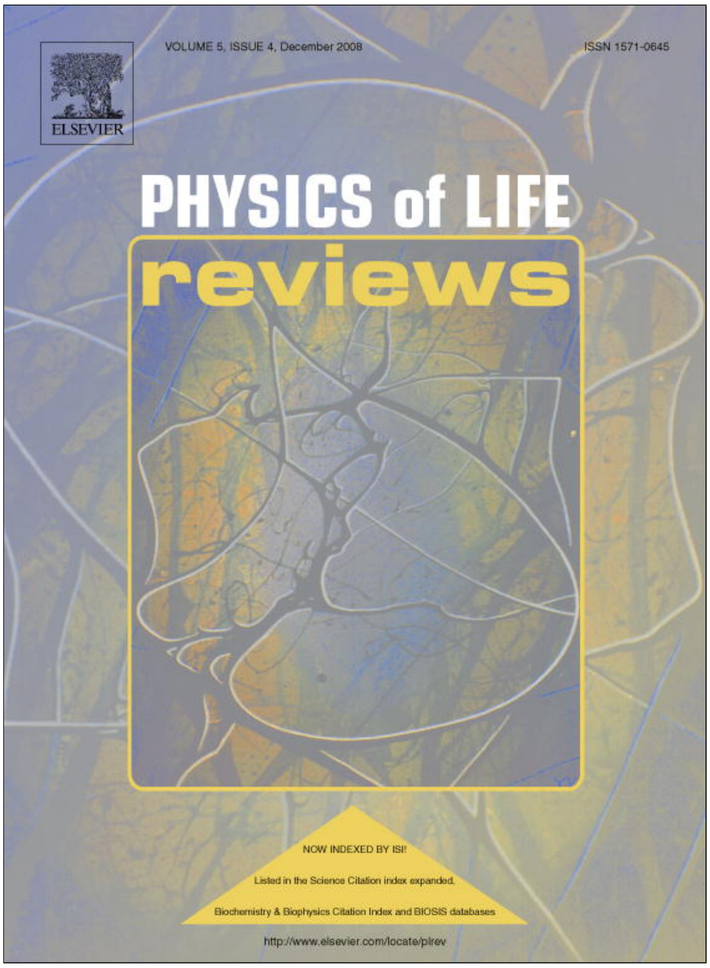
VOLUME 5, ISSUE 4, December 2008
ISSN 1571-0645
ELSEVIER
PHYSICS of LIFE
reviews
NOW INDEXED BY ISI!
Listed in the Science Citation index expanded,
Biochemistry & Biophysics Citation Index and BIOSIS databases
http://www.elsevier.com/locate/plrev

Available online at www.sciencedirect.com
ScienceDirect
Physics of Life Reviews 5 (2008) 207–224
PHYSICS of LIFE reviews
www.elsevier.com/locate/plrev

Review

Still a bridge too far? Biolinguistic questions for grounding language on brains

Massimo Piattelli-Palmarini, Juan Uriagereka *

Cognitive Science, Linguistics, and Psychology at the University of Arizona, Department of Linguistics, University of Maryland, 1401 Marie Mount Hall, College Park, United States

Received 16 May 2008; received in revised form 28 July 2008; accepted 29 July 2008

Available online 9 August 2008

Communicated by L. Perlovsky

Abstract

In this paper we discuss how Fibonacci growth patterns are apparent in the structure of human language. We moreover show how familiar dynamics yielding these sorts of patterns in nature may be taken to apply, at some level of abstraction, for the human faculty of language. The overall picture casts doubts on any simplistic treatment of language behavior, of the sort stemming from classical behaviorism in psychology (which is again popular in contemporary computational models). Instead, it appears to be more profitable to study language as a complex dynamic system, emerging in human brains for physical reasons which are yet to be fully comprehended, but which in the end disfavor any dualistic approach to the study of mind in general, and the human mind in particular.

Keywords: Language; Syntax; Fibonacci; Structure; Complexity

Contents

* Corresponding author.
E-mail address: juan@umd.edu (J. Uriagereka).

