of minimization is unclear to us. One could, of course, argue that the heuristic itself is the proper solution to the FSP, but we know of no evidence that children pursue a minimization objective in the first place.

Finally, the reviewer writes that "a specialized feature minimization algorithm is likely to yield better results, for example by prioritizing coarse features like consonantal over more fine-grained ones like strident. Features are not the same as the arbitrary sets in the set cover problem, some features are more important than others. […] I fully expect the problem to be highly tractable under these parameters". There are two possible ways to interpret this suggestion. First, one could interpret this suggestion as type of greedy heuristic, in which "coarse" features are given higher precedence than "fine-grained" ones. Alternatively, many feature systems consist of specifications which are proper subsets of each other, and one could potentially take advantage of this type of subsumption structure while solving a set cover problem. For example, if all +STRIDENT segments are +CONSONANTAL but not all +CONSONANTAL segments are +STRIDENT, as Hayes (2009) assumes, then the extension of [+STRIDENT] is a proper subset of the extension of [+CONSONANTAL]. Translating this to the minimization problem, if the intersection contains +STRIDENT one can safely discard +CONSONANTAL from consideration. However, such structure has no effect on the overall tractability of the minimization objective. We have chosen not to pursue either suggestion, since we are reluctant at this early stage to commit to any account which closely depends on the vagaries of specific featural systems.

5.2 Minimization bias

The bias in favor of minimal description can add unnecessary, unmotivated complexity to phonological analyses. We demonstrate this by reviewing one detail of Georgian /ɫ/-fronting and introducing Malayalam palatalization. Like our discussion of non-uniqueness, these case studies are meant to be exemplary rather than exhaustive, and we suspect many examples of a similar form can be found in the literature.

5.2.1 Georgian again

Above, we followed Robins & Waterson, henceforth R&W, in their assumption that the Georgian low vowel phoneme they write as /a/—which is not a trigger for /ɫ/-fronting—is +BACK. R&W write that "L was pronounced with a decided 'dark,' velarized quality in all positions except before *e* and *i*; its dark quality before *a* is the phonological justification for grouping *a* with *o* and *u* as a back vowel." (R&W: 63). However, they also describe the *a* phoneme as being "generally of front quality” (R&W: 59). Perhaps R&W imagine this segment is essentially /ɑ/ but somehow realized as the front low vowel [a]. The following ordered rules implement this analysis.

(21) Georgian /ɫ/-fronting (after R&W; to be revised):

$$\begin{bmatrix} +\textsc{Lateral} \\ \dots \end{bmatrix} \rightarrow \left\{ -\textsc{Back} \right\} \ / \ \_\_ \ \left[-\textsc{Back} \right]$$

(22) Georgian redundancy rule (to be rejected):

$$\begin{bmatrix} +\textsc{Back} \\ +\textsc{Low} \end{bmatrix} \rightarrow \left\{ -\textsc{Back} \right\}$$

We are aware of no evidence for this two-stage analysis of /ɫ/-fronting, and many phonologists, starting with Kiparsky (1968), propose principles intended to rule out exactly this kind of analysis because of its apparent “abstractness”. Regardless of the merits of these particular arguments, our assumption of a universal transduction between the featural output of the grammar and speech (§2.1.1) rules out the possibility that (22) is a rule of language-specific phonetic interpretation. So why do R&W pursue this /ɑ/ ⇝ [a] analysis in the first place? We believe the answer is minimization bias *avant la lettre*. If the Georgian low vowel is the +Back /ɑ/ when (21) applies, then the trigger for fronting must only be the “minimal” –Back. If, however, the vowel is the –Back /a/, then the trigger must also be –Low, and the rule requires another specified feature. Under our stated assumptions though, this argument has no force, and our rejection of minimization seems to argue against the analysis given above. Instead, we assume this segment is underlyingly the –Back /a/, just as it is on the surface, and propose the /ɫ/-fronting rule is conditioned by both –Back, –Low, and other features shared by /i, e/. Of course, this simpler analysis has no use for (22).

5.2.2 Malayalam

In Malayalam, simple velar onsets palatalize after /i, e, a/ but not after /u, o/.

