

---

## *Conquer Primal Fear: Phonological Features are Innate and Substance-Free*

---

### 1. INTRODUCTION

When a honeybee finds food, it returns to the hive to perform the famous waggle dance that conveys to hivemates the direction and distance of the food source.<sup>1</sup> The direction is the angle from the hive with respect to the sun at the time of dancing (not at the time the food was found), and distance is approximately reflected in the duration of each waggle interval, also taking into account the required effort due to wind speed and other factors. Somehow, the bee’s observations of the location of food are encoded as information in the bee’s nervous system. The angle between the bee’s figure-eight loop and the down direction provided by gravity matches the angle between the direction of the food source and the direction of the sun, which is how the bees know where they have to go. The number of waggles corresponds to distance. The hivemates that observe the dance end up with (approximately) the same knowledge about where the food is as the reporting scout. In both cases, the information is encoded in the neurology of an individual bee-brain as a symbolic representation that constitutes that bee’s knowledge of the location of the food source. Are those symbolic representations substantive? Are they made of sun, or wind, or angles, or waggles?

---

<sup>1</sup>This paragraph draws on Gallistel (2000) as well as the documentary *The Waggle Dance of the Honeybee* at <https://youtu.be/bFDGPgXtK-U> by the Georgia Tech College of Computing.

These questions are misguided. The symbols by which the bees store knowledge are transduced from sensory experience (such as observations of sun, wind and flowers) into the nervous system of the scout; then they are transduced out of the scout in the form of directed waggles by what Gallistel calls one of the bee’s ‘behavioral read-out systems’; then they are transduced into the brains of other bees by some kind of write-in system; and then those bees can use that information to find the food. The symbols themselves do not have nectarous, visual, geometric, waggly or any other kind of property of the world outside of brain neurology.

Phonological symbols work the same way: they are neurally encoded, and they are neither acoustic nor articulatory. Our main goal in this paper is to sketch a neurobiologically grounded view of the symbolic primes of phonological representations, features (Section 4). Before doing so, we discuss the nature of phonological computation (Section 2) and we argue that the innateness of features remains the best hypothesis for their origin in individuals (Section 3). We offer these concrete arguments about various aspects of phonological *primes* in the hope of mitigating an apparent pandemic of *primal* fear in the field—a resistance to positing innate primes born of an unjustified and tenacious empiricist bias. In brief, we present a model of phonological competence in which both the computational part and the representational part are substance-free, and we argue that the primes of phonological representations, features, are innate brain-based symbols related to—but distinct from—phonetic substance.

## 2. PHONOLOGICAL COMPUTATION IS SUBSTANCE FREE

Using an analogy from visual perception, we attempt here to clarify what was originally meant by [Hale and Reiss \(2000, 2008\)](#) in using the term “substance-free” in the context of phonological computation.

Consider the two sides of Figure 1. It seems fair to describe the left side as representing a dark gray rectangular object partially occluding a light gray one. The right side appears to depict a light gray rectangular object partially occluding a dark gray one.



**Figure 1:** Occlusion with dark and light gray regions.

These descriptions and our ability to see these scenes obviously rely on differences of ‘visual substance’. For each side of Figure 1, your mind ‘sees’ that there are three regions, and that the two non-contiguous regions have the same hue/texture/texture. Your mind also ‘infers’ that the same-looking regions are parts of the same object, occluded by the other object. This process of inference is substance-free, in that the merging of two regions of appearance  $x$  into a single object partially occluded by an object of appearance  $y$  is not dependent on the properties of  $x$  and  $y$ , as long as they are distinct to some degree.<sup>2</sup> As the two sides in Fig 1 show,  $x$  and  $y$  can be switched without affecting the inference. ‘Light gray occludes dark gray’ involves the same abstract inference as ‘Dark gray occludes light gray.’ The interpretation of object constituency is substance-free in the same sense that we intend the notion to apply


---

<sup>2</sup>There are complications of course, such as ones involving apparent light sources, as in the checker shadow illusion: [https://en.wikipedia.org/wiki/Checker\\_shadow\\_illusion](https://en.wikipedia.org/wiki/Checker_shadow_illusion).

4 CONQUER PRIMAL FEAR

in phonology. Deletion of +ROUND before –HIGH is no different computationally from the deletion of –HIGH before +ROUND.

We are not claiming that absolute and relative properties of  $x$  and  $y$  can never bias certain interpretations of foreground vs. background. However, it is clear that there are infinitely many substance-y ways in which  $x$  and  $y$  can differ (e.g. in terms of color, texture and patterns), so the operation used by the mind to decide that one object is occluding another must be substance-free. Furthermore, it is not necessary that each region has a constant ‘look’—the two regions of an occluded object can be colored with a gradient surface, for example. Here’s a couple of random examples:



**Figure 2:** Occlusion with more complex surfaces

Obviously, such examples can be multiplied without bounds.

Optimality Theory (OT) (e.g. [Prince and Smolensky, 1993](#); [McCarthy and Prince, 1993](#)) is substance-free in that the EVAL function counts markedness violations the same way irrespective of whether they occur in a constraint referring to +NASAL or –LABIAL. However, OT, in the ‘classic’ version and in many commonly adopted later versions, is not substance-free in that it encodes putative markedness in CON, the universal constraint set. There are constraints against front rounded vowels and against voiced obstruents in codas, but there are no constraints against back rounded vowels or voiceless obstruents in codas. Substance-free phonology posits that a neuroscientist of the future could program a human to have a grammar with coda *voicing*, despite the phonetic unlikelihood of such a system arising naturally. This would not involve changing the built-in Universal Grammar (UG). In contrast, a classic OT

phonology with universal innate constraints does not allow the possibility of coda voicing unless the neuroscientist is able to rewrite CON, a posited component of UG itself.

Let’s suppose for the moment that features have a ‘semantics’ via transduction to and from acoustics and articulation (see Section 4 for concrete discussion). Even so, phonological computation remains substance-free in the same sense that formal symbol manipulation is typically understood:

- (1) Brian Cantwell Smith (1996), *On the Origin of Objects*

Because formal symbol manipulation is usually defined as ‘manipulation of symbols independent of their interpretation,’ some people believe that the formal symbol manipulation construal of computation does not rest on a theory of semantics. But that is simply an elementary, though apparently very common, conceptual mistake. Symbols must have a semantics—i.e., have an actual interpretation, be interpretable, whatever—in order for there to be something substantive for their formal manipulation to proceed independently of.

In other words, we claim that phonology is substance-free in the same sense that a rule of inference like *modus ponens* is substance free:

- (2) If  $P$  then  $Q$   
 $\frac{P}{\therefore Q}$

This rule of inference applies regardless of whether the contents of  $P$  and  $Q$  refer to mortals, dogs, quarks, coronals or deadly sins.

Now, a given language may have a rule that involves deleting the valued feature +VOICED in syllable codas; and another language may have a rule deleting

+LATERAL from a liquid after a front vowel; but the mechanism of deletion will presumably be the same in each case, for example via set subtraction (Bale et al., 2014; Bale and Reiss, 2018). The process of phonological acquisition involves determining what valued features (from an innately determined set; see Section 3) are present in the rules of a language, but the kinds of operations available for rule construction rely in no way upon which particular features occur in a rule. UG does not determine that some operations are for  $\pm$ VOICED and others for  $\pm$ LATERAL. That’s what we mean when we say that phonological computation is substance free.

### 3. FEATURES ARE INNATE

Jerry Fodor (1980) famously demonstrated that there are no serious theories of learning from scratch. Learning is always a form of hypothesis confirmation that depends on some pre-existing, built-in primes, which cannot themselves be learned. Echoing Fodor, Jackendoff (1990, p. 40) says that “in any computational theory, learning can consist only of creating novel combination of primitives already innately available.” Fodor’s argument is often ignored, but has never been refuted. Even among scholars familiar with nativist ideas, in discussions of the source of cognitive structures, there appears to be a bias favoring learning from a blank starting state over nativist explanation.<sup>3</sup> Our position, in contrast, is that some version of nativism should generally be taken as the null hypothesis in light of Fodor’s logical demonstration. Given this perspective, we will not offer a full rebuttal to every claim put forth *against* the innateness of a set of specific (but still to be determined) universal phonological features, but rather poke, in a “brief, cavalier and dogmatic” (Tolman, 1948, 207) fashion, at some of the ideas circulating in the anti-nativist camps.