1571-0645/$ – see front matter
doi:10.1016/j.plrev.2008.07.002

1. Introduction

The scientific study of language has made relatively little progress as the description of a behavior; instead, most insights have come from focusing it as the analysis of *structures*, the majority of which are in some sense hidden and quite abstract. From that perspective, the issue is how human minds manage to have access to said structures, in many instances without learning in a literal sense playing any significant part in the process. This is to the point that humans seem to literally *grow* into the acquisition of central linguistic patterns, up to a critical age when the process virtually ends. If this is the case, we need to study growth patterns in nature, and see how they apply to language.

Early linguistic investigations led to an elegant and compact formulation of the problem of what may be seen as 'natural knowledge'. Many so-called mental modules thus described already resist any simple-minded emergence in classical adaptationist terms. But the challenge became all the more extreme with the advent of the Minimalist Program ([9] and elsewhere), with its suggestion that linguistic structuring obeys some kind of optimum (at any rate, an optimal compromise) at the interface between the interpretation and the externalization of language. Interestingly, optima turn out to be quite important when considering growth patterns in general.

We find it natural to conceive that program in terms of broadly construed physics, a dimension that is alien to standard accounts in terms of natural selection. The present paper argues for this approach, illustrating it with the presence of Fibonacci growth patterns in language, which we take to be a signal case. (Space limitations prevent us from going into several other optimal solutions in this realm.) What could it mean to ground such properties on a human brain? We wager that they are the result of the brain's very physical conditions, in some sense to be fully investigated and understood. Structural properties of linguistic behaviors should then, more generally, follow as a deductive consequence of quite abstract and all-important brain micro-properties.

2. A primer on Fibonacci

Eyes, limbs, digestive systems, all sorts of organ networks rest on genetic and proteomic interactions. But these are ultimately physico-chemical as such, and thus—quite aside from the concrete intricacies that depend on which particular dependencies are turned on or off in given interactive cascades—the 'genetic soup' must satisfy even more basic, and really non-genetic, constraints imposed by underlying dynamic conditions. Consider for example optimal growth structures in organisms, carrying them from a single cell to a complex network thereof, or from an axial state to a skeleton or nuanced limb. A priori, such growing patterns could have been vastly different from species to species. Indeed, tetrapods are not the same as mollusks, or the bilateral body plan of an endoskeletal or exoskeletal organism is patently specific in each case, and differs even more from the radial symmetry witnessed in non-skeletal organisms. However, at a sufficient level of abstraction regularities emerge below these differences. Thus, for instance, all of those groups mentioned present Fibonacci (henceforth, F) patterns: either a number of features falling into the series 1, 1, 2, 3, 5, 8, ... or logarithmic growth based on the limit of the ratio between successive terms in the F series (1.618033988..., the so-called golden expression or ϕ). The patterns witnessed go from the organization of skin pores (in tetrapods) to the way in which shells grow (in mollusks), extending beyond the animal kingdom into the overall body-plan of plants. What could it mean, for whatever problem might be common to such different organisms, to be solved in specific F terms by natural selection acting on randomly generated forms? The search space to converge in the witnessed solution would be so huge that a random walk through it would take more generations than could exist—so this cannot be the locus for an explanation to the remarkable convergence that interests us here.

The patterns we are after present a characteristic optimality to them, which might suggest that it is, after all, a result of natural selection. However, the sort of optimality that natural selection could yield is tied-up to contextual specifications of a functional sort. A structure, in these terms, is optimal *for a function*, and therefore it wins the evolutionary race simply because alternative structures are less fit for that function in its context. This is quite different, in principle, from the optimality we seek to understand. To be concrete about it in terms of F patterns, bear this in mind: We know that the golden expression is an irrational number requiring large denominators to approximate rationally—it is a 'very irrational' number. This is significant for 'mini-max' packaging problems, seeking to optimize the maximum number of regular features within a minimal amount of space. Spacing such elements at an angle apart that is maximally irrational entails that they overlap as little as possible. Moreover, this solution can be expressed in aggregative terms, between successive elements in an additive series; so it is ideal for biological entities dealing with tissue growth. This does not seem to be related to the function that the grown tissue happens to have: the solution is mathematically the same whether one is dealing with skeletal frameworks, metabolic regulation, photo-synthesis, ventilation or anything else. So just as the justified contextual optimality of a given structure may argue for its adaptive value, to the extent that we find structures that can be characterized as optimal *irrespective of any functional correlate*, the opposite conclusion should be driven: their optimality must obey a different cause, not bio-genetic interactions winnowed by natural selection. In our view, the more that optimality in any biological network can be separated from its putative function, the more we should suspect that the process behind the abstract form follow from physico-chemical invariants.