(23) Malayalam palatalization:

a.	kutːikʲə	‘child (dat.)’	(Mandal 2023)
	wekːʲal	‘cooking’	
	kanakʲam	‘gold’	
b.	en̪t̪ukoɳʈaːɳə	what.with.it	(Mohanan 1996: 425)
	ɦirɨd̪roːgat̪t̪inor̪u	heart.disease.Acc.one	
	kuːʈut̪alaːɭukaɭe	more.people.Acc	

Mohanan & Mohanan (1984: 586, fn. 24), henceforth M&M, formulate palatalization as copying of –Back from the vowel trigger to the velar target, except they take /a/ to be –Back.

(24) Malayalam palatalization (after M&M; to be revised):

$$\begin{bmatrix} -\textsc{High} \\ -\textsc{Continuant} \end{bmatrix} \rightarrow \left\{-\textsc{Back}\right\} \;/\; \overset{\substack{\mathrm{V} \\ |}}{\begin{bmatrix} -\textsc{Back} \end{bmatrix}} \;\_\_$$

This rule is "assimilatory" in the sense that –Back is present both in target and trigger. Yet, M&M (*loc. cit.*) write that the segment transcribed above as [a] is "[p]honetically…in fact a back vowel". Presumably, M&M are assuming that this vowel is underlyingly –Back /a/ at the time which palatalization applies, but realized in all cases as the +Back [ɑ].

(25) Malayalam redundancy rule (to be rejected):

$$\begin{bmatrix} -\textsc{Back} \\ +\textsc{Low} \end{bmatrix} \rightarrow \left\{+\textsc{Back}\right\}$$

In other words, M&M's treatment of Malayalam palatalization is more or less the mirror image of R&W's analysis of Georgian /ɫ/-fronting, and it is subject to the exact same critique: it introduces unnecessary derivational complexity, namely (25). As a simpler alternative we propose that the low vowel is +Back /ɑ/, as it is on the surface, and restate the rule of palatalization as follows:

(26) Malayalam palatalization (revised; after Mandal 2023):

$$\begin{bmatrix} -\textsc{High} \\ -\textsc{Continuant} \end{bmatrix} \rightarrow \left\{-\textsc{Back}\right\} \;/\; \overset{\substack{\mathrm{V} \\ |}}{\begin{bmatrix} -\textsc{Round} \end{bmatrix}} \;\_\_$$

Such an approach eliminates the need for (25). The fact that this rule is no longer conceived as an assimilation rule is not a problem for the substance-free approach we adopt. In fact, the revised Malayalam palatalization rule is a clear manifestation of one of the fundamental ideas of substance-free phonology:

(27) Substance-Freeness of Structural Change (Dabbous et al. 2024): The features added to segments by the application of a rule need not be found in the rule environment.

There is no need to reify—or even to define—taxonomic artifacts like assimilation, dissimilation, fortition, palatalization, etc. Environments need not provide the features to the target since the rules are computationally arbitrary in this sense.

6 Conclusions

We have defined the FSP, and shown that despite the received wisdom, the minimization of features in rules and representations during the course of acquisition is computationally intractable, and there may not be a unique minimal solution. One alternative proposal, maximization via intersection, is shown to be empirically adequate, tractable, and guaranteed to provide a unique solution. It is quite likely that intersection is not the only tractable solution, but it is the simplest satisfactory solution we can conceive of at present. Alternative proposals should demonstrate empirical adequacy (and convergence guarantees, if relevant), tractability (ideally, polynomial time complexity), and uniqueness. Potential alternatives that meet these conditions may require greater complexity than intersection computations, so such proposals should also justify whatever additional resources they require. Until such proposals are made and until they are properly analyzed, intersection is the only game in town. Our explicit claims could potentially guide future behavioral and neuroimaging studies as new technologies become available.