---

<sup>3</sup>See our paper, Volenec and Reiss (2020), for examples of denial of the relevance of the Argument from the Poverty of the Stimulus by phonologists.

### 3.1 A little UG goes a long way

One recurrent objection to positing innate features is that we’d need too many of them—it is implausible that UG should contain so much information. Such claims tend to be vague about how much is too much, and typically don’t even discuss basic notions like orders of magnitude in relating the genome to information structures.<sup>4</sup>

Phonologists tend to work with 15-30 features, and in this context positing a thousand features seems radical, but merely raising the number to, say, 50 is trivial in terms of the amount of encoded information involved (since adding one bit of information doubles the number of distinctions that can be encoded). As we’ll see when we do the basic math, the consequences of such a minor change are tremendous.

If Universal Grammar (UG) provides a set of 20 binary features as the building blocks of phonological segments, then UG defines *intensionally* a very large set of possible segments.<sup>5</sup> If each feature must be present in each segment with some value (e.g. +ROUND or –ROUND), the number of segments is about a million:  $2^{20} = 1,048,576$ . If underspecification is allowed, that is, if segments can be

---

<sup>4</sup>We are outside of our intellectual comfort zone here, but it is worth pointing out that the human genome (the totality of our DNA) consists of six billion individual nucleotide bases (<https://www.genome.gov/human-genome-project/Completion-FAQ>). So having a few base pairs which would code for the proteins that build the brain tissue whose activity corresponds to features doesn’t seem unreasonable (especially if that same brain tissue is involved in other functions). The burden of proof for the claim that that is too much innate information is on those suffering from primal fear.

<sup>5</sup>This, of course, does not mean that children are born with lots of segments encoded in their brains. It means that they have the capacity to build a lexicon containing morphemes built from segments, each of which is one of the many, many possible combinations of features. UG intensionally defines this (finite, but large) set of segments in the same sense that an English-type grammar intensionally defines an infinite set of sentences.

‘incomplete’, lacking a value for some of the 20 features, then the number of segments rises to the billions:  $3^{20} = 3,486,784,401$ . In light of these basic calculations (Reiss, 2012; Matamoros and Reiss, 2016; Bale and Reiss, 2018), the null hypothesis for phonological acquisition should be that a small amount of universal, innate knowledge is sufficient to account for the potential diversity of phonological inventories. This conclusion is at odds with the contrary suggestion that to learn any possible language to which they were exposed, learners would “need a (very large) *a priori* set of possible features to choose from” (Cowper and Hall, 2014). In fact, with 20 features, the number of segment inventories that UG intensionally defines is  $2^{1,048,576}$  (which Google calculator returns as “infinity”). You get this number by considering how many ways there are of answering “yes” or “no” about whether each segment is found in a particular language—that’s just the size of the power set (the set of all subsets) of the set of all possible segments. The number of particles in the universe is, by some estimates, only around the order of  $2^{285}$  (around  $10^{85}$ ). So, our simple calculations show that claims like that of Cowper and Hall (2014) are factually incorrect. The burden of proof is on those who deny universal, innate features to show that a fairly small set is *not* sufficient to account for linguistic diversity.

This perspective on the combinatorics of innate primitives of phonology is consistent with Gallistel and King’s discussion of how combinatoric explosion solves the problem of the power of the representational capacity of biological systems:

[T]he number of symbols that might have to be realized in a representational system [like the brain-cr&vv] with any real power is for all practical purposes infinite; it vastly exceeds the number of elementary particles in the universe, which is roughly  $10^{85}$  (or  $2^{285}$ ), give or take a few orders of magnitude. This is an example of the difference between the infinitude of the possible and the finitude of the actual, a distinction of enormous importance in the design of practical computing machines. A computing machine



can only have a finite number of actual symbols in it, but it must be so constructed that the set of possible symbols from which those actual symbols come is *essentially infinite*. (By ‘essentially infinite’ we will always mean greater than the number of elementary particles in the universe; in other words, not physically realizable.) This means that the machine cannot come with all of the symbols it will ever need already formed. It must be able to construct them as it needs them—as it encounters new referents.

[Gallistel and King 2009, p.74]

So on the one hand, we must have some innate primes on the basis of which more complex data structures like segments are constructed during the growth of language in the brain, while on the other hand the number of these primes can be fairly small due to combinatoric explosion. Our perspective on combinatorics is thus consistent with Chomsky’s (2007) point that “the less attributed to genetic information (in our case, the topic of UG) for determining the development of an organism, the more feasible the study of its evolution”. We don’t need a lot of UG to get a lot of empirical coverage.

We have made this combinatoric argument repeatedly, yet, apparently not convincingly, so let’s bring the discussion back down to something more concrete. With just the five most common features used to discuss vowel systems—HIGH, LOW, BACK, ROUND, ATR—we can define  $3^5 = 243$  different vowels (well, 216 if you insist on ruling out any superset of  $\{+HIGH, +LOW\}$ ).<sup>6</sup> If you imagine a five-vowel system as containing one member from each of five regions of the vowel space, then with a single partitioning of the space, there are over 40 choices for each region, each of the

---

<sup>6</sup>There are nine logical possibilities for these two features, given underspecification: all combinations of ‘+’, ‘-’ and ‘absence’ for the two features. Excluding co-occurrence of +HIGH and +LOW leaves eight possibilities. Multiply eight by  $3^3 = 27$  possibilities for BACK, ROUND and ATR to get 216.

five vowels. Let’s say we cut that in half under the assumption that the vowels in a five-vowel system are not going to be too close to each other—no two vowels will lie near a common region boundary. This leaves us with 20 vowels per region. What this means is that what we phonologists tend to reify as ‘the common five-vowel system’ is really a set of  $20^5 = 3.2$  million vowel systems.

If the set of ‘basic’ vowel features turns out to be a bit larger than 5, say, 10, then the number of definable vowels in each region rises to  $3^{10}$ , which is about sixty-thousand. Making the same cuts as above (ruling out co-occurrence of +HIGH and +LOW, and then cutting the remainder in half), we still get about twenty-six thousand vowels, so over five thousand per region. So then we get an even larger typology of five vowel systems, namely  $5000^5$  which is more than  $3 \times 10^{18}$  distinct vowel systems.

This discussion has been concerned with the variety of vowels attributable to featural differences, but there are other factors relevant to observed differences in outputs between languages or even closely related dialects. Such factors again do not require that each variety have a language-specific phonetic module. There can be interactions among the transduction of featural representations and the transduction of other properties such as pitch or stress. If, descriptively, you find languages with two consistently different [i]’s, it may be that they are featurally different, but it can also be that they are the output of I-languages with, say, different pitch accent systems. The interaction of different pitch accent systems with the exact same feature combination will give rise to different ”pronunciations of [i]”. If we have these several kinds of explanation for variation in, say, vowel pronunciation across languages, why would we want to posit language specific *phonetic* learning as well (see [Hale et al. 2007](#))?

### 3.2 Claims about ‘emergence’

We will briefly address two claims in the literature concerning putative ‘emergence’<sup>7</sup> of features. [Mielke \(2008\)](#) argues that features are not innate and thus must be learned in the course of language growth. The argument is based on the claim that there are many rules that do not refer to (phonetically) natural classes of segments. The discussion has several problems. First, this claim cannot be evaluated since Mielke does not provide a theory of what counts as a phonological rule. As we explain in detail in [Volenc and Reiss \(2020\)](#), positing/determining rules on an intuitive basis is scientifically inadequate. Relatedly, without natural classes of environments, defined by features, it is unclear how a learner could even detect the patterns that supposedly instantiate a putative rule. Finally, as we discuss in [Volenc and Reiss \(2020\)](#), even in a model in which rules are *defined* as referring to natural classes intensionally, the extensionally defined set of forms to which a rule applies may not constitute a natural class, because of bleeding rule interactions.