Traditional evolutionary accounts of a Neo-Darwinian sort have never taken into consideration the fact that what may be thought of as 'morpho-spaces' in the biological realm are scarcely populated with respect to theoretically possible forms [19,36]. Rhetoric aside—the familiar mantra that *Form always follows Function*—canonical adaptationist accounts find difficulty explaining the actually observed forms on the basis of their own logic alone. It is just not obvious that any one concrete function can be ascribed to such forms, even if descriptively expressed in such generic terms as 'packaging': ultimately, a tree growing branches in time or a developing animal generating skin pores are not literally packing anything (a pore is a hole and a branch the outgrowth of a gem, and these features are not being packed within the organism). Moreover, if there is a discernible function for each relevant structure, it is entirely disparate, as we saw above for F patterns. In contrast, a physicalist approach seems relatively straightforward in this instance: F patterns of the relevant sort are patently found in nature—in systems that range from cosmology (spiral galaxies) to hydrodynamics (whirlpools)—and have been carefully created in the lab by Douady and Couder in an inorganic system of droplets falling on a flat surface [15].

These researchers used a magnetized ferro-fluid, which they slowly dropped on the center of a rotating paraffin plate, to witness the path followed by each successive drop as the plate filled with these elements. A way to abstractly visualize the Douady and Couder experiment, simplifying things a bit, is by realizing that it involves two opposing forces. On the one hand, we have a local repulsion (here, electro-magnetism) that keeps each drop separate from one another. But at the same time, we also have a more global attractive force (here, gravity) that keeps the system in place. The issue is how these opposing forces balance each other out, such that the largest number of repelling droplets can fit within the plate at any given time, as they fall onto it. It turns out that ϕ achieves this dynamic equilibrium. Descriptively, if the goal is to fit the largest number of drops, we want a division of the packing space that is not rational, for any rational division n would only work for sectioning the 360° of the circular space in n portions, which is only useful for n drops. *Which* irrational division is best? The 'more irrational' the division is, the more it will fit the present purposes. It appears that either ϕ is optimal in some relevant mathematical sense (as alluded to above) or else it happens to be the irrational solution that nature has found for its dimensions—apparently regardless of the scale of the problem. It is of course interesting to ponder exactly how it is that nature reaches this solution, in physical systems generally involving opposing forces as just described; but for our purposes here it suffices to say that it does, at least in some concrete instances.

Importantly, the Douady and Couder scenario [15] was purely physical and is, thus, not a structure that emerges in Darwinian conditions. That said, 'all' that needs to be assumed for biological instances is that organisms too participate in the relevant complex dynamic systems that one encounters in nature, involving roughly repulsive and attractive forces pulling in opposite directions, at whatever level turns out to be relevant. Of course, figuring out the details of this will be difficult, for we do not expect nature to reason globally—or for that matter attempting to resolve the riddle for each new living species (from mammals to viruses, really). But it probably is to be expected that F growth patterns in living entities will reduce to more elementary dynamics. The prospect certainly seems promising for body plans, and an entire branch of bio-physics is concerned with matters of this sort. Our proposal adds structured behavior to the mix, suggesting that there too a reduction in this direction should be sought.

That approach, which seeks the interplay between opposing 'forces', could be explored at three different dimensions. (i) At the time of the evolution of a given species, the physics of the context where this large-scale event took place may have resulted in systemic interactions of the right sort, which in turn somehow ended up coded in the genome corresponding to organisms of that species (perhaps specifying the timing of successive cell divisions and the cell-to-cell adhesion thresholds); this would be a philogenetic approach. (ii) The developmental pathway of an organism, which again will be channeled by the relevant physical constraints, coordinates the growth and allometric ratios of the different parts, in this instance requiring some epigenetic interplay between the genome originated at an evolutionary scale and the proteome deployed in the developing individual; this would be an ontogenetic approach. Finally, (iii) it could even be that these dynamics actually continue to unfold throughout the organism's (adult) life, regulating physiological processes of various, presently not well understood, sorts. Indeed, it is possible that what we experience as 'behavior' has components that owe their properties to *all three* of these factors, which range from the global characteristics of a phylum or species to the behavioral patterns in an individual's existence.

