OUP UNCORRECTED PROOF – FIRSTPROOFS, Mon Sep 04 2017, NEWGEN

4.1 REPRESENTATIONS AND RULES IN LANGUAGE

Ray Jackendoff

One thing that has always distinguished Dan Dennett's philosophy is his eagerness to get his hands dirty with the phenomena, to take things apart and see what's inside, to "turn the knobs" on the simplistic scenarios proffered by intuition pumps in the lore (Dennett, 2013). My goal here is to practice what he preaches.

The scenario whose knobs I'll be turning here is what might be thought of as the "lay cognitive scientist's" (and lay philosopher's) view of language: that sentences (one kind of mental representation) are constructed by combining words (another kind of mental representation) by means of rules. This scenario is of course an instance of the central stance of cognitive science: that mental activity can be described as computational operations over mental representations. But all too often this stance is given little more than lip service, and not much thought is devoted to the question of precisely what the mental representations and the operations on them *are*, beyond the demands of the experiment or computational model at hand.

In this chapter I want to wrestle with the nature of representations and rules in language, comparing a number of proposals in the literature to my own (Culicover & Jackendoff, 2005; Jackendoff, 2002). The outcome is a reconceptualization of how knowledge of language is encoded by the brain, with implications for cognitive science at large.

A road map: First, after a prelude about what a "mental representation" could be, we look at how words are encoded in the lexicon (or dictionary). We then examine three different ways to think about rules: the Procedural Strategy, the Declarative Strategy, and the No Rules Strategy. Turning the knobs, we explore a range of phenomena that lie between classical words and classical rules, concluding that the Declarative Strategy has a decided advantage in accounting for the

oso-9780199367511.indd 95 9/4/2017 11:12:19 PM

existence of such items. We then turn some different knobs, comparing fully productive rules—the classical rules, which work all the time—with partially productive rules, which work only *some* of the time. Again, the Declarative Strategy proves superior. Finally, we think briefly about how rules could be acquired, and show how the Declarative Strategy offers a more satisfying account (though still incomplete in many respects).

I recognize that my story here is somewhat technical for non-linguist audiences. Nevertheless, I encourage readers to bear with me. The details *do* make a difference.

1. What Is a Mental Representation?

I want to think of mental representations not as characterizations of whole-brain states (the way philosophers sometimes talk about knowing P as being in some particular brain state), but rather as data structures that pertain to a particular domain of brain activity. A model case is phonological structure in language. The sound structure of a word is encoded as a sequence of phonological segments, each of which can be characterized in terms of the values of a set of phonological distinctive features. For instance, the word *cat* is represented phonologically as the sequence of phones /k/, /æ/, and /t/, where /k/ in turn consists of the features [+consonant, -continuant, -voiced, +velar closure], and similarly for the other segments.

Phonological structure is a level of abstraction beyond the phonetic representation. In the context of beginning a stressed syllable, phonological /k/ is aspirated as phonetic [k']. And the exact phonetic value of phonological /æ/ depends on the consonant following within the same syllable: longer (and in some dialects, a bit higher) if followed by a voiced consonant (as in *pad*), shorter (and lower) if followed by an unvoiced consonant (as in *pat*). And in turn phonetic representation itself is a sequence of discrete sounds, and is therefore is a major abstraction from the continuously varying acoustic signal.

When one hears the word *cat*, the distinctive features are not present in the acoustic signal; they must be computed on the basis of the acoustic signal. (This is one important reason why automatic speech recognition has proven so difficult.) Yet linguists consider these features "psychologically real" because of the role they play in determining the patterns of human languages. A feature doesn't "represent" anything in the physical world, it is not a "symbol" for anything; rather, it imposes an equivalence class on heard and produced sounds that classifies them for linguistic purposes, that is, for human interests. The same goes in the syntactic domain, if not more so, for features such as Verb, Noun Phrase, and Suffix.

Just to be clear: the letter "k" *does* represent something. It *is* a symbol for something, namely the phonological segment /k/. Similarly, the *notation* [-voiced] represents something, namely the value of a feature in a hypothesized mental representation. But the phonological segment and the feature themselves do *not* "represent"—they are not symbols for anything.[1] They get their values only by virtue of their roles in the system of language. One might think of the phonological structure of a word as a sort of virtual object, like a melody or a baseball score or a turn (as in taking turns), or a dollar (as in Dennett's,2013, "How Much Is That in Real Money?"). These are all "cognitively real," that is, they play a role in human construal of (or understanding of) the physical and social world—"institutional facts" in Searle's (1995) terms. But they are not physically real in any useful sense. Again, the same arguments pertain to syntactic categories and syntactic features.

In addition to the varieties of data structures, an essential component of mental representation is the links *among* different spaces of data structures. The word *cat* and the non-word *smat* both have phonological structures. But *cat* is more than that: its phonological structure /kæt/ is linked in long-term memory with a syntactic category and a meaning. In the domain of meaning, the relevant mental representations differentiate concepts into objects, events, properties, places, and so on; the objects are differentiated into animate and inanimate and characterized in terms of shape and function; and so on down to the concept of a cat somewhere in the taxonomy, itself further differentiated into kinds of cats and cats one has known. The concept of a cat (or some version of it) could exist in the mind of a nonlinguistic organism such as an ape or a baby, without being linked to a phonological structure, that is, without being part of a word. Alternatively, in the mind of a speaker of another language, this concept could be linked to a different phonological structure (or, in a sign language speaker, to a gestural structure).

In other words, phonological structures and meanings are independent domains of mental representations, and the links *among* mental domains are as vitally important as the structures themselves. In particular, the space of phonological structures is useful precisely because it offers a rich domain that can be linked to concepts and thereby render them expressible. More generally, without such linkages or "interfaces," a domain of mental representation would be thoroughly encapsulated, and hence useless to the organism in regimenting perception and determining behavior.

1. But they are discrete. The term *symbolic* in cognitive science has frequently conflated being discrete with being a symbol for something.

It is a habit in much of cognitive science to depict an interface as an arrow connecting two boxes. What this obscures is that this process is not like sending water through a pipe, where water goes in one end and comes out the other. For instance, one cannot just "send" phonological representations to semantics. Phonological representations per se are completely unintelligible to semantics; and retinal images per se mean nothing to "higher-level" visual representations such as those involved in face recognition. Hence, even if one notates the interface with an arrow, it has to signify a complex correspondence or conversion between one sort of data structure and another, and this conversion is an essential component of mental computation.

In order to describe cognition, then, it is important to characterize both the data structures in the head and the interfaces among them. At least for the foreseeable future, we can do so only through reverse engineering, through developing theories of mental representations that enable us to understand subjects' behavior (including their reports of their experience—Dennett's (1991) heterophenomenology—which is also a kind of behavior).

In short, a theory of mental computation has to characterize the repertoire of domains (or "levels") of mental representation, such as phonological structure and whatever kinds of structures are involved in encoding meaning. This repertoire may be different for different species. For instance, we probably should not expect kangaroos to have phonological structure in their repertoire. But we (or at least I) would suspect that all humans share a common repertoire of domains of mental representation.

This is not to deny that individuals may differ significantly in the richness of their representations in particular domains. A professional soccer player must have much more highly articulated motor representations than I do, but I certainly don't lack motor representations altogether. In fact, I probably have myself developed some highly specialized motor representations connected with playing the clarinet.[2]

There need to be (at least) two kinds of computations over mental representations. One kind remains within a domain: computing a new representation in the domain on the basis of existing ones. For example, the phonology of a suffix can be adjusted, depending on what it is attached to: the plural of *cat* is /kæts/, with an /s/, but the plural of *dog* is /dɔgz/, with a /z/. Similarly, mental rotation involves computing new visual representations on the basis of existing ones.

2. Are individuals with autism entirely lacking in whatever representations are responsible for theory of mind, or are their representations just sparser?

The other kind of computation uses the interface links between two domains D_1 and D_2 to compute a structure in D_2 on the basis of a structure in D_1. This is the sort of computation involved in any sort of perception, for instance using a phonological input to compute a sentence meaning from the meanings of the individual words and their phonological linear order. Speaking requires a mapping in the opposite direction; reading uses visual representations to compute linguistic structures (primarily phonological), via interface links between the two domains.

Hand in hand with identifying different domains of mental representation and their interfaces, another issue must be addressed: What determines the potential repertoire of data structures *within* each of the domains? And if there is an interface between two domains, what determines the potential repertoire of links? If a domain is finite, the repertoire of structures can be enumerated. But in most any domain of interest, such as vision, language, music, social cognition, or action, an enumeration is impossible. The repertoire is potentially open-ended: the organism is capable of creating novel data structures. Yet such new structures are not unconstrained; they fall into patterns. So an essential issue for cognitive theory is to characterize these patterns. In the domains involved in the language faculty, these amount to what Chomsky (1986) calls "knowledge of language." The issue I will be concerned with here is how these patterns are to be characterized formally.