Another kind of argument against nativist feature theory is that the substantive correlates of (some) features reflect perceptual categories shared with other species. What such discussions fail to point out is that it is only a subset of such shared

---

<sup>7</sup>‘Emergence’ is a poor choice of terms for the assumption that features are not innate, but rather learned on the basis of exposure to external stimuli. It stems from a tacit assumption that ‘innate’ means ‘present at birth’, which is clearly incorrect. In humans, teeth are innate, yet they are not present at birth, they ‘emerge’ several months after birth; puberty is innate, but not present at birth, it ‘emerges’ later in life, and so on. So something that is innate can also ‘emerge’ during the growth of an organism, which is clearly not what the so-called ‘emergent’ theories of features are supposed to convey.

biologically determined perceptual categories that get used in human phonology.<sup>8</sup> We have a name for that subset: *phonological features*. Thus, features perhaps did emerge from acoustic or perceptual categories in the course of evolution, but that is actually an argument *against* ‘emergence’ in each individual human’s development. Evolution involves changes to the genome.

Our view is not original: [Chomsky and Halle \(1965\)](#) concluded in their response to [Householder \(1965\)](#), who claimed that features ought to be posited on a language-specific basis, that “all phonology breaks down if we do not assume analysis [...] in terms of universal [...] features” (p. 119) and that “the assumption of a universal feature structure is made (often only implicitly) in every approach to phonology that is known, and clearly cannot be avoided” (p. 127).

### 3.3 Be careful what you wish for

Before presenting our view of the neurobiological basis of features we want to point out some consequences that follow if we turn out to be wrong. If you are dead set on rejecting our nativist null hypothesis, you’d better be ready to face the consequences.

#### 3.3.1 The scholar’s complaint

First, if there are no innate features, then a tremendous portion of the phonology literature is completely worthless, because there can be no such thing as, say, ATR harmony or ROUND harmony. This is because without innate features there will be ‘what we’ll call ATR in Bongo’, and there will be ‘what we’ll call ATR in Dagik’ and there will be ‘what we’ll call ATR in Yoruba’, and there will be ‘what we’ll call ATR in Québec French’, but there won’t be a coherent general notion of ‘ATR harmony’

---

<sup>8</sup>And a tiny subset at that. For example, the human auditory system can discriminate between more than 1000 levels on a single pitch scale ([Fastl and Zwicker \(2006\)](#)), but we do not have 1000 different levels for vowel height (which correlates with F1).

because there will be no universal ATR feature as part of the innate human language faculty. All those papers on ATR harmony will be about completely different things, by definition.<sup>9</sup>

### 3.3.2 The learner’s complaint

Second, the idea that phonological feature space is just a learned structure imposed on the space of auditory perception potentially raises more problems than it solves. Under the innate features view, the child of a bilingual mother is constrained (i.e. guided by a hand of nature that cannot be disobeyed) to parse all input in terms of combinations of a small set of predetermined categories. However, as the existence of language-independent voiceprint technology demonstrates, there are purely acoustic properties of specific human voices that are detectable regardless of which language is being used. How does a child learn to ignore these properties that tell him that mommy is mommy, or that mommy is happy, and use the ones that allow him to discover French vs. Italian patterns? Pure acoustics provides lots of ways (maybe infinitely many) to categorize signals, most of them linguistically irrelevant, so without innate features, we need another way to explain how kids find the particular structure that works.

Given the astronomical (or maybe infinite) number of ways in which a raw acoustic signal can be categorized, it is miraculous that we get any cross-linguistic generalizations at all, even more so that they are so robust, and that there has never

---

<sup>9</sup>Aside from the fact that we *do* believe innate features, as non-substance abusers, we are not bothered by the fact that ATR harmony has no *formal* properties that set it aside from, say ROUND harmony. We are *not* claiming that any particular feature system in the literature is correct and that the literature is therefore *not* largely worthless—we are just saying that this problem becomes one of principle if there are no pre-determined features at all.

been a need to posit more than about thirty features to account for the attested patterns. All this seems much less miraculous if we assume that humans are innately equipped with a relatively small set of features.

### 3.3.3 The syntactician’s complaint

Third, *Why not?* Unless you have an alternate account, maybe involving “ocillators and energy levels” (Tilsen 2019), of why languages have things like negative polarity items and *wh*-movement, you probably believe in innate aspects of syntax and semantics. Isn’t it odd to accept that evolution could have managed to ‘nativize’ those really abstract categories (NPIs don’t all have the same formant transitions cross-linguistically), but reject the idea that the abstract (but less so) phonological features could also be innate? Or maybe you’re going to have to tell your syntactician friends that all *their* features, like *wh*-, really don’t exist either, and that cross-linguistic recurrent patterning is just a mirage.

### 3.3.4 More on our null hypothesis

Fourth, a further *Why not?* point relates to the pernicious empiricist prejudice noted above that positing that a cognitive capacity is learned from scratch is somehow ‘better’ than positing that it reflects some innate endowment: nativism is rejected as the null hypothesis. However, it behooves us to recall that there is very strong evidence for innate knowledge in many domains and many species. For example, Jon Rawski (p.c.) directs our attention to [Zador \(2019\)](#):

A striking example of a complex innate behavior in mammals is burrowing: Closely related species of deer mice differ dramatically in the burrows they build with respect to the length and complexity of the tunnels. These innate tendencies are independent of parenting: Mice of one species reared by foster mothers of the other species build burrows like those of their

biological parents. Thus, it appears that a large component of an animals behavioral repertoire is not the result of clever learning algorithms supervised or unsupervised but rather of behavior programs already present at birth.

Lorenz and Tinbergen’s classic example of innate knowledge is the reaction of newly hatched chicks of various bird species to a dummy figure moving over their cage. Moving in one direction, the figure appears to have a short neck and a long tail, like a hawk or some other bird of prey; moving in the other direction the figure appears to have a long neck and a short tail, more like a harmless goose or other non-prey species. The ‘hawk’ elicits a panicked reaction in the lab-hatched chicks, but the ‘goose’ does not (see [Schleidt et al. 2011](#) for discussion and references). Anti-nativists have to explain why they are ready to attribute highly specific innate structures to bird brains, but resistant to attributing innate structures like a phonological feature set to human brains.

There are countless more examples. Turtles orient themselves toward water immediately after hatching even though they’ve never seen water and water is not in their line of vision when they hatch; spiders raised in isolation weave a web just like spiders who had a chance to see a web being woven; ants display the path integration ability (dead reckoning) in absence of relevant experience; honeybees raised in isolation use waggle dance to signal the direction of a food source like other honeybees, and so on. Even when some learning is involved, it is necessary to posit an innate apparatus. There are species of ants and bees that navigate on the basis of the position of the sun in the sky, and part of this capacity appears to be learned to account for variations in latitude. If you ask a typical linguistics undergraduate (at least one who isn’t studying in Tromsø and who hasn’t been a member of the Nome (*sic*) Chomsky Fan Club) to point in the direction of the sun at midnight, they will typically stare at you in a slack-jawed muddle. However, those sun-navigating ants and bees, when

presented at midnight with a bright lamp, navigate as if the light is due north, which of course is the direction of the sun at midnight (in the Northern Hemisphere). The insects’ innate endowment ‘completes the circle’ of the sun’s movement from east in the morning to south at noon to west in the evening, allowing the bug to ‘know’ that the sun is to the north at midnight ([Jander \(1957\)](#)). Do we really want to deny innate endowment to humans, who have about 86 billion neurons ([Lent et al. 2012](#)), but grant it to ants with a mere two hundred fifty thousand?

### 3.3.5 The breather’s complaint

Fifth, [Lenneberg et al. \(1967, Chapter 3\)](#) points out that “it is quite clear that breathing undergoes peculiar changes during speech” due to “respiratory peculiarities that have evolved” which are specific to speech. Speech is speech, and not grunting or whistling, by virtue of being an externalization of a phonological (featural) representation. Feature nativists (like us) view the evolution of speech-specific respiratory mechanisms as dependent on the existence of speech, and therefore, on the existence of features. If one denies the existence of innate features, one needs to imagine that what ends up as a speech-specific respiratory mechanism in each individual could have evolved without the existence of a set of features present in each individual. It seems a bit odd to think that children are not innately endowed with features, upon which speech is predicated, but have an innate evolved system of respiration for speech. Implausibly, one has to posit a innate respiratory system that has evolved to support a feature system that itself is not innate. It would be like positing an innate evolved mechanism that is only manifested for the learned process of moving chess pieces.