We make an effort to outline those basic ideas because they are central to the core reasoning of this paper, which is concerned with what may be thought of as *the growth of form in behavior*. Form is an abstract notion, so we expect its manifestation in matters that are not obviously 'material'. That is, while form exists in the classical physical sense, if there is something to the notion 'body-plan' in the terms sketched above, we also expect form to be lurking inside the regular behaviors of living creatures emerging from those plans. That is the intuition behind rationalist psychology, which got a major boost in the mid-twentieth century from Noam Chomsky's mentalist approaches to grammar. In a nutshell, Chomsky's scientific production over more than half a century presented a long and detailed argument that the linguistic behavior that human children grow into is highly structured, and this structure cannot be summarily dismissed as a mere support for a communicative function, or some such functionalist notion. Just as we expect this for directly observable forms in biology, so too we expect general factors at play in abstract linguistic forms: surely there may be organism-internal and external natural selection at various levels or contingent interactions between many levels within the organism (which Lewontin [32] calls developmental noise); but also, crucially, ultimately physico-chemical demands roughly as discussed above. More specifically for our purposes, if we find an F-pattern in language structuring, we hardly expect natural selection to dictate that such a form be what it is, having emerged as a result of some unclear function. We find it more profitable to assume that underlying growth processes emerge because of some interplay between opposing 'forces' of nature, whatever those turn out to be.

3. F-patterns in behavior

F patterns are found in many aspects of human behavior. The most obvious is music, the classic diatonic scale being comprised of 8 notes, of which the 5th and 3rd create the chord foundation. The scale is based on the whole tone, 2 steps removed from the root tone (cf. the disposition of the 13 piano keys between any two C's—counting both of these, in the way they are used in chords—with its 8 white and 5 black keys, the latter split in groups of 3 and 2). Note also that the dominant fifth is the eighth note in the chromatic scale, the basis of musical coloring. Aside from this, F patterns can also be found in the timing of musical compositions; for instance, the climax is often placed at a point that relates the piece's length in terms of the golden proportion. (Whether any of this is done consciously by classical composers is, naturally, very hard to establish.)

F patterns are not restricted to music; it is hard to find an artistic domain where they do not show up. Even the Greek letter ϕ (pronounced *phi*) is used to denote the golden mean because of its systematic use by Phidias. Certainly the world of aesthetics—explicitly since the Renaissance, but implicitly long before—is largely based on the golden mean, or reactions to and derivations from it. This is not the place to review the vast literature on this topic, and our only concern is the obvious one for cognitive psychology: Why or How does the human mind 'gravitate' towards this proportion, making it central to its most precious creations? This is particularly interesting inasmuch as nature more generally seems to have found the proportion, as has been noted by scholars from Pythagoras to Roger Penrose, through such intellectual luminaries as Euclid, DaVinci, Kepler, Goethe, D'Arcy Thompson, Alan Turing and Noam Chomsky.

We believe we have found F patterns in language. This was first shown for the ontology of phonological syllables in [43, Chapter 6]. To see why these may be profitably described in F terms, observe that, simplifying matters to just consonant (C) and vowel (V) combinations (i.e., disregarding consonant clusters as in *pranks*, or concentrating on consonantal anchors at either the nucleus or coda of any given syllable) syllabic patterns in the world's languages are as in (1)—in order of frequency:

(1) CV (and its stressed variant CVV), CVC (and its stressed variant CVVC), V and VC.