But first, just a little more context is in order. First, as Chomsky (1965) stresses, there is a further issue of how the patterns that constitute "knowledge of language" get into the brain, that is, how knowledge of a language is *acquired*. Recurring disputes within linguistics, as well as between linguistics and much of cognitive science, concern the cognitive resources necessary for acquiring language—so-called Universal Grammar—and how *they* are to be characterized. Are some of these resources specific to the language faculty, or are they all domain-general? I won't rehearse the debate here (but see Jackendoff, 2002, chapter 4). In any event, whatever these resources are, whether language-specific or domain-general, they have to be wired into the brain prior to learning, since they form the basis for learning.

A further and hugely difficult issue is how all these mental representations and computations over them are instantiated in neural tissue—what Marr (1982) calls the "implementational level." This is more than just a matter of finding brain localization; ideally it involves showing exactly how neurons encode, store, and compute these data structures and their links. And beyond this, there is the issue for developmental biology of how a developing brain comes to have the potential for computation that it does, under genetic guidance and in response to environmental input.

2. Three Models of Linguistic Rules

How are the patterns of mental representations for language to be characterized formally? I want to contrast three major approaches, which I'll call the Procedural Strategy, the Declarative Strategy, and the No Rules Strategy. This section lays them out; the rest of the paper is devoted to deciding among them.

The Procedural Strategy (also called "proof-theoretic" or "generative enumerative," Pullum, 2013) defines the potential repertoire of linguistic representations by means of *derivations* that build linguistic expressions step by step. This is the approach of classical generative grammar and all its descendants in the narrow Chomskyan tradition, including the Minimalist Program (Chomsky 1995, 2002) and Distributed Morphology (Halle & Marantz, 1993; Siddiqi forthcoming) as prominent contemporary subtraditions. In this approach, rules of grammar are taken to be procedures that are available for building structure. For instance, the phrase structure rule (1a) is an instruction to rewrite a VP symbol as a sequence V—NP, or to expand VP into the tree (1b).

(1) a. VP→V – NP

b.

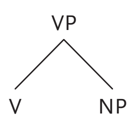

Similarly, the rule determining the form of English plurals can be stated informally as a directive: "To form the plural of a noun, add /z/." The open-endedness of linguistic representations comes from the fact that some rules of grammar are recursive, that is, the rewriting of some symbols such as S and NP can eventually result in an output containing another instance of the same symbol.

Since the early days of generative grammar, it has been recognized that these procedures cannot possibly correspond to the way language is actually processed in the brain. A speaker surely does not create a sentence by successively expanding S into a syntactic tree,[3] then "sending" this (a) to semantics to determine what it means and (b) to phonology to determine how to say it. Standard generative grammar therefore adopts the stance that the linguist's grammar should be thought of not as an active procedure, but simply as a convenient way of defining a set, parallel to inductive definitions of number in mathematics. In turn, this stance leads to the familiar firewall between competence (the characterization of linguistic patterns) and performance (the actual processes taking place in the brain), going back to Chomsky (1965).

3. Or, in the Minimalist Program, starting with a collection of words and building upward to S; or, in Distributed Morphology, deriving the completed syntactic tree first and then populating it with words.

The Declarative Strategy (some versions of which are called "model-theoretic" by Pullum, 2013) defines the repertoire of possible linguistic representations in terms of a set of *well-formedness conditions* on structure. This approach encompasses frameworks such as Lexical-Functional Grammar (LFG; Bresnan, 2001), Head-Driven Phrase Structure Grammar (HPSG; Pollard & Sag, 1994), Construction Grammar (Goldberg, 1995; Hoffman & Trousdale, 2013), Autolexical Syntax (Sadock, 1991), and the Parallel Architecture (Culicover & Jackendoff, 2005; Jackendoff, 2002), among others.

On this approach, a phrase or sentence as a whole is legitimate if all its parts are *licensed* by one of the well-formedness conditions. For example, the rule for the English transitive verb phrase is not a rewriting rule, but simply the piece of syntactic structure (1b) itself: a syntactic schema or "treelet." Any part of a syntactic structure that conforms to this structure is thereby licensed by the rule. Similarly, the English plural rule is stated as a *schema* that licenses plural nouns, along the lines of (2).[4]

(2)

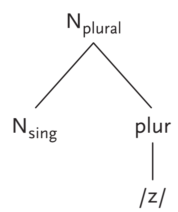

Here, the open-endedness of linguistic forms is a consequence of what might be called "structural recursion": a complex syntactic (or semantic) structure, licensed by a particular schema, can contain subparts that are licensed by the very same schema.

Well-formedness conditions can also implement correspondences between two domains. For instance, one such condition stipulates that an agent of an action (a semantic configuration) is typically expressed by the subject of the verb that denotes the action (a syntactic configuration).

Unlike the Procedural Strategy, the Declarative Strategy ascribes no implicit directionality to linguistic representations. There is no "input" or "output"; nothing is "sent" to another component for further steps of derivation. For instance, it is impossible to stipulate that words are "inserted" into sentences "before" or "after" some step of derivation.

4. I am abstracting away here from the variation between /z/, /s/, and /əz/. In a procedural approach, the plural has an "underlying form," say /z/, and the grammar has rules that in the course of derivation change this to /s/ or /əz/ in the appropriate contexts. In a declarative framework, the plural morpheme simply has three forms that are licensed in different contexts.

Paradoxically, this makes a declarative grammar potentially more amenable to being embedded in a model of language processing. In processing, the grammar is being used to build up a linguistic structure online, over some time course (Jackendoff, 2002, chapter 7; Jackendoff, 2007). A procedural grammar determines an order in which rules are to be applied; and as seen above, this order makes little sense in terms of actual sentence processing. By contrast, in a declarative grammar, the processor builds structure by "clipping together" stored pieces of structure in any order that is convenient, using a procedure called unification (Shieber, 1986), which is arguably domain-general (Jackendoff, 2011).[5] In particular, in sentence comprehension, phonological structure can be taken as "input" and semantics as "output"; and in sentence production, the reverse. A declarative grammar doesn't care in which direction it is put to use.

The No Rules Strategy encompasses a variety of theories, including connectionism (Christiansen & Chater, 1999; Elman, 1990; Rumelhart & McClelland, 1986; Seidenberg & MacDonald, 1999), item- or exemplar-based theories (Bod, 2006; Frank, Bod, & Christiansen, 2012; Pierrehumbert, 2002;), and statistically based approaches such as Latent Semantic Analysis (Landauer & Dumais, 1997). What they have in common is a rejection of the idea that the repertoire of linguistic forms is determined by any sort of abstract schemas such as phrase structure rules. Rather, knowledge of language consists of a list or an amalgam of forms one has encountered, and the open-endedness of language is ascribed to processes that construct or interpret novel expressions by analogy to stored forms.

This approach faces an inherent logical problem. Even if one's knowledge of language lacks rules or schemas, one still has to somehow be able to encode the potentially unlimited variety of instances one encounters and remembers. So the basic question arises: In what format are instances encoded? In order to store and compare tokens, they must be coded in terms of a set of features and/or a set of analog dimensions of variation. The basic tenet of the No Rules Strategy is that such a repertoire cannot be explicit, because that would amount to having rules that define possible instances. So the characterization of the repertoire has to be implicit, built into the architecture in such a way that the distinct domains of mental structure emerge from the input.[6] The crucial question is how

5. Unification is a sort of Boolean union defined over hierarchical structures rather than sets. For instance, the unification of the string ABC with the string BCD is the string ABCD, overlapping at BC—rather than, say [ABC][BCD] or some other combination.

6. Just saying the dimensions of variation are "emergent properties" is not enough. For instance, some aspects of phonology are motivated by (and hence "emerge from") acoustic and motor affordances. But they are not inevitable. German and Dutch have a rule of "final devoicing," such that some words end in an unvoiced consonant at the end of a word, but use the corresponding voiced consonant when followed by a vowel suffix; hence German *Hand*, "hand," is

phonological tokens (for example) come to have such different dimensions of similarity from, say, semantic tokens.