## 3.4 Reiteration

Never in the history of philosophy or science has it been demonstrated that a true blank slate can learn. Even when a prominent paper like [Silver et al. \(2017\)](#) makes



PHONOLOGICAL FEATURES

explicit reference to a “tabula rasa”, the ‘tabula’ is anything but ‘rasa’—it includes loads of built-in assumptions and preconfigured learning machinery. Where does that machinery come from? It has not been learned, but has been built into the system by human engineers. True blank slates, like rocks, don’t learn; that which can learn has some kind of a learning mechanism which itself hasn’t been learned; therefore the mechanism is bestowed by biology (or by an intelligent programmer)—it’s innate.

Innateness is omnipresent in the biological world—unless you believe that genetics doesn’t exist. Since brain development is guided by genetics and since the growth of the language faculty displays all the relevant traits of other genetically-driven developmental processes like puberty and menopause, we should expect that phonology has an innate core just as much as we think that puberty has an innate core. Clearly, the superficial aspects of both puberty and phonological development can be determined by external factors (in phonology, this amounts to ‘learning’), but the core general properties of both of these processes are given by biology, not learned by observation. In the case of phonology, they cannot be learned purely by observation because the sorting of experience (which, objectively, is a welter of undifferentiated stimuli) into categories implies the prior existence of at least some categories. Otherwise, either the process of categorization could not even start or each individual would end up with a completely different and incommensurable categorization based on random factors such as moods and recent experiences; this does not happen with phonological features. Therefore, the null hypothesis should be that all *elementary* units of language—those irreducible units from which all other, more complex representations are built—are given *a priori*, i.e., provided by genetics and

that I-languages are constructed through the interaction of those innate primes with environmental factors.<sup>10</sup>

#### 4. THE NATURE OF PHONOLOGICAL FEATURES AND THEIR RELATION TO PHONETIC SUBSTANCE

Phonological computations operate over atomic units out of which phonological representations are built. At the lowest level of granularity these units are phonological features. In this section we clarify how features are conceptualized in Substance Free Phonology, focusing primarily on their ontology and their relationship to phonetic substance.

The first and most basic question that can be asked about features is What are they? Following a long tradition in generative phonology, we take features to be abstract mental units that have a lawful but highly indirect relation to phonetic substance.

Such considerations [that languages do not make free use of acoustic values or articulatory properties] were much in our minds [...] when Jakobson, Fant and I were working on *Preliminaries to Speech Analysis*, and it was these considerations that led us to draw a sharp distinction between distinctive features, which were abstract phonological entities, and their concrete articulatory and acoustic implementation. Thus, in *Preliminaries* we spoke not of ‘articulatory features or of ‘acoustic features, but of ‘articulatory and/or ‘acoustic correlates of particular distinctive features. [Halle, 1983/2002, 108-9]

---

<sup>10</sup>Again, it is conceivable that some innate features are not necessarily present in a fully mature form at birth, but the results of speech perception studies on infants like those discussed by Werker (1995) are not particularly supportive of that possibility.

PHONOLOGICAL FEATURES

The relation between a phonemic system [which is built out of bundles of features cr & vv] and the phonetic record [...] is remote and complex .  
[Chomsky, 1964]

As phonology in the 20th century progressed from the taxonomic and mostly anti-mentalistic structuralism to the cognitively and neurobiologically oriented generative perspective, attempts have been made to connect feature theory to human neural structures.

In articulatory terms each feature might be viewed as information the brain sends to the vocal apparatus to perform whatever operations are involved in the production of the sound, while acoustically a feature may be viewed as the information the brain looks for in the sound wave to identify a particular segment as an instance of a particular sound. [Kenstowicz and Kisseberth, 1979, 239]


[F]eatures correspond to controls in the central nervous system which are connected in specific ways to the human motor and auditory systems. In speech perception detectors sensitive to the properties [...] are activated, and appropriate information is provided to centers corresponding to the distinctive feature[s] [...]. This information is forwarded to higher centers in the nervous system where identification of the utterance takes place. In producing speech, instructions are sent from higher centers in the nervous system to the different feature[s] [...] about the utterance to be produced. The features then activate muscles that produce the states and configurations of different articulators [...]. [Halle, 1983/2002, 109-10]

Continuing this tradition and sharpening its main claim about the neural reality of features, we adopt the cognitive-neuroscience framework proposed by Gallistel

and King (2009). In this framework, the atomic elements of mental representations are called *symbols*. Symbols are “physical entities in a physically realized representational system” (Gallistel and King, 2009, 72), where the physical system in the case of phonological symbols, and all other cognitive symbols, is the human brain. Thus, phonological features are symbols realized in the human brain. The common properties of all neural symbols are (at least) *distinguishability*, *combinability* and *efficacy*.

The standard assumption in cognitive neuroscience is that different symbols are *distinguished* by place coding of neural activity, rate coding, time coding, or, most likely, some combination of those. Of course, we are still far from being able to state precisely how exactly features qua neural symbols are realized in the brain, but experimental studies are consistently emphasizing the importance of neural activity in the superior-most part of the superior temporal gyrus (STG), superior temporal sulcus (STS), and Brodmann (BA) areas 44 and 6 (Figure 3).

The representations of articulatory correlates of features are coded in the posterior inferior frontal gyrus of the left hemisphere, traditionally known as Brocas area (Okada et al. (2018)). More specifically, Hickok (2012, 138) reports that pars opercularis (BA 44) and the ventral-most part of BA 6 store articulatory programs needed to reach the auditory targets imposed by features (roughly corresponding to the blue circle in Figure 3). These auditory targets, i.e., the representations of auditory correlates of features, are coded in the STG and the STS (roughly corresponding to the red circle in Figure 3). Mesgarani et al. (2014) showed that acoustic phonetic information is represented in the STS and is distributed along five distinct areas, each corresponding to a general manner of articulation parameter. By measuring the responses in implanted electrical cortical grids placed along the superior-most part of the temporal gyrus, they found that one electrode responded selectively to stops, one to sibilant fricatives, one to low back vowels, one to high front vowels



**Figure 3:** Regions of the human brain (left hemisphere) assumed relevant for the coding of phonological features.



and a palatal glide, and one to nasals. Similarly, [Bouchard et al. \(2013\)](#) constructed an auditory-based place of articulation cortical map in the STG, confirming labial, coronal and dorsal place features with different electrodes, and cutting across various manner classifications. Using magnetoencephalography, [Scharinger et al. \(2012\)](#) localized three vowel features—height, frontness and roundness—in different parts of the STG.

Features qua neural symbols also meet the criterion of *combinability*. A hallmark of modern phonology is the notion that features can be grouped into sets to construct higher-level, non-atomic data structures. An unordered, unstructured set of features constitutes a phonological *segment* (see [Volenc and Reiss, 2020](#), §4), while

a particular organization of segments constitutes a data structure of the next higher taxonomic level, namely a *syllable*. This combinability of features allows phonology to construct complex symbols from an inventory of simple parts, and provides an explanation for the so-called *natural class behavior*—different structures can behave alike because they contain identical substructures.

As we saw above, features are also an *efficacious* way of coding information since their combining leads to combinatoric explosion. For example, if we assume that the brain stores and uses only 30 binary features with the possibility of underspecification, then from this small set of primitive symbols we can construct  $3^{30}$  or about 206 trillion different segments. Of course, the richness that arises from feature combinability should not be taken to imply that any particular I-language should come close to exploiting the full range of possibilities. Instead, what we expect to find in particular I-languages is in line with the traditional view of feature combination: “no language has as many phonemes as there are possible combinations of the utilized distinctive features [Halle \(1954\)](#). A corollary of this combinatoric explosion is that such richness goes a long way towards eliminating the need for a phonetics module specific to each language (as in [Keating, 1985, 1990](#)), which simplifies the sequence of conceptual steps needed to account for the externalization of language (see [Volenec and Reiss, 2020](#), for elaboration).