There do not seem to be any syllable templates like CVVVC or VVC, for instance, and the question is why only the basic patterns in (1) emerge, and with that (descending) frequency. In particular, as is usefully summarized in [2], many unrelated languages exhibit all of these syllabic possibilities, varying mainly in the nuanced consonant clusters or vocalic colorings they allow and which are now being set aside (e.g., the Spanish ***va*** 'goes' vs. the Dutch ***vee*** 'cattle', both instances of the broad CV(V) template; or the Spanish ***sin*** 'without' vs. the Dutch ***baard*** 'beard',

```
a.  +-   +-+   +--+   +-++-   +--++-   +--++-+   +--++--+
    1    1     +-+-   +--+-   +-+--+   +-++--+   +--++-+-
               2      +-+-+   +-++-+   +--+--+   +-+-+-+-
                      3       +--+-+   +-++-+-   +-+-+--+
                              +-+-+-   +--+-+-   +-++-++-
                              5        +-+-++-   +--+-++-
                                       +-+--+-   +-++--+-
                                       +-+-+-+   +--+--+-
                                       8         +-++-+-+
                                                 +--+-+-+
                                                 +-+-++-+
                                                 +-+--+-+
                                                 +-+-+-+-
                                                 13
..............................................................
                                                 13
                                                 -+-+-+-+
                                                 -+-++-+-
                                                 -+-+--+-
                                                 -++-+-+-
                                       8         -+--+-+-
                                       -+-+-+-   -++-++-+
                                       -+-++-+   -+--++-+
                              5        -+-+--+   -++-+--+
                              -+-+-+   -++-+-+   -+--+--+
                      3       -++-+-   -+--+-+   -+-+-++-
               2      -+-+-   -+--+-   -++-++-   -+-+-+-+
    1    1     -+-+   -++-+   -+-++-   -+--++-   -++--+-+
b.  -+   -+-   -++-   -+--+   -++--+   -++--+-   -++--++-
```

Fig. 1. F patterns emerging from the F game, for 2, 3, 4, 5, 6, 7 and 8 symbols.

both instances of the broad CV(V)C template). Then again, many unrelated languages entirely lack syllables without an onset (the V(C) template), while keeping all other forms present (e.g., we have words like *baa-l'aa-?as* 'bread' in the native American language Klamath, spoken in Oregon, but nothing equivalent to the Spanish ***al*** 'to-the' or the Dutch ***ui*** 'onion', a situation that extends to the Vietnamese Sedang and many other languages). Finally, in a few languages only templates of the CV variety (stressed or unstressed) are present. For example, the way to say 'whiskey' in Hawaiian is *We-ke-ke* and natives refer to the Arabela language, spoken in Peru, as *Ta-*p_w*e-yo-*k_w*a-ka*—without final consonants, ever).

For clarity, let's represent 'spaces' with a minus sign and 'boundaries' on those spaces with a plus sign, letting these representations freely interact according to two simple computational rules to generate symbol strings, as follows:

(2) F game.
  (i) Starting with either a + or a −,
  (ii) Go on to concatenate it to another + or a −, with one condition:
  (iii) Avoid combining identical symbols, unless they are adjacent to a different symbol.

The results of this game, starting with a space or with a boundary and adding successive symbols (one, two, three, etc. up to seven symbols in this concrete instance), are as in Fig. 1. The possible combinations as various elements are

```
a. +-    +-+   +--+    +-+ +-    +--+ +-     +--+ +-+      +--+ +--+
               +- +-   +-- +-    +- +--+     +-+ +--+      +--+ +- +-
                       +- +-+    +-+ +-+     +-- +--+      +- +- +- +-
                                 +-- +-+     +-+ +- +-     +- +- +--+
                                 +- +- +-    +-- +- +-     +-+ +-+ +-
                                             +- +-+ +-     +-- +-+ +-
                                             +- +-- +-     +-+ +-- +-
                                             +- +- +-+     +-- +-- +-
                                                           +-+ +- +-+
                                                           +-- +- +-+
                                                           +- +-+ +-+
                                                           +- +-- +-+
                                                           +- +- +- +-
                                                           - +- +- +-+
                                                           - +-+ +- +-
                                                           - +- +-- +-
                                                           -+ +- +- +-
                                                           - +-- +- +-
                                             - +- +- +-    -+ +-+ +-+
                                             - +-+ +-+     - +--+ +-+
                                             - +- +--+     -+ +- +--+
                                 - +- +-+    -+ +- +-+     - +-- +--+
                                 -+ +- +-    - +-- +-+     - +- +-+ +-
                       - +- +-   - +-- +-    -+ +-+ +-     - +- +- +-+
               - +-+   -+ +-+    - +-+ +-    - +--+ +-     -+ +-- +-+
b. -+    -+-   -+ +-   - +--+    -+ +--+     -+ +-- +-     -+ +--+ +-
```

Fig. 2. Patterns emerging from adding linguistic conditions on the F game.

added yield different arrays of spaces and boundaries, and it is easy to see that the number of possible combinations as the added elements grow falls within the Fibonacci series.