In the Procedural and Declarative Strategies, the variables in rules are the essential source of open-endedness. VP, for instance, stands for *any* VP, and N_{plur} stands for *any* plural noun. The No Rules Strategy will have none of this. It insists that knowledge of language involves no abstraction of the sort captured by variables. Unfortunately, Marcus's (1998, 2001) critique of connectionist approaches shows that a network of the usual sort is in principle incapable of learning and encoding the sorts of relations usually taken to involve two variables—relations as simple as "X is identical to Y"—in such a way that they can be extended to novel instances outside the boundaries of the training set. Yet two-place relations are endemic to linguistic structure: "X rhymes with Y," "X agrees with Y in number," "X is coreferential with Y," and so on. Hence there are strong reasons to believe that connectionist machinery is not up to the task of learning language in general. (And I believe similar arguments apply to other varieties of the No Rules Strategy. See also Jackendoff, 2002, section 3.5.)

The No Rules Strategy has been most successful in dealing with the structure of individual words, starting with Rumelhart and McClelland's (1986) well-known treatment of the English past tense, and in dealing with the impressive degree of frequency- and contextually-induced micro-variation in phonetic realization (as documented by Bybee, 2010) and lexical access (as studied by Baayen, 2007). However, this strategy has been explored and tested almost exclusively on relatively simple morphological phenomena such as English past tense formation and on the syntax of toy languages with minimal structure and a minimal vocabulary. To my knowledge, the approach has not been scaled up to full complex sentences such as this one, whose structure is relatively well understood in terms of either procedural rules or declarative schemas. The basic difficulty seems to be that a No Rules Strategy offers no way to encode and build hierarchical structure, a fatal flaw in a theory of language.[7]

---

pronounced /hant/, but its plural *Hände* is pronounced /hendə/, with phonetic [d] rather than [t]. This alternation can be motivated on physiological grounds: devoicing is a natural outcome of lowered air pressure at the end of a stressed syllable. One might therefore want to argue that there is no need for phonology—that apparently phonological principles are "emergent" from physiology. However, final devoicing is still a phonological principle—part of the grammar of German and Dutch—not a purely physiological one. If it were purely physiological, English would have to do it too, and it doesn't: *hand* is pronounced with a voiced [d]. So the physiology motivates but does not determine phonology (Blevins, 2004).

7. Frank, Bod, and Christiansen (2012) attempt to argue that language lacks hierarchical syntactic structure, but they build hierarchy implicitly into their parsing procedure, and they deal with only the simplest nonrecursive sentence structures.

I have tried in this section to cite relatively pure examples of each strategy. But there are hybrids as well. Tree Adjoining Grammar (Joshi, 1987) is declarative, in the sense that its building blocks are pieces of structure. But the combination of these pieces is still conceived of procedurally. Optimality Theory (Prince & Smolensky, 2004) is based on well-formedness constraints rather than rewriting rules, so in that sense it is declarative, but it still has "input" and "output"—procedural notions—as essential theoretical constructs.

3. Words and Rules Belong in the Same Bucket

How to adjudicate among these three approaches? This section offers an argument for the Declarative Strategy and against the Procedural Strategy; we return to the No Rules Strategy in subsequent sections. The argument "turns the knobs" on the distinction between words and rules, spotlighting linguistic phenomena that fall in the cracks between the two.

A typical dictionary of a foreign language has several hundred pages of words plus a small separate section devoted to grammar. From the very beginning, generative grammar inherited this traditional distinction, assuming (or even asserting—Chomsky, 1965) that the lexicon and the grammar are fundamentally different kinds of mental representations: words are declarative and rules of grammar are procedural. I want to argue against this assumption and show that in giving it up, there is nothing to lose (aside from tradition), and there is actually a great deal to gain.

The argument goes by a slippery slope: there are things that have to be stored in the lexicon, in declarative format, that are progressively more and more rule-like, so there seems less and less reason to distinguish them from things that everyone accepts as rules. Not only is there no clear place to draw the line between words and rules, there is no need to: when you get to the bottom of the slippery slope, you discover it's not so bad down there after all. (Versions of this argument have been made by HPSG, Cognitive Grammar (Langacker, 1987), and Construction Grammar, as well as by me, in Jackendoff, 2002).

Starting with words: In any mentalist theory, words are stored in memory, presumably in some declarative form: the phonological structure /kæt/ is linked to some semantic structure CAT (which I leave uncharacterized here) and some syntactic features such as part of speech, as sketched in (3). (I don't think anyone believes that the word *cat* is a procedure to derive CAT from /kæt/ or vice versa.)

(3) Phonology: /kæt/
Syntax: +N, singular
Semantics: CAT

What about rules? The Procedural Strategy says they are procedures; the Declarative Strategy says they are pieces of structure that license parts of composed structures. But of course, the rules of the grammar, like the words, have to be present in memory in some form or another, either explicitly—or if only implicitly, then in the structure of the processor. Either way, we might want to say the lexicon and the grammar are stored differently in the brain. For instance, the lexicon might be stored in the temporal lobe and the grammar stored or otherwise instantiated in the frontal lobe, more or less as proposed by Ullman 2004.

But let us look at this distinction more closely. Example (3) is just the stereotypical case of a word. There are other variants. Words like those in (4) have phonology and semantics but no identifiable part of speech. They can occur alone as a full utterance, and they occur in combination only in things like *Hello, Bill* and in quotative contexts like (5), where anything at all can be inserted, even an expression in another language.

(4) *Phonology and semantics, no syntax:*
hello, ouch, yes, oops, dammit, upsey-daisy, allakazam, feh, uh-oh, shucks

(5) "Hello," she said. (cf. *"Shema Yisroel," she said.*)

English also has a few words such as those in (6) that have phonological and syntactic features but no semantic features. They function just as "grammatical glue."

(6) *Phonology and syntax, no semantics:*

a.	epenthetic *it*:	**It**'s noisy in here.
b.	*do*-support *do*:	I **did**n't see her.
c.	*N of NP*:	a picture **of** Bill
d.	subordinating *that*:	I know **that** you came.

Nonsense words like those in (7) have phonology, but no syntax or semantics at all; their linguistic function is just to fill up metrical space in poetry or songs. If you know the song, you recognize these the same way you recognize words.

(7) *Phonology, no syntax or semantics*
fa-la-la, hey diddle diddle, e-i-e-i-o, brrr-raka-taka, inka-dinka-doo, rickety-tickety-tin

So we find all different combinations of phonological, syntactic, and semantic features in stored lexical items.

What else has to be stored in the lexicon? It is clearly necessary to store *idioms* and other fixed expressions in some form or another. For instance, *kick the bucket* has phonological structure, plus a semantic structure approximately equivalent to *die*, plus the syntactic structure of a verb phrase. These have to be stored as a unit because the meaning is noncompositional.

(8) Phonology: /kɪk#ðə#bʌkət/
Syntax: $[_{VP}$ V $[_{NP}$ Det N]]
Semantics: DIE (X)

We know that (8) has the structure of a verb phrase, and is not simply an undigested string, because *kick* inflects just like an ordinary verb, e.g., *he kick**ed** the bucket*, not **he kick-the-bucket**ed***.

Another reason to think idioms have internal syntactic structure is that, just like verbs, they can have argument structure. For instance, the idioms in (9) take a freely chosen direct object, just like ordinary transitive verbs, and this direct object goes exactly where a direct object *should* go—which happens to be in the middle of the idiom.

(9) take NP for granted
put NP on ice
give NP the once-over

A theory of knowledge of language cannot shrug idioms off as some kind of exceptional aberration. English has thousands of idioms, perhaps tens of thousands. Hence it is necessary to give up the assumption that declarative memory for language—the lexicon—is simply a list of words: it also contains items with internal syntactic structure.

Other phenomena use standard syntax, but to unusual semantic ends. Consider (10).

(10) Willy whistled his way down the hall.

Syntactically, this sentence has a typical verb phrase. However:

- The normal direct object of the verb *whistle* denotes an acoustic phenomenon (*whistle a tune* or *a signal*), but instead, (10) has this strange phrase *his way* in object position.

- *Whistle* also doesn't usually occur with a path expression like *down the hall*. (**Willy whistled down the hall*.)
- The meaning of the sentence involves Willy *going* down the hall, even though there is no verb of motion.
- What *whistle* is doing is describing the manner or concurrent activity with which he goes down the hall.

In other words, (10) has a non-canonical mapping between syntax and semantics, marked by the phrase *his way* in object position. Any verb of the right semantic type is possible: you can drink your way across the country or knit your way through a conference. This sort of phenomenon might be called a "constructional idiom," and the case in (10) might be sketched informally as (11).

(11) *Way*-construction:
Syntax/Phonology: [$_{VP}$ V *pro*'s way PP]
Semantics: 'NP goes PP while/by V-ing'

(12) gives two other mappings of this sort, with different semantics, marked by *away* and *head* (or other body part) *off*.