We reiterate that the view that children "get rich quick": encoding of rules and representations does not proceed from sparse to rich. This dovetails with the considerations of phonological acquisition based on the subset principle of Hale & Reiss (2003), and the analysis of child speech as reflecting a massive degree of performance effects compared to adult speech, rather than "child phonology", simplified representations, or the absence of marked structures *vis-à-vis* adults (Hale & Reiss 1998; 2008). This is not to suggest that children's lexical representations are immediately faithful copies of adult representations; the point is that children must have immediate access to all the features provided by UG, so they can get rich representations quickly. The infeasibility of minimization procedures suggests that children don't just get rich quick, but they stay rich, too.

We argue it is a widespread category error among linguists and other cognitive scientists to equate the application of Occam's Razor in the work of a scientist with the elimination of redundant features from the formulation of a phonological rule or lexical entry. Unlike the scientist, the scope and limits of the learner's grammatical apparatus is defined in advance; this knowledge, not the Razor, restricts the hypothesis space for the learner. A hypothetical child who chooses to hitch a ride on Shannon's bandwagon and minimize feature specifications would be faced with an intractable problem.

Abbreviations

Abbreviations used: ACC.: accusative; DAT.: dative.

Acknowledgements

We thank audiences at the 47th Penn Linguistic Colloquium (PLC), the 46th Generative Linguistics in the Old World Colloquium (GLOW), the MorrisHalle@100 Conference, Stony Brook University, and Rutgers University.

Competing interests

The authors have no competing interests to declare.

References

Archangeli, Diana. 1988. Aspects of underspecification theory. *Phonology* 5(2). 183–207. DOI: https://doi.org/10.1017/S0952675700002268

Archangeli, Diana & Pulleyblank, Douglas. 2015. Phonology without universal grammar. *Frontiers in Psychology* 6. 12–29. DOI: https://doi.org/10.3389/fpsyg.2015.01229

Bale, Alan & Papillon, Maxime & Reiss, Charles. 2014. Targeting underspecified segments: a formal analysis of feature-changing and feature-filling rules. *Lingua* 148. 240–253. DOI: https://doi.org/10.1016/j.lingua.2014.05.015

Bale, Alan & Reiss, Charles. 2018. *Phonology: A formal introduction*. MIT Press.

Bale, Alan & Reiss, Charles & Shen, David Ta-Chun. 2020. Sets, rules and natural classes: { } vs. [ ]. *Loquens* 6(2). e06. DOI: https://doi.org/10.3989/loquens.2019.065

Chaitin, Gregory. 2006. The limits of reason. *Scientific American* 294(3). 74–81. DOI: https://doi.org/10.1038/scientificamerican0306-74

Chandlee, Jane & Eyraud, Rémi & Heinz, Jeffrey. 2014. Learning strictly local subsequential functions. *Transactions of the Association for Computational Linguistics* 2. 491–503. DOI: https://doi.org/10.1162/tacl_a_00198

Chen, Hubie & Hulden, Mans. 2018. The computational complexity of distinctive feature minimization in phonology. In *Proceedings of NAACL-HLT 2018*, 542–547. DOI: https://doi.org/10.18653/v1/N18-2086

Cherry, E. Colin & Halle, Morris & Jakobson, Roman. 1953. Towards the logical description of languages in their phonemic aspect. *Language* 29(1). 34–46. DOI: https://doi.org/10.2307/410451

Chomsky, Noam. 1966. Explanatory models in linguistics. *Studies in Logic and the Foundations of Mathematics* 44. 528–550. DOI: https://doi.org/10.1016/S0049-237X(09)70617-2

Chomsky, Noam. 1993. The Pentagon system. *Z Magazine* 6(2). 25–35.

Chomsky, Noam. 2021. Simplicity and the form of grammars. *Journal of Language Modelling* 9(1). 5–15. DOI: https://doi.org/10.15398/jlm.v9i1.257

Chomsky, Noam & Halle, Morris. 1965. Some controversial questions in phonological theory. *Journal of Linguistics* 1(2). 97–138. DOI: https://doi.org/10.1017/S0022226700001134

Chomsky, Noam & Halle, Morris. 1968. *The sound pattern of English*. Harper & Row.

Dabbous, Rim & Leduc, Marjorie & Reiss, Charles & Shen, David Ta-Chun. 2024. Locality is epiphenomenal: Adjacency is opaqueness. In *Proceedings of the 2024 annual conference of the Canadian Linguistic Association*.