Recent neuropsychological studies have shed light on some other aspects of features that are significant for phonological theory, namely their discreteness, binarity and potential underspecification. By eliciting magnetic mismatch fields in an oddball paradigm, [Phillips et al. \(2000\)](#) have shown that the left hemisphere auditory cortex has access to representations of discrete and binary phonological categories. In other words, their study has “demonstrate[d] the all-or-nothing property of phonological category membership (p. 1038), where this category membership is determined on the basis of phonological features and not on the basis of general categorical auditory



PHONOLOGICAL FEATURES

perception. This finding contradicts the possibility of gradient articulatory gestures serving as basic units in phonological computation (as in [Browman and Goldstein, 1989](#)). Furthermore, [Scharinger et al. \(2016\)](#) found that a less specified vowel compared to its more specified counterpart resulted in stronger activation in the left STS, thus providing some insight into the neural underpinnings of phonological underspecification.

Even though some progress has been made in discovering the neural reality of features, we are still far from being able to refer to features by stating their neurobiological substrate and therefore have to resort to using symbols (labels) to refer to symbols. So when we write, for example, LABIAL, we use a sequence of letters to form a symbol for a particular feature, which in turn is also a symbol, just in the brain. In other words, LABIAL is a non-neural symbol for a neural symbol. We, the researchers, need these phonetic labels to know what we are talking about, the brain does not. The brain does not need such phonetic labels because the transduction algorithms at the phonology-phonetics interface (see below) interpret the identity of a feature by the place of the neural activity in the brain and its temporal properties ([Khalighinejad et al., 2017](#)). This is similar to how a computer does not retrieve the identity of a symbol solely on the basis of its form (1s and 0s), but rather by combining the information about the form with the location and context in the memory ([Gallistel and King, 2009](#), 73). Possibly, the actual form of all features is the same neural spike (i.e., an action potential). But more importantly, the unique location of the spike and the rate of its repetition is how the transducer determines the identity of the feature and knows which neuromuscular schema (e.g., labiality and not, say, nasality) to assign to it. It can of course be debated whether it is misleading or not to use phonetic labels such as LABIAL to refer to features qua neural symbols and whether there is a better solution to this. But a decision on this issue has no bearing



on the actual nature of features: the neural symbol is, of course, the same irrespective of whether we refer to it as LABIAL or by using a non-phonetic label.

The symbolic nature of something is that it stands for something else, something that is not the same as the symbol. That for which a symbol stands, that which it represents, is variably called its *referent* or *correlate*, or *the represented*. Phonological features are symbols that refer to aspects of speech. At this point, it is of utmost importance not to “make the common mistake of confusing the symbol with what it represents” (Gallistel and King, 2009, 56); “the tendency to confuse symbols with the things they refer to is so pervasive that it must be continually cautioned against (p. 62). There is a connection between phonological features and speech, but this connection is highly complex and indirect (see below), and features do not encode speech-related information in any straightforward way. In linguistics, information related to speech is called *phonetic substance*. It is the totality of the articulatory, acoustic and auditory properties and processes that constitute speech. For example, properties and processes of speech such as movements of the tongue, values of formants, loudness, duration expressed in milliseconds etc. fall under the rubric of substance. Since features are symbols physically realized in the brain, they cannot contain phonetic substance. In other words, features are substance-free. Believing that features are substance or that they contain substance is just an instance of the aforementioned mistake of confusing the symbol with what it represents.<sup>11</sup>

It should also be noted that there is ample experimental evidence for the substance-free nature of features. Phillips et al. (2000, 1040) have concluded that “there is good

---

<sup>11</sup>In some earlier work, we referred to features as *substantive* to indicate the regularity of the transduction processes—the feature we call +HIGH cannot be transduced in into, say, spreading of the vocal folds, which is correlated with a feature we typically call –VOICE. The use of ‘substantive’ here leads to confusion and thus turns out to have been a bad choice of terminology.



reason to distinguish the acoustic and phonetic representations that underlie categorical perception from the discrete phonological category representations involved in lexical storage and phonological computation, and that when it comes to phonological computation, “all within-category contrasts are lost: e.g., all different tokens of /d/ are treated by phonological processes as exactly the same, irrespective of the phonetic substance that is indirectly associated with the bundle of features that we conventionally label as /d/. [Magrassi et al. \(2015, 1\)](#) have shown that “the activity of language areas is organized in terms of features even “when language is generated mentally before any utterance is produced or heard, i.e., when there is no phonetic substance whatsoever. Similarly, [Okada et al. \(2018\)](#) have conducted an fMRI investigation of silent word sequence production (i.e., the subjects read words in their minds) where the stimuli (different words displayed one after another on a screen) varied in the degree of feature overlap in consonant onset position. The experiment confirmed a featural organization of investigated word sequences in absence of overt speech. These neuropsychological studies suggest that phonological features cannot be equated with the phonetic correlates that are typically associated with them, i.e., that features and phonetic substance are two different things.

The results of these studies should not surprise us, as they merely reflect a more general principle in cognitive neuroscience. In cognitive domains that are unrelated to language, a decoupling of a mental/brain representation from the stimuli used to elicit it has frequently been demonstrated. In a series of experiments ([Quiroga et al. 2005](#); [Cerf et al. 2010](#)), [Itzhak Fried](#) and colleagues have shown that an individual neuron in the medial temporal lobe is consistently selectively activated when—to take one specific example—a mental representation of the actress Halle Berry is invoked. Crucially, the same ‘Halle Berry neuron’ was activated in the following disparate situations: when Berry was shown in a photograph, when a video of her was played, when she was seen dressed as Catwoman, when she was seen without



26

CONQUER PRIMAL FEAR

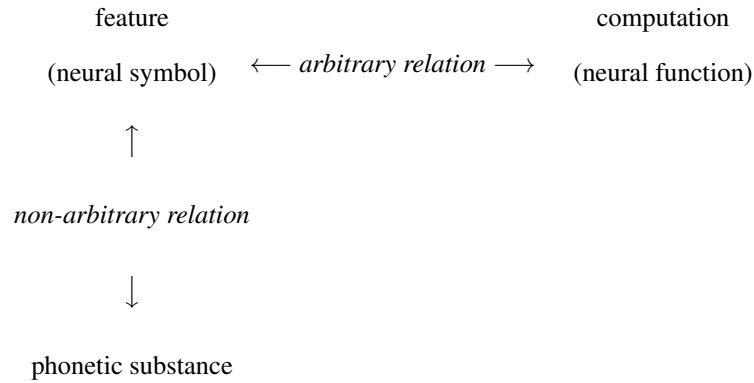
the costume in various instances (with various outfits, hairstyles and accessories), when only her name was written on a screen, when her name was said out loud, and when the participants imagined her. Such a neuron corresponds to a mental representation that is clearly not tied to any particular sensory modality, and whose relation to substance/stimulus is highly indirect. We assume that the same principle holds for phonological primes.

We do not think that the relation between any given feature and its correlate is random or arbitrary. If it were, then any feature could in principle be realized by any possible human articulation, just as the concept/signified DOG can in principle be assigned to any possible signifier. If such Saussurean arbitrariness were applicable to the realization of features, then it would be possible that +ROUND sometimes gets realized as a lowering of the velum, sometimes as a raising of the tongue dorsum, and so on. But there seems to be no convincing evidence for anything like that, so the simplest assumption is that it is not the case. Instead, we assume that there is a non-arbitrary, lawful relation between features and their correlates. The nature of this relation is described in more detail below, but at this point it is important to emphasize that the lawful relation between features and their correlates is *phonologically* irrelevant, as weve stated in Section 2. That means that phonological computation treats features as invariant categories, manipulating them in an arbitrary way irrespective of the variability in their realization in speech and irrespective of phonetic substance in general. The relation between features and phonological computation is thus completely independent from the relation between features and phonetic substance, as can be seen in the diagram in (3).