(12) a. Time-*away* construction:
Syntax/Phonology: [$_{VP}$ V [$_{NP}$ (*time*)] away]
Semantics: 'NP spends/wastes NP V-ing'
e.g. *Fred drank the afternoon <u>away</u>*. (= 'Fred spent/wasted the afternoon drinking')

b. *Head off* construction:
Syntax/Phonology: [$_{VP}$ V *pro*'s head/tush/butt off]
Semantics: 'NP V-s a lot/intensely'
e.g. *Suzie sang <u>her head off</u>*. (= 'Suzie sang a lot/intensely')

Knowing these constructions is part of knowing English. For each one, a speaker has to learn and store something about its syntactic structure, something about how its constituents correspond to semantics in other than the normal way, and something about the phonology of the designated elements *way, away*, and *head off* that signal that something unusual is going on.

Other constructions of this sort, such as the four in (13), have no distinguishing phonological content, so they're not very wordlike at all.

(13) a. Sound+motion construction
$[_{VP}$ V PP]
The bus rumbled around the corner. (= 'The bus went around the corner, rumbling')

b. Inverted NP construction:
$[_{NP}$ a/this/that N of an N]
that travesty of a theory (= 'that theory, which is a travesty')

c. Light verb construction:
$[_{VP}$ V NP NP/PP]
Pat gave Kim a hug. (= 'Pat hugged Kim')

d. Casual paratactic conditional:
[S, S]
You break my chair, I break your arm. (= 'If you break my chair, I will break your arm')
(Edward Merrin, June 12, 2014)

The bus rumbled around the corner means that the bus *went* around the corner—even though there is no verb of motion—and it was simultaneously making rumbling sounds. In *that travesty of a theory*, the syntactic head is *travesty*, and *theory* is a modifier; but semantically, *theory* is the referent of the expression and *travesty* is a modifier. In the light verb construction, the direct object *a hug* provides semantic content that is normally expressed by a verb, and the verb is more or less a dummy that provides argument positions and a host for tense. Finally, the paratactic conditional has no *if* or *then*, but the conditional meaning is perfectly clear.

Again, a speaker has to store knowledge of the constructional idioms in (13) as an association between a syntactic complex and a semantic complex. But the basic formalism is the same as for words—except that, as in idioms, the syntax is composite, not just a couple of features, and this time there is no phonology. The tradition of Construction Grammar (Goldberg, 1995; Hoffmann & Trousdale, 2013) has studied a large number of constructions of the type in (12) and (13), as have I (Culicover & Jackendoff, 2005; Jackendoff, 2010). Languages are full of them.

We now reach the crucial point in the argument. Using the very same formalism, we can state a language's phrase structure rules as syntactic schemas, without any association to phonology or semantics. For example, (14) is a partial rule for the transitive verb phrase in English. Tomasello (2003) calls this sort of thing a "slot-and-frame schema," using the notation in (14a); Janet Fodor (1998) calls it a "treelet," in the alternative notation (14b), which is our original structure for the VP in (1b).

(14) a. Slot-and-frame schema:
$[_{VP}$ V – (NP) …]

b. Treelet:

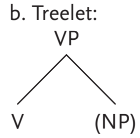

In light of what we have seen up to this point, (14) should be thought of as another type of lexical item. There are lexical items with just phonology like *fiddlededee*, items with just phonology and syntax like the *do* of do-support, items with just phonology and semantics like *hello*, and items such as idioms and meaningful constructions with complex syntactic structure. (14) happens to be an item with only syntactic structure—just one more logical possibility in the system. Its variables license novel, freely combined VPs in syntax.

At this point we have slid to the bottom of the slippery slope from words to rules, and it actually makes some things work better. The lexicon—the collection of declarative knowledge of language—includes both words like *cat* (3) and phrase structure rules like that for VP (14), as well as all sorts of things in between, such as idioms and meaningful constructions. The standard approach of generative grammar, in which words are declarative knowledge and rules are procedural knowledge, has to draw an unnatural line somewhere among these intermediate cases (and in practice, it has ignored them). In contrast, in a declarative conception of rules of grammar, this continuum is perfectly natural.

This leads to the radical hypothesis that all phenomena that have previously been called rules of grammar can be reformulated as pieces of structure stored in the lexicon. In other words, knowledge of language is wholly declarative, and the only procedural component—the only "operation on mental representations" that assembles larger structures—is unification. Furthermore, unification can assemble structures within any domain of mental representation, so strictly speaking, it is not part of "knowledge of language" per se.

This hypothesis presents numerous challenges for linguistic theory. For one thing, the domains of morphology (word structure) and phonology (sound structure) have nearly always been worked out in procedural terms, and it remains to be demonstrated that they can be reformulated in terms of purely declarative rules (see Jackendoff & Audring, 2016, forthcoming, for a beginning). Moreover, within syntax itself, the procedural notions of movement and deletion have to be expunged from the syntactic armamentarium. So it is crucial to figure out how a declarative theory of rules deals with phenomena such as passive, subject-auxiliary inversion, wh-question formation, and so on, which have always seemed very natural under a movement hypothesis. There is no space here to go through the arguments, but HPSG, LFG, Simpler Syntax (Culicover & Jackendoff, 2005), and various other frameworks are all attempts to work this out, with a good deal of success.

4. Full and Partial Productivity

Twirling a different knob on the lay cognitive scientist's conception of language, we next consider rules governing *morphology*, the interior structure of words. Traditional grammar tends to view morphology procedurally. For instance, it is natural to think of deriving the word *procedurally* by adding the suffix –*ly* to the word *procedural*, which in turn is derived by adding the suffix –*al* to the word *procedure*, which is itself derived by adding –*ure* to *proceed*. The earliest generative treatments of morphology, such as Lees (1960), took over this assumption, deriving morphological structure by means of transformations. For instance, Lees had transformations that turned the sentence *John constructed a theory* into the NP *John's construction of a theory.*

The difficulty with this approach is that transformations are supposed to be completely regular and productive, but lots of morphology *isn't* regular or productive. Three symptoms appear over and over again.

> Symptom 1: In many cases, although evidently a rule is involved, the instances of the rule still have to be listed individually. For example, many "derived" nouns consist of a verb plus the suffix –*ion*, and many more consist of a verb plus –*al*. But these suffixes cannot be applied indiscriminately, in the fashion that standard rules would dictate: *confusion* and *refusal* are real words, but the perfectly plausible **refusion* and **confusal* are not. These idiosyncratic facts must be recorded somewhere in the grammar or the lexicon.
>
> Symptom 2: Stereotypically, so-called "derived" items have words as their roots: *construction* is based on *construct* and *permission* is based on *permit*. But there happen to be many "derived" items without a lexical word as their "root," such as *commotion*, *tradition, ambition*, and *retribution*. This is not a matter of rare exceptions. For instance, among the hundreds of English adjectives that end in –*ous*, over a third of them lack a lexical root, for example *tremendous, gorgeous, salacious*, *impetuous*, and *supercilious*.
>
> Symptom 3: The meanings of so-called "derived" items are often partly idiosyncratic with respect to the meaning of their roots, so they cannot be generated by a general meaning-blind procedure. Consider for instance the difference between the meanings of *recite* and *recital*, or *proverb* and *proverbial*. Again, this phenomenon is extremely widespread.[8]

I will use the term "partially productive" for patterns with some combination of these three symptoms. Such rules contrast with fully productive patterns, which

8. A common impulse is to appeal to the history of English, for example, borrowings from Latin, to account for these facts. However, speakers of English are typically unaware of the origin of the words they learn. So on the whole, the history of the language plays no role in its speakers' knowledge.

apply to *every* eligible form aside from listed exceptions, and whose meanings are completely predictable. Speakers are happy to apply fully productive patterns to novel items, for example unhesitatingly giving the plural of *wug* as *wugs* (Berko, 1958).

Chomsky and Halle (1968) and Lakoff (1970) retained the Procedural Strategy, and accounted for partially productive phenomena by means of "exception features." For instance, *refuse* would be marked as unable to undergo the *–ion* transformation, and *tremend* would be marked as obligatorily undergoing the *–ous* transformation. While this approach deals with symptoms 1 and 2 by brute force,[9] it does not address symptom 3.

Chomsky's (1970) solution to partial productivity, the so-called Lexicalist Hypothesis, was to put partially productive morphology "in the lexicon," prior to inserting words into sentence structures. He was not very specific about how this would work; but at the time, people thought about his proposal in terms of "lexical rules"—rules in the lexicon distinct from rules of phrasal grammar. For instance, Wasow (1977) produced widely accepted arguments that there is a distinction between fully productive syntactic passives such as (15a), produced by syntactic rules, and partially productive adjectival passives such as (15b), produced by lexical rules.