It should be emphasized that there is nothing particularly surprising about this: the reason F patterns emerge here is because of the way in which the F game in (2) is set up. The next step in our reasoning, however, is rather more interesting. Suppose we adapt these patterns (generated over abstract spaces and boundaries thereof) to linguistic conditions—which is roughly the equivalent of having an F pattern in nature emerge specifically for photosynthesis *or* ventilation purposes, etc. So concretely for our linguistic purposes:

(3) Linguistic conditions
  (i) Nucleus constraint: Look for a maximal space. Then,
  (ii) Onset constraint: Try to assign an onset boundary to that space. Then,
  (iii) Coda constraint: Try to assign a coda boundary to that space.

(3) is an optimization algorithm, trying to make bounded spaces as large as possible (3i), and as delimited as possible (the terms *onset* and *coda* are the standard way to designate the beginning and end of a syllable, if they are consonants). We return in Section 7 to a possible rationale for the Onset and Coda constraints, in terms that go beyond phonological theory (these conditions may now be conceived in operative terms, because their postulation in that particular order yields the right results). All of this has the consequences in Fig. 2.

Thus observe how, first, the algorithm attempts to find maximal spaces (combinations of '−' elements); next it attempts to delimit that maximal space in terms of an onset (if this is possible); finally the algorithm tries to find a coda for the delimited spaces. In a few circumstances, as it turns out, the remaining space is a single '−' (not a maximal space), and in fact without either an onset or a coda—but this is a relatively rare circumstance.

Importantly, only six groupings emerge from applying the linguistic conditions in (3) to the F game in (2). Readers can try further combinations of successive (eight, nine, ten, etc.) symbols under these conditions. Then further combinations within the F series will emerge (twenty one, thirty four, etc.), but when the linguistic conditions are applied to these new objects, no more further combinations will arise. This is the entire set:

(4) a. $+-$ b. $+--$ c. $+-+$ d. $+--+$ e. $-$ f. $-+$

Now if we replace the minus representation with a vowel symbol and the plus with a consonant, these patterns correspond to the ontology in (1). Moreover, observe the number of occurrences for each type within the strings above (symbols are replaced as mentioned):

(5) a. (a) grouping: CV 37 (times), CVC 21, CVV 11, CVVC 10, V 0, VC 0.
(b) grouping: CV 37 (times), CVC 17, CVV 7, CVVC 8; V 19, VC 13.

Readers can check that the (a) grouping is generated by starting the game with a boundary—which can be rationalized as a punctuated element in an open space—while the (b) grouping emerges from starting the game with a space—which can be rationalized as an in-principle boundless topology. In either instance, after applying the linguistic conditions in (3), the CV pattern emerges 37 times, the maximum. At the opposite extreme we have the V(C) pattern, which does not emerge in the (a) grouping, and does moderately in the (b) grouping (the VC pattern is the least common, emerging only 13 times). In between is the CVC pattern. This correlates with the frequency order in (1) (understanding that a sequence of two vowels corresponds to a stressed syllable; e.g., CVV is the stressed variant of CV).

Perhaps it should be emphasized that the languages where the facts in (1) obtain can be unrelated, historically or geographically—so none of these generalizations seem attributable to the vagaries of history. Even more significantly, Perlmutter [35] argued that syllabic distributions of the sort alluded to obtain in signed languages too (with no sound), where hand movements are the equivalent of open spaces, and hand positions of boundaries thereof. We have no space to elaborate on this matter, but apparently certain impossible combinations in signed language are the equivalent of unacceptable combinations in syllabic structure—for instance, the fact that the syllable, say, *baar* is grammatically possible in English (quite irrespective of whether it actually shows up in any give word of the language), while would-be syllables like **aar* (with no onset) or **rba* (with that ordering of consonants, from the more to the least sonorant preceding the vowel) are not. Remarkably, abstractly similar limitations, this time involving hand movements and positions, obtain for signed languages too. If this is correct, and the reason why such a limitation emerges is not a low-level demand associated to the speech organs, then the ontology of syllables cannot be blamed on such articulatory or perceptual facts of the oral sort, and should instead be understood in terms as abstract as just presented.

4. Generalizing F conditions: the base

Syllables can emerge as properties of two factors 'pulling' the system in opposite directions: 'repulsion' forces, of the sort in (2), that generate F patterns more generally, and more specific 'gluing' forces, as in (3), that result in discrete units of various shapes. In other words, a syllable is a mini-max compromise, the 'max' aspect being determined by general F conditions, and the 'mini' one by the linguistic specificities that (3) dictates.