Davenport, Michael & Hannahs, S. J. 2020. *Introducing phonetics and phonology*. Routledge 4th edn. DOI: https://doi.org/10.4324/9781351042789

Dell, François. 1980. *Generative phonology and French phonology*. Cambridge University Press.

Eisner, Jason. 1997. Efficient generation in primitive Optimality Theory. In *35th annual meeting of the Association for Computational Linguistics and 8th conference of the European chapter of the Association for Computational Linguistics*, 313–320. DOI: https://doi.org/10.3115/976909.979657

Flemming, Edward. 2004. Deriving natural classes in phonology. *Lingua* 115(3). 287–309. DOI: https://doi.org/10.1016/j.lingua.2003.10.005

Fodor, Jerry. 1980. Reply to Putnam. In Piattelli-Palmarini, Massimo (ed.), *Language and learning: The debate between Jean Piaget and Noam Chomsky*, 325–334. Harvard University Press.

Frixione, Marcello. 2001. Tractable competence. *Minds and Machines* 11. 379–397. DOI: https://doi.org/10.1023/A:1017503201702

Gopnik, Alison. 2003. The theory theory as an alternative to the innateness hypothesis. In Antony, Louise & Hornstein, Norbert (eds.), *Chomsky and his critics*, 238–254. Wiley-Blackwell. DOI: https://doi.org/10.1002/9780470690024.ch10

Gorman, Kyle & Reiss, Charles. 2025a. How not to acquire exchange processes in Logical Phonology. In *Proceedings of the 2025 annual conference of the Canadian Linguistic Association*.

Gorman, Kyle & Reiss, Charles. 2025b. Metaphony in Substance Free Logical Phonology. *Phonology* in press.

Gorman, Kyle & Reiss, Charles. 2025c. Natural class reasoning in segment deletion rules. In *Proceedings of the 56th annual meeting of the North East Linguistics Society*. In press.

Hale, Mark & Reiss, Charles. 1998. Formal and empirical arguments concerning phonological acquisition. *Linguistic Inquiry* 29(4). 656–683. DOI: https://doi.org/10.1162/002438998553914

Hale, Mark & Reiss, Charles. 2003. The subset principle in phonology: Why the tabula can't be rasa. *Journal of Linguistics* 39(2). 219–244. DOI: https://doi.org/10.1017/S0022226703002019

Hale, Mark & Reiss, Charles. 2008. *The phonological enterprise*. Oxford University Press. DOI: https://doi.org/10.1093/oso/9780199533961.001.0001

Halle, Morris. 1961. On the role of simplicity in linguistic description. In *Structure of language and its mathematical aspects: Proceedings of the XIIth symposium in applied mathematics*, 89–94. DOI: https://doi.org/10.1090/psapm/012/9978

Halle, Morris. 1975. Confessio grammatici. *Language* 51(3). 525–535. DOI: https://doi.org/10.2307/412886

Halle, Morris. 1978. Knowledge unlearned and untaught: What speakers know about the sounds of their language. In Halle, Morris & Bresnan, Joan & Miller, George A. (eds.), *Linguistic theory and psychological reality*, 294–303. MIT Press.

Halle, Morris. 2019. Aspects of the morphophonology of the verb in Latin and in German and English. *Linguistic Inquiry* 50(1). 3–12. DOI: https://doi.org/10.1162/ling_a_00313

Halle, Morris & Clements, G. N. 1983. *Problem book in phonology: A workbook for introductory courses in linguistics and in modern phonology*. MIT Press.

Han, Chung-hye & Lidz, Jeffrey & Musolino, Julien. 2007. Verb-raising and grammar competition in Korean: Evidence from negation and quantifier scope. *Linguistic Inquiry* 38(1). 1–47. DOI: https://doi.org/10.1162/ling.2007.38.1.1

Harms, Robert T. 1968. *Introduction to phonological theory*. Prentice-Hall.

Hayes, Bruce. 2009. *Introductory phonology*. John Wiley & Sons.