(15) a. Pat was hugged by Kim.
b. Pat was very surprised at the news.

Some subsequent syntactic frameworks such as Lexical Functional Grammar (Bresnan, 1982, 2001) and Head-Driven Phrase Structure (Pollard & Sag, 1987) have taken this distinction as foundational, as have some morphological theories (e.g., Anderson, 1992, Stump, 1990).

Some approaches to morphological rules (e.g., Aronoff, 1976; Halle, 1973; Halle & Marantz, 1993; Lieber, 1992) take them to be encoded as procedural rules that create morphologically complex words such as *procedurally* from simpler words or stems stored in the lexicon. However, if one takes symptoms 1–3 seriously, a procedural account of partially productive rules is problematic in the same way as Lakoff's account. First, it still has to somehow list which affixes can be added onto *refuse* and which ones onto *confuse*. Second, a word like *tremendous* has no root in the lexicon to which the suffix *–ous* can be added, so it is

9. Actually, to a first approximation only. For example, *ambi* would have to be marked to obligatorily undergo *either* the *–tion* transformation (to form *ambition*) or the *–ous* transformation (to form *ambitious*). In other words, one would have to invoke Boolean combinations of exception features.

unclear how to generate it (unless **tremend* is listed as requiring a suffix). Finally, putting *recite* and –*al* together does not alone specify what *recital* means.[10]

Other approaches (e.g., Blevins, 2006; Bochner, 1993; Booij, 2010; Jackendoff, 1975) take lexical rules to be declarative patterns that characterize relations among full words stored in the lexicon. For instance, both *refuse* and *refusal* are stored in the lexicon, and the –*al* rule does not *derive* one from the other, but rather stipulates how these forms are related. The –*al* rule itself is stated as a schema containing variables: it expresses the possibility for a noun to consist of a verb plus the suffix –*al*, as in (16). (*V* in syntax is a variable for the verb, X is a variable for its meaning in semantics, and / . . . / is a variable for its phonological form.)

(16) Phonology: / . . . / – al
Syntax: $[_N$ V – suffix]
Semantics: ACT-OF X-ING *or* RESULT-OF X-ING

This approach fares somewhat better with the three symptoms of partial productivity.

> Symptom 1: Since the declarative theory lists all the forms that exist, the fact that there is no word **refusion* has nothing to do with features of *refuse*; there simply is no such word listed in the lexicon. To be sure, the lexical rule for –*ion* says that if there *were* such a word, it could be related to *refuse*; but it does not predict there to be such a word.
>
> Symptom 2: The fact that there is a word *tremendous* but no **tremend* is the mirror image of this situation: the –*ous* rule says that if there were such a word as **tremend*, it could be related to *tremendous*. But again, it does not *require* there to be such a word.

10. Some strands of generative grammar, such as Hale and Keyser (1993, 2002), revert to accounting for morphologically complex items through syntactic operations such as head-raising. For instance, they derive *saddle the horse* from something like *put a saddle on the horse*. But their derivational mechanism has the very same three problems as Lees's approach: how is application of the rules restricted to only certain items, how is something derived when there is no root to derive it from, and how are semantic idiosyncrasies accounted for? To my knowledge, none of these three problems has been seriously addressed within the Hale & Keyser approach.

Distributed Morphology (Halle & Marantz, 1993; Harley 2014; Siddiqi forthcoming) accounts for symptoms 1–3 by appealing to a Vocabulary list of existing forms and an Encyclopedia list of word meanings—more or less the position taken here, but distributed across a procedural derivation. However, I am not aware of any formal characterization of the Vocabulary or the Encyclopedia within this framework.

> Symptom 3: The fact that *recital* means something other than 'act of reciting' is tolerated by the lexical rule, which motivates but does not fully determine the meaning relation between the two words. In other words, lexical rules can be considerably more flexible if stated declaratively than if stated procedurally.

How does a morphological schema relate forms such as *recite* and *recital*? An approach that has been broadly adopted (HPSG: Pollard & Sag, 1994; Construction Grammar: Goldberg, 1995, 2006; Construction Morphology: Booij 2010; Parallel Architecture: Jackendoff, 2002) treats words and morphological schemas as nodes in an inheritance hierarchy. A word "inherits" structure from smaller words and more abstract schemas that predict parts of its form. The predicted parts of its structure are redundant, and therefore in some sense "cost less": there is less to learn.[11] For instance, as shown in (17), *recitation* inherits structure both from the word *recite* and from the deverbal *-tion* schema—but it also has idiosyncratic semantic structure of its own which must be "paid for." An item such as *commotion*, which lacks a lexical root, also inherits structure from the *-tion* schema. But since there is no root *commote*, it has to "pay" for its pseudo-root. As shown by Bochner (1993), schemas can also inherit structure from more abstract schemas. For instance, the *–tion* schema inherits structure from the more general nominalization schema, which is also ancestor to other nominal schemas such as *–al*.

(17)

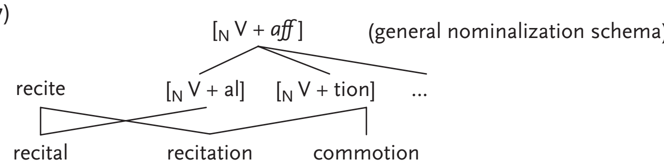

One thing that makes the notion of inheritance hierarchies attractive is that the same sort of machinery is needed for the relations among concepts: the concept of *animal* is ancestor of concepts like *cat* and *dog; dog* is the ancestor of *poodle* and *collie*, and so on, as in (18) (Murphy, 2002). In other words, this sort of taxonomic organization stands a chance of being domain-general, hence not specifically part of "knowledge of language."

11. The notion of "cost" is played out differently in different theories. See Bochner, 1993; Jackendoff 1975; Jackendoff and Audring (forthcoming).

(18)

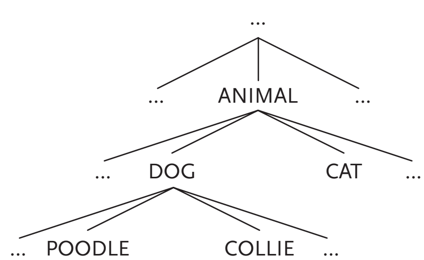

Inheritance hierarchies raise many tricky questions—both formally and psycholinguistically (Jackendoff & Audring, 2016, forthcoming). Setting these issues aside, what is interesting is that they allow us to formalize a morphological rule like the *–al* rule as a piece of linguistic structure containing variables, as in (16). There is no need for a new kind of formal rule: like the idioms and constructions discussed in section 4, these schemas are lexical entries too.

However, we face a new problem: the very same format can be used for *fully* productive morphology. For instance, English present participles, which are completely regular, can be defined by the schema (19) (the meaning is assigned by other principles).

(19) Phonology: / . . . / – ing
Syntax: $[_{\text{Vprespart}}$ V – suffix]

For this sort of fully productive morphology, the standard assumption is that *not* every form is stored as a descendant in the inheritance hierarchy. When speakers encounter a new verb like *skype*, they know automatically that its present participle is *skyping*, and someone hearing this for the first time doesn't bat an eyelash. So this case looks more like standard syntax, and it should be "in the grammar," not "in the lexicon"—just like the schema *[$_{VP}$ V- NP]*. But if partially productive lexical rules have the same format as fully productive syntactic rules, how can we tell them apart? Should the fully productive rules be "in the grammar" and the partially productive rules be "in the lexicon," as Chomsky posited?

A radical answer, proposed by the connectionists (e.g., Bybee & McClelland, 2005; Rumelhart & McClelland, 1986), is that there are *no* explicit rules or schemas for morphological patterns—either partially or fully productive; there are only statistical tendencies based on the amalgamation of learned examples. As we saw in section 3, the larger claim of the No Rules Strategy is that no rules are necessary for *any* patterns in language: memory encodes only individual words (at best), and novel inputs are understood simply by analogy with or by interpolation among memorized instances.

The psycholinguists, led by Steven Pinker (e.g., Pinker, 1999; Pinker & Prince, 1988), dispute the No Rules position on psycholinguistic grounds that we need not go into here, as well as on theoretical grounds, some of which were mentioned in section 3. They argue instead for a hybrid "dual-route" model, in which rules are indeed necessary for fully productive patterns like *–ing* words. However, they grant the connectionists' position for partially productive patterns like *–al* words: these have no rules, only statistical patterns of instances.[12]

6. Dual Route—But Where?

So we are now faced with a three-way choice between Procedural, Declarative, and No Rules strategies, in two distinct domains: fully productive and partially productive patterns. Table 1 sums up the situation. Three of the six possibilities have already been ruled out by considerations raised in sections 3 and 4.