As Uriagereka [43, Chapter 6] indicated, an 'equilibrium' of that sort should emerge elsewhere in language. This is perhaps not surprising in the case of poetry, as Duckworth [16] had already argued that Virgil made systematic use of F patterns, a claim often made for other poets. But matters go well beyond literary exercises. The most rigorous demonstration of F patterns in *natural* (i.e. non-poetic) language is due to Idsardi [25], for the study of metrical feet. A metrical foot is a unit of prosody, determining the relation between beats and slacks in a given sentence, which **yields** a characte**ris**tic **rhyth**mic **pat**tern—here: ba-bum-babababa-bum-ba-bum-ba-bum-ba. As it turns out, such patterns are flexible within certain limits, with secondary beats emerging in sequences involving several slacks (e.g., a cha-*rac*-te-**ris-**tic), this being a function of neighboring beats. Thus compare: **yields** a cha-*rac*-te-**ris**tic **pat**tern vs. **em**-pha-*si*-zes a *cha*-rac-te**ris**tic **pat**tern. Idsardi proves how the number of such (different) parsings into metrical feet of a string of length n feet is the Fibonacci number $Fib(2n)$.

Most striking is, however, the fact that F patterns generalize beyond phonology, as established by Boeckx et al., [7], Soschen [39] and Medeiros [34]. We can exemplify the matter with the latter work, as the most comprehensive. Observe a typical syntactic 'molecule' [9] as in (6). An object of this sort is called an XP (P for Phrase, X as a free variable over syntactic categories) because it applies regardless of the lexical category of the 'head' (noun, verb, adjective, etc.)—all the English examples in (7), in a sort of schema that is expanded across languages:

(6)

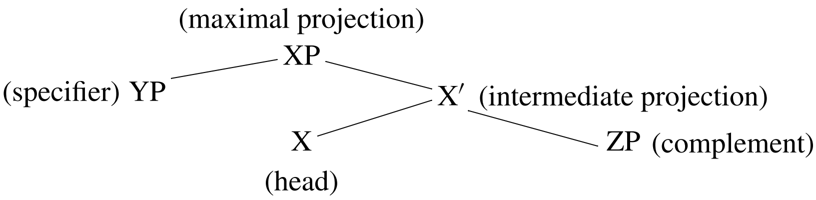

(7) [$_{\text{TP}}$[$_{\text{DP}}$ Brutus] [$_{\text{T}'}$[$_{\text{T}}$ shall] [$_{\text{V}'}$[$_{\text{V}}$ envy] [$_{\text{DP}}$ Caesar]]]]
Brutus shall envy Caesar.

a. [$_{\text{DP}}$[$_{\text{DP}}$ Brutus] [$_{\text{D}'}$[$_{\text{D}}$'s] [$_{\text{N}'}$[$_{\text{N}}$ envy] [$_{\text{P}'}$[$_{\text{P}}$ of] [$_{\text{DP}}$ Caesar]]]]]
Brutus's envy of Caesar

b. [$_{\text{A}'}$[$_{\text{A}}$ Envious] [$_{\text{P}'}$[$_{\text{P}}$ of] [$_{\text{DP}}$ Caesar]]]
Envious of Caesar

Just as, above, we disregarded secondary phonemes in consonant clusters, here too we are disregarding adverbial, adjectival and otherwise optional modifiers, which are customarily assumed to be generated 'at a separate dimension', thus not affecting the core structure in (6). Small differences still exist among the various examples above (e.g., names like *Brutus* achieve their maximal status as DPs without taking any dependent phrases, unlike nouns like *envy*; complements to such nouns require a dummy preposition *of*, unlike corresponding complements to corresponding verbs like *to envy*), but the overall pattern is known to hold quite regularly, applying even to more abstract syntactic structures (the so-called functional projections, D for Determiner, T for Tense, and all the way up to the maximal structure that an embedding verb like *think* takes as complement, the CP or Complementizer Phrase). Crucially, an XP as in (6) can be, in its turn, the ZP or YP of yet a larger structure that replicates—for example, the VP projection in (7a) is itself the complement ZP of a Tense projection, both obeying the general schema under discussion. This recursive process has no known limits in syntax, though other factors (memory, attention, death...) will de facto set some exogenous limitations—as is the case in any natural growth system in the known universe. This standard syntactic 'molecule' suggests a host of interesting additional considerations.