Heinz, Jeffrey. 2010. Learning long-distance phonotactics. *Linguistic Inquiry* 41(4). 623–661. DOI: https://doi.org/10.1162/LING_a_00015

Heinz, Jeffrey & Kobele, Gregory M. & Riggle, Jason. 2009. Evaluating the complexity of Optimality Theory. *Linguistic Inquiry* 40(2). 277–288. DOI: https://doi.org/10.1162/ling.2009.40.2.277

Hyman, Larry M. 1975. *Phonology: Theory and analysis*. Holt, Rinehart and Winston.

Idsardi, William. 2006. A simple proof that Optimality Theory is computationally intractable. *Linguistic Inquiry* 37(2). 271–275. DOI: https://doi.org/10.1162/ling.2006.37.2.271

Inkelas, Sharon. 1995. The consequences of optimization for underspecification. In *Proceedings of the North East Linguistic Society 25*, 287–302.

Karp, Richard M. 1972. Reducibility among combinatorial problems. In Miller, Raymond E. & Thatcher, James W. (eds.), *Complexity of computer computations*, 85–103. Plenum. DOI: https://doi.org/10.1007/978-1-4684-2001-2_9

Kenstowicz, Michael & Kisseberth, Charles. 1979. *Generative phonology: Theory and method*. Academic Press.

Kiparsky, Paul. 1968. *How abstract is phonology?* Indiana University Linguistics Club.

Kiparsky, Paul. 1973. Phonological representations: Abstractness, opacity, and global rules. In Fujimura, Osamu (ed.), *Three dimensions of linguistic theory*, 57–86. TEC Company.

Kiparsky, Paul. 1991. Economy and the construction of the Śivasūtras. In Deshpande, Madhav & Bhate, Saroja (eds.), *Pāṇinian studies: Professor S. D. Joshi felicitation volume*, 239–261. University of Michigan Center for South and Southeast Asian Studies.

Kiparsky, Paul. 1993. Blocking in non-derived enviroments. In Hargus, Sharon & Kaisse, Ellen (eds.), *Studies in Lexical Phonology*, 277–313. Academic Press. DOI: https://doi.org/10.1016/B978-0-12-325071-1.50016-9

Mandal, Sayantan. 2023. The cognitive phonetics of velar palatalization. In *Proceedings of WECOL 2022*, 8–19.

Marr, David. 1982. *Vision: A computational investigation into the human representation and processing of visual information*. Freeman.

Mascaró, Joan. 1976. *Catalan phonology and the phonological cycle*. Massachusetts Institute of Technology dissertation.

Mielke, Jeff. 2008. *The emergence of distinctive features*. Oxford University Press. DOI: https://doi.org/10.1093/oso/9780199207916.001.0001

Mohanan, K. P. 1996. Malayalam writing. In Daniels, Peter T. & Bright, William (eds.), *The world's writing systems*, 420–425. Oxford University Press.

Mohanan, K. P. & Mohanan, Tara. 1984. Lexical phonology of the consonant system in Malayalam. *Linguistic Inquiry* 15(4). 575–602.

Odden, David. 2013. *Introducing phonology*. Cambridge University Press 2nd edn. DOI: https://doi.org/10.1017/CBO9781139381727

Polka, Linda & Werker, Janet F. 1994. Developmental changes in perception of nonnative vowel contrasts. *Journal of Experimental Psychology: Human Perception and Performance* 20(2). 421–435. DOI: https://doi.org/10.1037/0096-1523.20.2.421

Rasin, Ezer & Berger, Iddo & Lan, Nur & Shefi, Itamar & Katzir, Roni. 2021. Approaching explanatory adequacy in phonology using Minimum Description Length. *Journal of Language Modelling* 9(1). 17–66. DOI: https://doi.org/10.15398/jlm.v9i1.266

Reiss, Charles. 2017. Substance free phonology. In Hannahs, S. J. & Bosch, Anna (eds.), *The routledge handbook of phonological theory*, 425–452. Routledge. DOI: https://doi.org/10.4324/9781315675428-15