Some theories are homogeneous or uniform, in that they use the same types of rule for fully and partially productive patterns. For instance, Lees's (1960) theory is homogeneously procedural; the connectionist theory is homogeneously No Rules (in fact, it denies that there is any distinction between fully and partially productive patterns).

Other theories are hybrid, using different types of rules for the two. For instance, in Chomsky's Lexicalist Hypothesis, fully productive rules are in the grammar and are procedural, while partially productive rules are in the lexicon and ambiguously procedural or declarative. Pinker's dual-route theory is another hybrid. It takes the fully productive rules to be procedural: "To form the past tense of a verb, add *–d*"; but there are no rules for partially productive patterns. Many contemporary morphologists, for example, Spencer (2013), adopt a similar hybrid position.

Jackendoff (2002) essentially adopts Pinker's hybrid position, but substitutes declarative rules like (19) for Pinker's fully productive procedural rules. I have since become skeptical of this position, because it entails that, for instance, there is no *–al* suffix in the grammar of English, only a lot of *–al* words. This now strikes me as too deflationary. I would like to be able to say that knowledge of English includes an explicit declarative *–al* schema along the lines of (16).

Of the remaining options in Table 1, the choice boils down to whether partially productive patterns are best accounted for in a declarative or a No Rules approach. If the former, there is a homogeneously declarative account of all rules of language.[13]

12. I reluctantly fault Pinker for choosing to fight the connectionists on their own ground, without forcing them to face up to the larger challenges of hierarchical structure in phrasal syntax. Even if one can eliminate rules for English past tenses, this does not show that they can be eliminated for verb phrases and relative clauses. Pinker could have invoked these considerations in his argument, opening the debate immediately to a broader domain.

13. Pending, of course, a demonstration that all morphology and phonology can be treated this way; see Jackendoff & Audring (forthcoming) for a start.

Table 1. Possible positions on the nature of rules

	Fully productive patterns	Partially productive patterns
Procedural	Not possible (shown in section 4)	Not possible (shown in section 5)
Declarative	Possible (shown in section 4)	Possible
No Rules	Not possible (shown in section 4)	Possible

If the latter, we retain Pinker's dual route account, but state fully productive phenomena declaratively instead of procedurally.

In order to decide between these two options, I need to work out the declarative option a bit further. So for the moment let us make the assumption that there *are* partially productive schemas like the –*al* rule (16), and ask how they could differ from fully productive schemas like the –*ing* rule (19).

Recall that Chomsky (1970) put the fully productive rules "in the grammar" and the partially productive rules "in the lexicon." But in the approach developed in section 4, fully productive rules are schemas in the lexicon, just like partially productive rules. So Chomsky's solution is not available to us. Moreover, since we are *assuming* that there are schemas for partially productive rules, Pinker's solution is not open to us either. What options *are* available?

One possibility would be to simply tag schemas with the features [*fully productive*] or [*partially productive*]. A partially productive schema such as the –*al* schema would require an enumeration of its instances in the lexicon. A fully productive schema such as the –*ing* schema would not have such a constraint: it would allow free application to novel forms that are not listed.

I wish to propose something a bit more delicate: the distinction between partial and full productivity is marked, not on entire schemas, but on the variables within schemas. I'll call a variable marked for full productivity an *open* variable, and one marked for partial productivity a *closed* variable, as illustrated in (20).

(20) a. (Partially productive –*al* rule)
Phonology: $/ \ldots /_{closed}$ – al
Syntax: $[_N\, V_{closed}$ – suffix$]$
Semantics: ACT-OF X_{closed} -ING *or* RESULT-OF X_{closed}-ING
b. (Fully productive –*ing* rule)
Syntax: $[_{Vprespart}\, V_{open}$ – suffix$]$
Phonology: $/ \ldots /_{open}$ – ing

A schema with a closed variable captures generalizations among listed items, through the inheritance hierarchy. A schema with an open variable can be used to create forms on the fly, such as *imprinting* and *differentiating*—and even novel forms such as *wugs* and *skyping*. But in addition, a schema with an open variable can serve as an ancestor of stored items. For instance, the productive English plural schema is an ancestor of stored items like *scissors* and *trousers*. In other words, open variables can do what closed variables can do—*plus* create new forms. In a sense, one can think of fully productive rules as partially productive rules that have "gone viral."

Turning the knobs again, here is one piece of evidence for this solution. English has four different principles for naming geographic features, illustrated in (21). Each one involves a name and a descriptive term (underlined) for the type of geographical feature.

(21) a. Arrowhead Lake, Biscayne Bay, Loon Mountain, Claypit Pond, Wissahickon Creek
b. Lake Michigan, Mount Everest, Cape Cod
c. the Indian Ocean, the Black Sea, the Hudson River, the Ventura Freeway, the San Andreas Fault
d. the Bay of Fundy, the Gulf of Aqaba, the Isle of Man, the Cape of Good Hope

Speakers of English know hundreds of place names of the sorts illustrated in (21). Their structure is inherited from the schemas in (22).

(22) a. [$_{N}$ Name$_{open}$ N$_{closed}$]
b. [$_{N}$ N$_{closed}$ Name$_{open}$]
c. [$_{NP}$ the Name$_{open}$ N$_{closed}$]
d. [$_{NP}$ the N$_{closed}$ of Name$_{open}$]

These schemas can be used freely to create new place names. If you want to name a mountain for Dan Dennett, you know immediately that you can call it *Dennett Mountain* (schema 22a) or *Mount Dan* (schema 22b). So the Name variable is completely open.

On the other hand, which of the four schemas has to be used depends on the word for the type of feature—whether it's *lake* or *ocean* or *mountain* or *mount*. Speakers of English know which words go in which schemas: for instance, our mountain can't be called **Daniel Mount* (schema 22b) or **the Mountain of Dennett* (schema 22d). So the variable for the type of geographic feature is only partially productive: the instances have to be learned one by one. Overall, then, these schemas have one variable of each type, and it is impossible to mark the schema as a whole as fully productive or not.

More generally, this case presents evidence against Chomsky's hypothesis that partial productivity is "in the lexicon" and full productivity is "in the syntax." In order to formalize this little subgrammar of place names, partial and full productivity must be in the same component of the grammar, as notated in structural schemas like (20) and (22).

I also conclude from this case that the difference between full and partial productivity—Pinker's dual route—need not be a matter of whether there is a rule or not, but can lie simply in the type of variable within schemas in the lexicon. In fact, one might say that there are indeed two routes: closed variables make use of one of these routes, and the open variables make use not just of the *other* route, but of *both*.

7. Aren't Morphology and Syntax Different?

It might be objected that this solution slides over an important difference between morphology and syntax: *morphology* may be full of these crazy partially productive patterns, but surely *syntax* isn't! If, as just proposed, the distinction between full and partial productivity in morphology is localized in the variables of schemas, then, since syntactic rules are also schemas, they could in principle be partially productive as well. At least according to the lay cognitive scientist—but also according to Chomsky's Lexicalist Hypothesis, this is wrongheaded, missing the point of what syntax is about.

One could resolve the issue by brute force: one could simply stipulate that variables in syntactic rules are all fully productive, but that those in morphological rules can be of either flavor. But this would not be the right resolution. Although syntax does tend to be more reliably productive than morphology, it has some strange little pockets that are genuinely partially productive. Consider again the schemas in (22c,d). We can tell these are patterns of phrasal syntax, because adjectives can be inserted in their usual position within NP: *the lordly Hudson River, the dangerous Bay of Fundy*. Yet, as just shown, one of the variables in each schema is closed.

Another such case is part of the English determiner system. There is no productive rule behind the list of the determiner patterns in (23).[14]

14. Plurals show further eccentricity:

(i) a. such/what trees
  b. *that beautiful trees (cf. *trees that beautiful*)
  c. *how beautiful trees (cf. **trees how beautiful*)

(23) a. such a tree
b. what a tree
c. that beautiful a tree
d. how tall a tree
e. *very tall a tree
f. *more beautiful a tree
g. *twenty feet tall a tree

(24) shows the inheritance hierarchy for this construction.

(24)

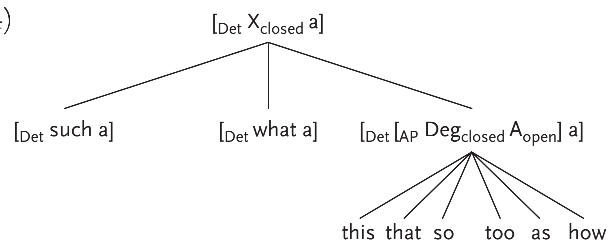

The overall pattern is a determiner of the form *X a*, which attaches to a noun. This pattern is partially productive, and it has a small number of instances including *such a, what a*, and a sub-pattern containing an adjective phrase. Within the adjective phrase, the choice of degree word is partially productive, restricted to *this, that, so, as, too* and *how*. But the choice of adjective is fully productive: we can say things like *how tantalizing an idea* or *as ridiculous a movie as I've ever seen*.