(3) Phonology (horizontal) and the phonology-phonetics interface (vertical)



PHONOLOGICAL FEATURES





The diagram in (3) also reflects the distinction between computation and transduction, which we will make use of below. Computation—corresponding to the horizontal relationship in (3)—is the formal manipulation (reordering, regrouping, deletion, addition etc.) of representational elements within a single cognitive system, and without a change in the representational alphabet characteristic for that system. Transduction corresponding to the vertical relationship in (2) is a process of transmuting an element in one form into a distinct form, that is, a mapping between dissimilar formats.

As substance-free symbols, features do not contain information on the temporal coordination of muscle contractions, on the spectral configuration of the acoustic target to be reached, and so on. Yet without this information, the respiratory, phonatory and articulatory systems cannot produce speech. The sensorimotor (SM) system<sup>12</sup> which is in charge of speech production requires information about substance and time in order to arrange the articulatory score (Guenther 2016), therefore this information has to be integrated into a representation before being fed to the SM system.

---

<sup>12</sup>Since the term sensorimotor system is used ambiguously in the linguistic literature, it should be noted that we use that term to denote a large-scale brain network that includes the pre-central and post-central gyrus and the supplementary motor area.



We thus posit a transduction component that connects phonological competence with the vastly different SM system. Our theory of that component and the component itself are called *Cognitive Phonetics* (CP; [Volenec and Reiss 2018](#)).

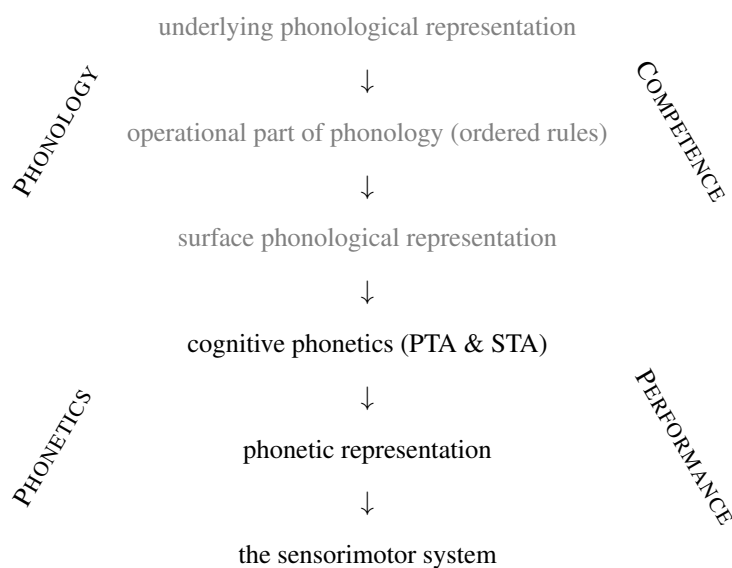
CP proposes that the phonology-phonetics interface consists of at least two transduction procedures that convert the substance-free output of phonology into a representational format that contains substantive information required by the SM system in order to externalize language through speech. The inputs to CP are the outputs of phonology, i.e., surface phonological representations (SRs). SRs are strings of segments, each of which is a set of features. Each feature of SRs is transduced and subsequently receives interpretation by the SM system (cf. [Lenneberg et al., 1967](#), §3). This transduction is carried out by two algorithms (cf. [Marr, 1982](#), 23-24). The *paradigmatic transduction algorithm* (PTA) takes a feature (a symbol in the brain) and relates it to a motor program which specifies the muscles that need to be contracted in order to produce an appropriate acoustic effect. The *syntagmatic transduction algorithm* (STA) determines the temporal organization of the neuromuscular activity specified by the PTA. In simpler terms, PTA assigns muscle activity to each feature, STA distributes that activity temporally. These transduction algorithms yield an output representation of CP, which then feeds the SM system. The output of CP is called the *phonetic representation* (PR),<sup>13</sup> and it can be defined as a complex array of temporally coordinated neuromuscular commands that activate muscles involved in speech production.

---

<sup>13</sup>In generative linguistics literature, the output of phonological grammar, what we call a surface phonological representation, is sometimes called a phonetic representation (cf. [Hall 2001: 29](#)). Note that we use the term very differentlyour phonetic representations are not part of grammar and do not contain the phonological features that are in URs and SRs.

The standard schema of phonological competence can now be expanded to accommodate the transduction performed by CP into the more complete ‘speech chain’ in (4):

(4) Phonology and Phonetics



The gray parts of the schema represent phonological competence, while the black parts correspond to the initial *phonetic* steps in speech production. That is, the difference in shading parallels the competence/performance dichotomy: phonology is competence, cognitive phonetics is (one component of) performance.

To clarify the effects of PTA and STA, we can explore in some detail the transduction of a few hypothetical SRs (see Volenec and Reiss (2019) for further examples). We will see that PTA and STA have considerable implications: they open the possibility of elegantly accounting for subtle yet systematic interactions of two kinds of coarticulatory effects, which is only possible if we assume that the basic units of speech production are indeed transduced phonological features. Suppose that a hypothetical I-language contains SRs [lok] and [luk]. Each segment is a set of features,

and vowels [o] and [u] both contain the valued feature +ROUND, on which we will focus. One thing that should be noticed is that [o] and [u] are different in terms of height: [o] is –HIGH, [u] is +HIGH. The PTA takes a segment, scans its feature composition and determines the required muscular activity for the realization of every feature. Roughly, for +ROUND the PTA activates at least four muscles—*orbicularis oris*, *buccinator*, *mentalis*, and *levator labii superioris* (Seikel et al., 2009)—which leads to lip rounding. The difference in PTAs effect on –HIGH and +HIGH is that for the latter the algorithm raises the tongue body and the jaw, while it does not for the former. While transducing +ROUND, the PTA takes into account the specification for HIGH and assigns a slightly different lip rounding configuration for [o] than for [u]. Let us refer to a *transduced* feature, which we take to be the basic unit of speech production, as  $PR_F$ , where PR stands for phonetic representation and F stands for an individual (valued) feature. So,  $PR_{+ROUND}$  is the transduced feature +ROUND. We can now say that  $PR_{+ROUND}$  will be different for [o] because of its interaction with  $PR_{-HIGH}$  than for [u] because of its interaction with  $PR_{+HIGH}$ . Since these interactions involve transduced features within a single segment, [o] or [u], we can refer to these effects as *intrasegmental coarticulation*. The PTA accounts for intrasegmental coarticulation by assigning a different neuromuscular schema depending on the specification of features from the same segment.

Let us suppose further that, while determining the durational properties of transduced features, the STA temporally extends  $PR_{+ROUND}$  from the vowel onto the preceding consonant, i.e., in the anticipatory direction. This amounts to the more familiar *intersegmental coarticulation*, where transduced features from *different* segments interact. Returning to SRs [lok] and [luk], it is now apparent that (a)  $PR_{+ROUND}$  is different for [o] than for [u] due to its intrasegmental coarticulation with  $PR_{HIGH}$ ; (b) [l]s inherent  $PR_{-ROUND}$  is now temporally overlapping with the  $PR_{+ROUND}$  from the adjacent vowels because of intersegmental coarticulation. It is important to note that

the difference in  $PR_{+ROUND}$  from [o] and  $PR_{+ROUND}$  from [u] will be reflected on the preceding consonant: [l] in [lok] will be articulated differently with respect to lip rounding than [l] in [luk]. Thus, [l] simultaneously bears the effect of both intra- and intersegmental coarticulation. CP allows us to account for such subtle yet systematic phonetic variations in an explicit and straightforward way they follow automatically from PTA and STA which are independently motivated by the need for transduction (see [Volenec and Reiss, 2018](#), §5 for further implications of CP).

CPs transduction is deterministic, which means that it assigns the same neuromuscular schema to each feature every time that feature is transduced. This also includes all cases of feature combinations that lead to intra- and intersegmental coarticulation. CP thus makes another empirically testable prediction: In principle, given a full and correct list of features, it should be possible to exhaustively describe all possible intra- and intersegmental coarticulatory effects just by using the two algorithms proposed by CP.