Reiss, Charles & Volenec, Veno. 2022. Conquer primal fear: Phonological features are innate and substance free. *Canadian Journal of Linguistics* 67(4). 581–610. DOI: https://doi.org/10.1017/cnj.2022.35

Robins, R. H. & Waterson, Natalie. 1952. Notes on the phonetics of the Georgian word. *Bulletin of the School of Oriental and African Studies* 141(1). 55–72. DOI: https://doi.org/10.1017/S0041977X00084196

van Rooij, Iris. 2008. The tractable cognition thesis. *Cognitive Science* 32(6). 939–984. DOI: https://doi.org/10.1080/03640210801897856

Sandler, Wendy & Aronoff, Mark & Meir, Irit & Padden, Carol. 2011. The gradual emergence of phonological form in a new language. *Natural Language & Linguistic Theory* 29. 503–543. DOI: https://doi.org/10.1007/s11049-011-9128-2

Sen, Ranjan. 2024. The Greco-Roman tradition. In Dresher, B. Elan & van der Hulst, Harry (eds.), *The Oxford history of phonology*, 109–133. Oxford University Press. DOI: https://doi.org/10.1093/oso/9780198796800.003.0006

Shannon, Claude E. 1948. A mathematical system of communication. *Bell System Technical Journal* 27(3). 379–423. DOI: https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Shannon, Claude E. 1956. The bandwagon. *IRE Transactions on Information Theory* 2(1). 3. DOI: https://doi.org/10.1109/TIT.1956.1056774

Simon, Herbert A. 1956. Rational choice and the structure of the environment. *Psychological Review* 63(2). 129–138. DOI: https://doi.org/10.1037/h0042769

Slavík, Petr. 1997. A tight analysis of the greedy algorithm for set cover. *Journal of Algorithms* 25(2). 237–254. DOI: https://doi.org/10.1006/jagm.1997.0887

Spencer, Andrew. 1996. *Phonology: Theory and description*. Blackwell.

Tryon, D. T. 1970. *An introduction to Maranungku (Northern Australia)*. Australian National University.

Volenec, Veno & Reiss, Charles. 2018. Cognitive Phonetics: The transduction of distinctive features at the phonology-phonetics interface. *Biolinguistics* 11. 251–294. DOI: https://doi.org/10.5964/bioling.9089

Volenec, Veno & Reiss, Charles. 2019. The intervocalic palatal glide in Cognitive Phonetics. In *Proceedings of the forty-ninth meeting of the North East Linguistics Society*, 255–264.

Volenec, Veno & Reiss, Charles. 2020. Formal generative phonology. *Radical: A Journal of Phonology* 2. 1–148.

Volenec, Veno & Reiss, Charles. 2025. Cognitive Phonetics: The universal phonology-phonetics interface. In Fábregas, Antonio & Mayol, Laia & Prystauka, Yanina (eds.), *Cambridge handbook of linguistic interfaces*, in press. Cambridge University Press.

Werker, Janet F. & Gilbert, John H. V. & Humphrey, Keith & Tees, Richard C. 1981. Developmental aspects of cross-language speech perception. *Child Development* 52(1). 349–355. DOI: https://doi.org/10.2307/1129249

Werker, Janet F. & Lalonde, Chris E. 1988. Cross-language speech perception: Initial capabilities and developmental change. *Developmental Psychology* 24(2). 672–683. DOI: https://doi.org/10.1037//0012-1649.24.5.672

Werker, Janet F. & Tees, Richard C. 1984. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7(1). 49–63. DOI: https://doi.org/10.1016/S0163-6383(84)80022-3

Wheeler, Max W. 2005. *The phonology of Catalan*. Oxford University Press. DOI: https://doi.org/10.1093/oso/9780199258147.001.0001

Zimmer, Karl E. 1970. On the evaluation of alternative phonological descriptions. *Journal of Linguistics* 6(1). 89–98. DOI: https://doi.org/10.1017/S0022226700002383

Zsiga, Elizabeth C. 2012. *The sounds of language: An introduction to phonetics and phonology*. John Wiley & Sons. DOI: https://doi.org/10.1002/9781394260980