In short, partial productivity extends across the entire lexicon, including syntax; it is not specific to morphology, as people often think.[15] By marking variables open or closed, it is simple to make the distinction between syntactically productive and partially productive patterns. In contrast, it is uncomfortable at best to generate patterns such as (22c,d) and (23) procedurally.

I conclude that a uniformly declarative theory can still maintain the distinction that Chomsky's Lexicalist Hypothesis and Pinker's dual route are meant to capture. But the distinction is not whether there is a lexicon-grammar distinction, as Chomsky proposes, or whether there are rules or not, as Pinker puts it. It is simply how the variables in the rule are marked for productivity.

15. For other cases, see Culicover (1999); Culicover and Jackendoff (2005, section 1.5); Kay (2013); Taylor (2012). Kay advocates for a version of Pinker's position: grammar concerns fully productive constructions only, while partially productive constructions belong instead to "meta-grammar," which operates by analogy rather than by rules.

8. Learning Partially and Fully Productive Schemas

But is this the right approach? Or can we get away with No Rules for partially productive patterns? To answer this, let us turn one last knob, and consider how a learner acquires schemas. They are not present in the input. At best, what *is* present in the input (once the learner has some command of phonology) is examples of utterances, which a schema-less learner must acquire and store in memory individually. The usage-based theorists (e.g., Bybee, 2010; Taylor, 2012; Tomasello, 2003) are perfectly correct on this point. However, the end state of acquisition has to include schemas for fully productive patterns like NP and VP. Where do they come from?

Various researchers from various schools of thought (e.g., Albright & Hayes, 2003; Culicover & Nowak, 2003; Tomasello, 2003) have conjectured that an active process in the learner's mind/brain looks for similarities among stored items. When this process finds similar items (or a sufficient number of similar items), it extracts the similar parts and builds them into a schema that replaces the differing parts with a variable. For example, if the lexicon contains items ABC, ABD, and ABE, this process adds a schema ABX to the lexicon, and ABC, ABD, and ABE listed as its daughters in the inheritance hierarchy.[16] The process can continue recursively, constructing more and more general schemas.

This approach, while appealing, raises a further question that has not to my knowledge been addressed clearly in the literature. Suppose language learners deduce a linguistic rule from the primary linguistic input. How do they decide whether a particular pattern is fully or only partially productive, and what is at stake?

In a hybrid theory of any sort, this choice entails deciding which component of the grammar this pattern belongs in—a rather radical choice. In Chomsky's approach, the choice is between being "in the syntax" and being "in the lexicon." In Pinker's approach, the choice is whether the pattern is captured by a rule or whether there are merely associations. In contrast, the uniformly declarative approach proposed in section 6 makes the choice formally simple: the learner has to decide only which diacritic to place on a variable. That is, the choice is localized within a particular lexical entry—a much more straightforward distinction.

Culicover & Nowak (2003) and Tomasello (2003) propose (in our terms) that the learner initially treats a variable as closed, that is, that the default case is partial productivity. A partially productive schema would help learners understand and

16. There is the additional difficult problem of establishing the domain of the variable: Is it set to be maximally general, or to the minimum range needed to encompass the observed instances, or somewhere in between? See Albright and Hayes (2003).

store forms they have never heard before, but they would not be inclined to *produce* new forms without evidence for these forms in the primary linguistic data. This seems to accord with the lore about language acquisition: children, at least at first, are relatively conservative in their generalizations.

If the learner creates a schema and initially assumes it is partially productive, the next decision is whether to "promote" its variable from closed to open. Both kinds of variables can register patterns of relations among items in the lexicon, but a fully regular variable has two further capabilities. First, it can be used to create new forms such as *wugs* and *skypers* without prior evidence that they exist. Second, it eliminates the need to store new forms as they are heard. If one learns the word *patio*, one doesn't need to also learn (or construct) the word *patios*; it can be understood the next time on the basis of its composition. So at the cost of using the non-default value of the closed-open feature, one gains the benefits of greater flexibility and less burden on long-term memory.

Now here is a crucial point of logic. The process that upgrades a variable from closed to open is not omniscient. It has no way of knowing in advance whether a pattern is going to be fully productive and or only partially productive. And it can't wait until it has all the evidence: if it is considering setting up a schema, it has to store its hypothesis in order to test it against subsequent evidence.

What is the form of a hypothesis? A theory of mental representation requires us to ask this question—though it is rarely asked. In the present approach, the simplest way to encode a hypothesis is as a tentative schema. If the hypothesis is validated, the relevant variables in the schema are upgraded to open. But if it fails, that does not invalidate the original assessment that a pattern exists. Hence alongside the fully productive schemas, one should expect to find lots of partially productive schemas that have failed to be promoted. Since such schemas are capturing partial regularities, there is no reason to discard them. We should expect the lexicon to be littered with failed (but still not useless) hypotheses.

In short, partially productive schemas are potential steppingstones to fully productive rules. Fully productive schemas start life as partially productive schemas, and when they are promoted, they do not lose their original function of supporting inheritance hierarchies among listed items. Hence it stands to reason that there must be items with *just* the original function, namely partially productive schemas. This is altogether natural in a uniformly declarative theory of rules.

By contrast, a No Rules theory claims that there are no partially productive schemas that lead the way to the construction of regular rules. For hard-core No Rules theorists, this is not a problem, because in any event there are no regular rules to be constructed. But the failure of No Rules approaches to account for the free combinatoriality of language—which requires open variables—casts doubt on this option.

This conclusion also reinforces my doubts about Pinker's hybrid position, which grants the existence of rules for fully productive patterns but not for partially productive ones. If you need partially productive rules to attain fully productive rules, then you need partially productive rules, period.

In addition, this account of acquisition upends the notion, traditional within generative grammar and much of morphological theory, that only the fully productive rules are the real business of linguistic theory, and that the unruly "lexical rules" are relatively "peripheral." In the present approach, *all* patterns in language are expressed by schemas, of which a special subset happen to be fully productive. One cannot attain fully productive rules without a substrate of partially productive ones.

9. Ending

Returning to larger issues mentioned at the outset, the overall conclusion here is that:

- The mental representations involved in knowledge of language have the format of pieces of linguistic structure in phonology, semantics, and syntax (including both phrasal syntax and morphosyntax), plus linkages between them.
- What makes a piece of linguistic structure more rule-like is the presence of variables. Lexical items run the gamut from being fully specified (like *cat*) to being made up entirely of variables (like VP).
- Variables in either morphology or syntax can be marked as closed (partially productive) or as open (fully productive).
- The open-endedness of language is a consequence of combining pieces of structure that have open variables, using the domain-general operation of unification.

Looking to further horizons, this raises the question of whether domains of mental representation other than language can be characterized similarly. Candidate domains might include visual/spatial cognition, the formulation of action, and music. For instance, in music: How does one store one's knowledge of a melody? How is this knowledge supported by structural schemas at different levels of specificity, such as common metrical, melodic, and harmonic patterns? How does this knowledge allow one to appreciate ornamentation, elaboration, and variation on the melody? Among the structural schemas (if they exist), is there a counterpart of the fully/partially productive distinction? How are the structural schemas learned? And so on.

My working hypothesis is that the overall texture of different mental domains is largely the same, differing primarily in the types of units from which the domains are built and what other domains they connect with. Memory in all applicable domains is stored as pieces of structure—both fully specified items (like words) and underspecified items with variables (like rules and constructions). Open-endedness is achieved by unifying these pieces of structure. In this respect, the present argument, built on details of linguistic structure, presents a template and a challenge for cognitive science as a whole: Can this type of analysis serve to characterize all of cognition?

Acknowledgments

For the opportunity to vet many iterations of this material, I am grateful to members of the Tufts linguistics seminar: Ari Goldberg, Neil Cohn, Eva Wittenberg, Anita Peti-Stantić, Anastasia Smirnova, and Naomi Caselli. Thanks also to Jay Keyser for many enjoyable discussions, and particularly to Jenny Audring, my guide in the wilds of morphology. Finally, I owe a great debt of gratitude to Dan Dennett for an unending supply of intellectual adventures over the past 30 years.

Works Cited

Albright, A., & Hayes, B. (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition, 90*, 119–161.

Anderson, S. (1992). *A-morphous morphology*. Cambridge, UK: Cambridge University Press.