It should be stressed that CPs outputs, phonetic representations, should not be equated with actual articulatory movements or with the acoustic output of the human body. What is actually pronounced is further complicated in the process of language externalization by a great number of factors. Transduction is accompanied by other performance factors that have no bearing on either phonology or transduction, factors like loudness, muscle fatigue, degree of enunciation, interruptions due to sneezing, and many other situational effects, all of which will have an effect on the final output of the body, and will therefore make (co)articulatory variation seem even greater. For that reason, it is not the case that the articulatory and the concomitant acoustic substance will always be identical for each feature or feature combination. However, this apparent lack of invariance in the realization of a cognitively invariant category is not a matter of transduction, but rather is a result of accidental performance factors.

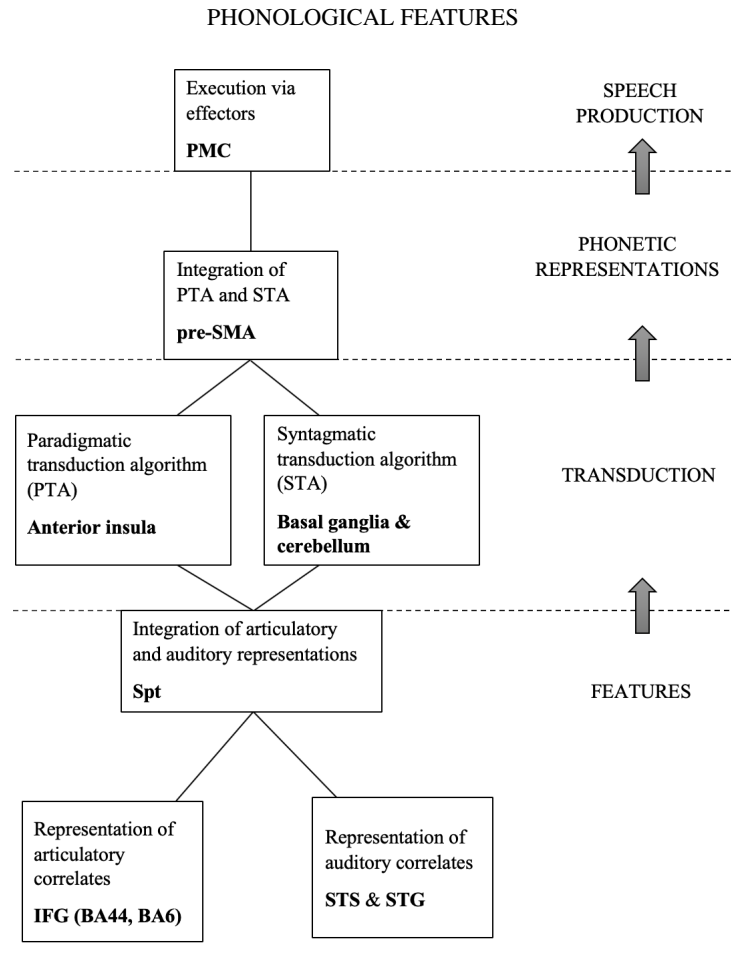
Recent neuroscience evidence is consistent with the idea that CP transduces features, which are to be understood as symbols in the brain, into temporally distributed neuromuscular activities (elements of PRs), thus relating phonological competence to the vastly different SM system (Dronkers, 1996; Hickok and Poeppel, 2007; Guenther, 2016; Eickhoff et al., 2009; Hickok, 2012; Okada et al., 2018). The activity in parts of the inferior frontal gyrus (IFG) corresponds to the representations of the articulatory correlates of features, while the activity in parts of the superior temporal gyrus (STG) and sulcus (STS) corresponds to the representations of the auditory correlates of features. An area in the Sylvian fissure at the boundary between the parietal and the temporal lobe (Spt) unifies these two aspects into a complete symbol, a feature. The symbols are sent to the anterior insula where the PTA is carried out, and to the cerebellum and the basal ganglia where the STA is carried out. The PTA and the STA are integrated in the anterior part of the supplementary motor area (pre-SMA) to form the phonetic representation, which is a set of neural signals that the primary motor cortex (PMC) sends to the effectors that produce speech. These neural processes are graphically represented in Figure 4.

Figure 5 summarizes the general architecture of the phonology-phonetics interface in the theory of Cognitive Phonetics (CP). To connect substance-free phonology with the substance-laden physiological phonetics, CP takes features of phonological surface representations (SRs) and relates them to neuromuscular activity (PTA) and arranges that activity temporally (STA), thus generating phonetic representations (PRs) that are directly interpretable by the sensorimotor (SM) system.

## 5. FINAL REMARKS

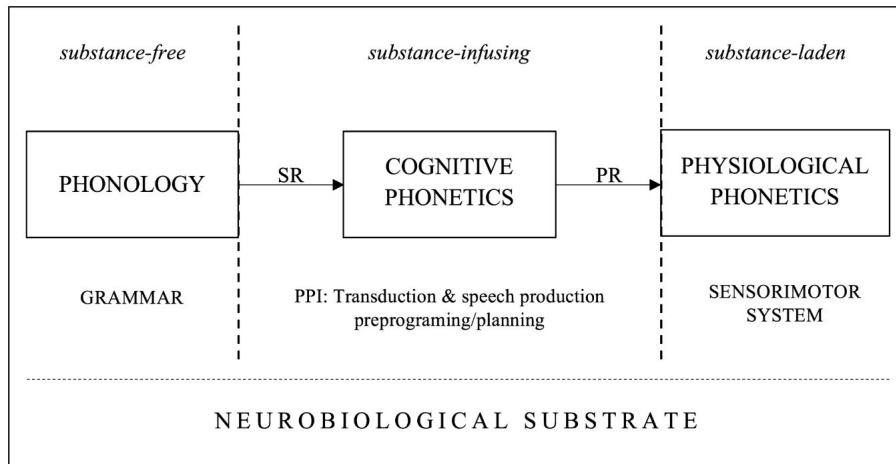
In lieu of a conclusion and still in Tolman’s “cavalier and dogmatic style, let us summarize our main points. The human language faculty exists in the human brain, and only there (*internalism*). Phonology is one part, or one *module*, in technical terms, of





**Figure 4:** Neural substrate of phonological features and of the phonology-phonetics interface.

that cognitive, brain-based faculty. This module consists of computational and representational aspects. The computations are ordered logical operations, rules, and they apply blind to phonetic substance. The representations are made from elementary units, features, which are symbols in the brain, in the sense of Gallistel and King (2009). These symbols are innate (*nativism*)—they are knowledge unlearned and untaught (Halle 1978), which is bestowed by human biology. They are also devoid of phonetic substance. Thus, neither the computational nor the representational aspect



**Figure 5:** The architecture of the phonology-phonetics interface in the theory of Cognitive Phonetics.

of phonology contains phonetic substance (*formalism*), which is why this approach is called *Substance Free Phonology*.

Our main hope with this paper is to mitigate a primal fear related to the innateness of phonological primes, or, more generally, related to the necessary innateness of all elementary (indivisible, irreducible) units of cognition. This fear is so pervasive and apparently so deeply ingrained in empiricist quarters that it prompted Lila Gleitman to remark ironically that empiricism is innate (quoted from Gallistel 2018, 276). Also, with our discussion of the relation between phonological competence and phonetic substance, we hope to have provided a further illustration of the necessity of getting behind the sense data of any type of expression in order to grasp the intuitive [] forms which alone give significance to such expression” (Sapir 1925, 45).

#### REFERENCES

Bale, Alan, Maxime Papillon, and Charles Reiss. 2014. Targeting underspecified segments: A formal analysis of feature changing and feature filling rules. *Lingua* 148:240–253.

- Bale, Alan, and Charles Reiss. 2018. *Phonology: A formal introduction*. Cambridge, MA: MIT Press.
- Bouchard, Kristofer E, Nima Mesgarani, Keith Johnson, and Edward F Chang. 2013. Functional organization of human sensorimotor cortex for speech articulation. *Nature* 495:327.
- Browman, Catherine P, and Louis Goldstein. 1989. Articulatory gestures as phonological units. *Phonology* 6:201–251.
- Chomsky, Noam. 1964. Formal discussion in response to W. Miller and S. Ervin. In *The Acquisition of Language*, ed. Ursula Bellugi and Roger Brown, 35–39. Chicago: University of Chicago Press.
- Chomsky, Noam. 2007. Approaching UG from below. In *Interfaces + recursion = language?: Chomsky’s minimalism and the view from syntax-semantics*, ed. Uli Sauerland and Hans-Martin Gärtner, 1–24. Berlin: Mouton de Gruyter.
- Chomsky, Noam, and Morris Halle. 1965. Some controversial questions in phonological theory. *Journal of Linguistics* 1:97–138.
- Cowper, Elizabeth, and Daniel Currie Hall. 2014. Reductio ad discrimen: Where features come from. *Nordlyd* 41:145–164.
- Dronkers, Nina F. 1996. A new brain region for coordinating speech articulation. *Nature* 384:159–161.
- Eickhoff, Simon B, Stefan Heim, Karl Zilles, and Katrin Amunts. 2009. A systems perspective on the effective connectivity of overt speech production. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367:2399–2421.
- Fastl, Hugo, and Eberhard Zwicker. 2006. *Psychoacoustics: facts and models*, volume 22. Springer Science & Business Media.
- Fodor, Jerry. 1980. On the impossibility of acquiring “more powerful” structures. In *Language Learning: The debate between Jean Piaget and Noam Chomsky*, ed. M. Piatelli-Palmarini, 142–162. Harvard University Press.
- Gallistel, C. R., and Adam Philip King. 2009. *Memory and the computational brain: why cognitive science will transform neuroscience*. Chichester, West Sussex, UK: Wiley-Blackwell.

- Gallistel, C Randy. 2000. The replacement of general-purpose learning models with adaptively specialized learning modules. *The new cognitive neurosciences* 1179–1191.
- Guenther, Frank H. 2016. *Neural control of speech*. MIT Press.
- Hale, M., M. Kissock, and C. Reiss. 2007. Microvariation, variation, and the features of universal grammar. *Lingua* 117:645–665.
- Hale, M., and C. Reiss. 2000. “substance abuse” and “dysfunctionalism”: Current trends in phonology. *Linguistic Inquiry* 31:157–169.
- Hale, Mark, and Charles Reiss. 2008. *The phonological enterprise*. Oxford: Oxford University Press.
- Halle, Morris. 1954. The strategy of phonemics. *Word* 10:197–209.
- Halle, Morris. 1983/2002. On distinctive features and their articulatory implementation. *From Memory to Speech and Back: Papers on Phonetics and Phonology 1954-2002* 3:105–121.
- Hickok, Gregory. 2012. Computational neuroanatomy of speech production. *Nature reviews neuroscience* 13:135–145.
- Hickok, Gregory, and David Poeppel. 2007. The cortical organization of speech processing. *Nature reviews neuroscience* 8:393–402.
- Householder, Fred W. 1965. On some recent claims in phonological theory. *Journal of linguistics* 1:13–34.
- Jackendoff, Ray. 1990. *Semantic Structures*. MIT Press.
- Jander, Rudolf. 1957. Die optische richtungsorientierung der roten waldameise (*formica ruesa* L.). *Zeitschrift für vergleichende Physiologie* 40:162–238.
- Keating, Patricia. 1985. Universal phonetics and the organization of grammars. v. fromkin (ed.) *phonetic linguistics*.
- Keating, Patricia A. 1990. Phonetic representations in a generative grammar. *Journal of phonetics* 18:321–334.
- Kenstowicz, Michael, and Charles Kisseberth. 1979. *Generative phonology: Description and theory*. Academic Press.
- Khalighinejad, Bahar, Guilherme Cruzatto da Silva, and Nima Mesgarani. 2017. Dynamic encoding of acoustic features in neural responses to continuous speech. *Journal of Neuroscience* 37:2176–2185.

- Lenneberg, Eric H., Noam Chomsky, and Otto Marx. 1967. *Biological foundations of language*. New York: Wiley.
- Lent, Roberto, Frederico AC Azevedo, Carlos H Andrade-Moraes, and Ana VO Pinto. 2012. How many neurons do you have? some dogmas of quantitative neuroscience under revision. *European Journal of Neuroscience* 35:1–9.
- Magrassi, Lorenzo, Giuseppe Aromataris, Alessandro Cabrini, Valerio Annovazzi-Lodi, and Andrea Moro. 2015. Sound representation in higher language areas during language generation. *Proceedings of the National Academy of Sciences* 112:1868–1873.
- Marr, David. 1982. *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.
- Matamoros, Camila, and Charles Reiss. 2016. Symbol taxonomy in biophonology. In *Biolinguistic investigations on the language faculty*, ed. Anna Maria Di Sciullo, 41–54. John Benjamins.
- McCarthy, John, and Alan Prince. 1993. Prosodic Morphology I: Constraint interaction and satisfaction. *Technical Report# 3, Rutgers University Center for Cognitive Science* .
- Mesgarani, Nima, Connie Cheung, Keith Johnson, and Edward F Chang. 2014. Phonetic feature encoding in human superior temporal gyrus. *Science* 343:1006–1010.
- Mielke, Jeffrey. 2008. *The emergence of distinctive features*. Oxford: Oxford University Press.
- Okada, Kayoko, William Matchin, and Gregory Hickok. 2018. Neural evidence for predictive coding in auditory cortex during speech production. *Psychonomic bulletin & review* 25:423–430.
- Phillips, Colin, Thomas Pellathy, Alec Marantz, Elron Yellin, Kenneth Wexler, David Poeppel, Martha McGinnis, and Timothy Roberts. 2000. Auditory cortex accesses phonological categories: an meg mismatch study. *Journal of Cognitive Neuroscience* 12:1038–1055.
- Prince, Alan, and Paul Smolensky. 1993. *Optimality Theory: Constraint interaction in generative grammar (technical report)*. New Brunswick, NJ: Rutgers Center for Cognitive Science.
- Reiss, Charles. 2012. Towards a bottom-up approach to phonological typology. *Towards a Biolinguistic Understanding of Grammar* 169–191.

- Scharinger, Mathias, Ulrike Domahs, Elise Klein, and Frank Domahs. 2016. Mental representations of vowel features asymmetrically modulate activity in superior temporal sulcus. *Brain and language* 163:42–49.
- Scharinger, Mathias, Philip J Monahan, and William J Idsardi. 2012. Asymmetries in the processing of vowel height. *Journal of Speech, Language, and Hearing Research* .
- Schleidt, Wolfgang, Michael D Shalter, and Humberto Moura-Neto. 2011. The hawk/goose story: the classical ethological experiments of lorenz and tinbergen, revisited. *Journal of comparative psychology* 125:121.
- Seikel, J, D King, and D Drumright. 2009. *Anatomy & physiology for speech, language, and hearing*: Cengage learning.
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of go without human knowledge. *Nature* 550:354–359. URL <https://doi.org/10.1038/nature24270>.
- Smith, Brian Cantwell. 1996. *On the origin of objects*. Cambridge, MA: MIT Press.
- Tolman, Edward C. 1948. Cognitive maps in rats and men. *Psychological Review* 55:189–208. URL <http://www.sciencedirect.com/science/article/B6X04-4NR7FD9-1/2/7245db5bc21ab41cf1bd0d5ac0b79ea6>.
- Volenec, Veno, and Charles Reiss. 2018. Cognitive Phonetics: The transduction of distinctive features at the phonology-phonetics interface. *Biolinguistics* 11.
- Volenec, Veno, and Charles Reiss. 2020. Formal generative phonology. *Radical: A Journal of Phonology* 2:1–65.
- Werker, Janet. 1995. Exploring developmental changes in cross-language speech perception. In *An invitation to cognitive science*, ed. Daniel N Osherson and Lila R Gleitman, volume 1, 87–106. Cambridge, Mass.: MIT Press, 2nd edition.
- Zador, Anthony M. 2019. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications* 10:3770. URL <https://doi.org/10.1038/s41467-019-11786-6>.