Aronoff, M. (1976). *Word formation in generative grammar*. Cambridge, MA: MIT Press.

Baayen, R. H. (2007). Storage and computation in the mental lexicon. In G. Jarema & G. Libben (Eds.), *The mental lexicon: Core perspectives* (pp. 81–104). Amsterdam, Holland: Elsevier.

Berko, J. (1958). The child's learning of English morphology. *Word*, *14*, 150–177.

Blevins, J. (2004). *Evolutionary phonology: The emergence of sound patterns*. Cambridge, UK: Cambridge University Press.

Blevins, J. (2006). Word-based morphology. *Journal of Linguistics*, *42*, 531–573.

Bochner, H. (1993). *Simplicity in generative morphology*. Berlin, Germany: Mouton de Gruyter.

Bod, R. (2006). Exemplar-based syntax: How to get productivity from examples. *The Linguistic Review*, *Vol. 23*, *Special Issue on Exemplar-Based Models of Language*

Booij, G. (2010). *Construction morphology*. Oxford, UK: Oxford University Press.

Bresnan, J. (Ed.). (1982). *The mental representation of grammatical relations*. Cambridge, MA: MIT Press.

Bresnan, J. (2001). *Lexical-functional syntax*. Oxford, UK: Blackwell.

Bybee, J. (2010). *Language, usage, and cognition*. Cambridge, England: Cambridge University Press.

Bybee, J., & McClelland, J. (2005). Alternatives to the combinatorial paradigm of linguistic theory based on domain general principles of human cognition. *Linguistic Review*, *22*, 381–410.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Chomsky, N. (1970). Remarks on nominalization. In R. Jacobs & P. Rosenbaum (Eds.), *Readings in English transformational grammar* (pp. 184–221). Waltham, MA: Ginn. Also in Chomsky, *Studies on semantics in generative grammar,* The Hague: Mouton, 1972.

Chomsky, N. (1986). *Knowledge of language: Its nature, origin, and use*. New York, NY: Praeger.

Chomsky, N. (1995). *The minimalist program*. Cambridge, MA: MIT Press.

Chomsky, N. (2002). *On nature and language*. Cambridge, England: Cambridge University Press.

Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York: Harper & Row.

Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, *23*, 157–205.

Culicover, P. (1999). *Syntactic nuts: Hard cases in syntax*. Oxford, UK: Oxford University Press.

Culicover, P., & Jackendoff, R. (2005). *Simpler syntax*. Oxford, UK: Oxford University Press.

Culicover, P., & Nowak, A. (2003). *Dynamical grammar*. Oxford, UK: Oxford University Press.

Dennett, D. C. (1991). *Consciousness explained*. Boston, MA: Little, Brown.

Dennett, D. C. (2013). *Intuition pumps and other tools for thinking*. New York, NY: Norton.

Elman, J. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.

Fodor, J. D. (1998). Parsing to learn. *Journal of Psycholinguistic Research*, *27*, 339–374.

Frank, S. L., Bod, R., & Christiansen, M. H. (2012). How hierarchical is language use? *Proceedings of the Royal Society B: Biological Sciences*, *279*, 4522–4531.

Goldberg, A. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago, IL: University of Chicago Press.

Goldberg, A. (2006). *Constructions at work: The nature of generalization in language*. Oxford, UK: Oxford University Press.

Hale, K., & Keyser, S. J. (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. J. Keyser (Eds.), *The view from building 20* (pp. 53–109). Cambridge, MA: MIT Press.

Hale, K., & Keyser, S. J. (2002). *Prolegomenon to a theory of argument structure*. Cambridge, MA: MIT Press.

Halle, M. (1973). Prolegomena to a theory of word-formation. *Linguistic Inquiry*, *4*, 3–16.

Halle, M., & Marantz, A. (1993). Distributed morphology. In K. Hale & S. J. Keyser (Eds.), *The view from building 20*, 111–176. Cambridge, MA: MIT Press.

Harley, H. (2014). On the identity of roots. *Theoretical Linguistics*, *40*, 225–276.

Hoffman, T., & Trousdale, G. (2013). *The Oxford handbook of construction grammar*. Oxford, UK: Oxford University Press.

Jackendoff, R. (1975). Morphological and semantic regularities in the lexicon. *Language*, *51*, 639–671.

Jackendoff, R. (2002). *Foundations of language*. Oxford, UK: Oxford University Press.

Jackendoff, R. (2007). A Parallel Architecture perspective on language processing. *Brain Research*, *1146*, 2–22.

Jackendoff, R. (2010). *Meaning and the lexicon*. Oxford, UK: Oxford University Press.

Jackendoff, R. (2011). What is the human language faculty? Two views. *Language*, *87*, 586–624.

Jackendoff, R., & Audring, J. (2016). Morphological schemas: Theoretical and psycholinguistic issues. *The Mental Lexicon*, *11*, 467–493.

Jackendoff, R., & Audring, J. (forthcoming). Relational morphology in the parallel architecture. In J. Audring & F. Masini (Eds.), *Oxford handbook of morphological theory*. Oxford, UK: Oxford University Press.

Joshi, A. (1987). An introduction to Tree-Adjoining Grammars. In A. Manaster-Ramer (Ed.), *Mathematics of language* (pp. 87–114). Amsterdam, North Holland: John Benjamins.

Kay, P. (2013). The limits of (Construction) grammar. In T. Hoffman & G. Trousdale (Eds.), *The Oxford handbook of construction grammar* (pp. 32–48). Oxford, UK: Oxford University Press.

Lakoff, G. (1970). *Irregularity in syntax*. New York, NY: Holt, Rinehart & Winston.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.

Langacker, R. (1987). *Foundations of cognitive grammar*. Stanford, CA: Stanford University Press.

Lees, R. B. (1960). *The grammar of English nominalizations*. The Hague, The Netherlands: Mouton.

Lieber, R. (1992). *Deconstructing morphology: Word formation in syntactic theory*. Chicago, IL: University of Chicago Press.

Marcus, G. (1998). Rethinking eliminative connectionism. *Cognitive Psychology*, *37*, 243–282.

Marcus, G. (2001). *The algebraic mind*. Cambridge, MA: MIT Press.

Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.

Murphy, G. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Pierrehumbert, J. (2002). Word-specific phonetics. In C. Gussenhoven & N. Warner (Eds.), *Laboratory phonology VII* (pp. 101–139). Berlin, Germany: Mouton de Gruyter.
Pinker, S. (1999). *Words and rules*. New York, NY: Basic Books.
Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, *26*, 195–267.
Pollard, C., & Sag, I. (1987). *Information-based syntax and semantics*. Stanford, CA: Center for the Study of Language and Information.
Pollard, C., & Sag, I. (1994). *Head-driven phrase structure grammar*. Chicago, IL: University of Chicago Press.
Prince, A., & Smolensky, P. (2004): *Optimality theory: Constraint interaction in generative grammar*. Oxford, UK: Blackwell.
Pullum, G. (2013). The central question in comparative syntactic metatheory. *Mind and Language*, *28*, 492–521.
Rumelhart, D., & McClelland, J. (1986). On learning the past tense of English verbs. In J. McClelland, D. Rumelhart, & the PDP Research Group, *Parallel distributed processing* (Vol. 2, pp. 216–271). Cambridge, MA: MIT Press.
Sadock, J. (1991). *Autolexical syntax*. Chicago, IL: University of Chicago Press.
Searle, J. (1995). *The construction of social reality*. New York, NY: Free Press.
Seidenberg, M., & McDonald, M. (1999). A probabilistic constraints approach to language acquisition and processing. *Cognitive Science*, *23*, 569–588.
Shieber, S. (1986). *An introduction to unification-based approaches to grammar*. Stanford, CA: Center for the Study of Language and Information.
Siddiqi, D. (forthcoming). Distributed morphology. In J. Audring & F. Masini (Eds.), *Oxford handbook of morphological theory*. Oxford, UK: Oxford University Press.

AQ: Please update "Siddiqi, D. (forthcoming)" if possible in proof stage (p. 49)

Spencer, A. (2013). *Lexical relatedness*. Oxford, UK: Oxford University Press.
Stump, G. (1990). *Inflectional morphology: A theory of paradigm structure*. Cambridge, England: Cambridge University Press.
Taylor, J. (2012). *The mental corpus*. Oxford, UK: Oxford University Press.
Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.
Ullman, M. T. (2004). Contributions of memory circuits to language: The declarative/procedural model, *Cognition*, *92*, 231–270.
Wasow, T. (1977). Transformations and the lexicon. In P. Culicover, T. Wasow, & A. Akmajian (Eds.), *Formal syntax* (pp. 327–360). New York, NY: Academic Press.