What follows next may appear cabalistic, but it is a simple extension of the fact that (6) is seriously assumed like the basic unit of language, and it is, as we see momentarily, an F unit. It is a well-known property of F patterns that, as they recur, they present a basic self-similarity, exhibiting the same F balance at all scales. The issue that we will now pursue is whether the structure in (6) (and its recursive identical expansion) is just a fact that empirical research has unearthed, or it can be explained on the basis of first principles: optimal morpho-dynamic criteria of the sort presented earlier. To pursue the latter route, we need to present some concrete instances.

5. Generalizing F conditions: the induction

Linguists call XP a *projection* of X. In transitive structures of the sort of *Romans hunted deer*, for example, by the present classification *deer* is a complement, while *Romans* is a specifier (recall the analogous (7a)). The association between X and its *complement* ZP is more direct than that between X (projected to X′) and its *specifier* YP, which can be shown through the fact that the complement may undergo 'noun-incorporation' (to yield expressions like *Romans deer-hunted*), while the same is not true about the corresponding specifier (cf. the parallel, but impossible: **Roman-hunted deer*). Now, it should be easy to see that the elegant, well-established, object in (6) exhibits F conditions: 1 molecule, 1 projection (X's), 2 dependents (YP, ZP), 3 terminals (YP, X, ZP), 5 total categories (XP, YP, X′, Z, ZP). This may seem like an artificial coding by the theorist, but notice that there is no natural grouping of four or six elements in (6), and of course 4 and 6 are not F-numbers. In any case, things get more interesting.

Medeiros invites us to examine optimal expansions of (6), as in (8) (where labeling is purely arbitrary). The expansion below is optimal in that the molecule in (6) is deployed systematically *in all branchings*, although this is not necessary in linguistic representations; we saw that already in (3), where not all dependents of X in the XP projection are always present (e.g., in (7c) there is no YP specifier).

(8)

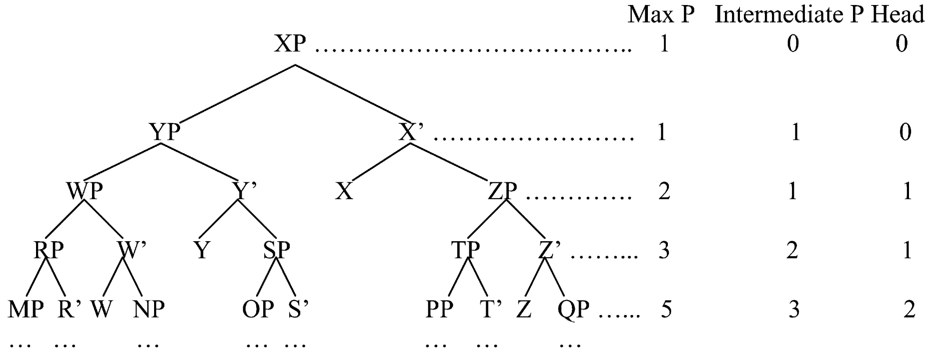

The unlimited 'tree' space made available by the unit structure in (6) contains, at each successive (full) expansion, successive F numbers of maximal projections, intermediate projections, and heads—the units of linguistic analysis. The reader will have noticed that at most two branches depart from each node. This is also a central property of all human languages, called 'binary branching', attested empirically since Kayne [28]. In a nutshell, syntactic dependencies involve precisely two elements (e.g., in this very sentence: [two + elements], [precisely + two-elements], [involve + precisely-two-elements], [syntactic + dependencies], [syntactic-dependencies + involve-precisely-two-elements]), although logically speaking any number of such dependencies could have been possible (e.g. it could have been that the right 'constituency' for this sentence were [syntactic-dependencies + involve + precisely-two-elements], but this is never witnessed in any of the world's languages studied thus far—and see (21) below). This condition is deduced from first principles of linguistic structuring and linearization by Kayne [29]. Moreover, Chomsky [10] has suggested that, since 2 is the minimum number of different lexical items that the syntactic engine can assemble (technically, *Merge*) into a larger unit, this must also be the maximum number. Within the Minimalist Program, everything that is necessary is also held to be sufficient, a point that we return to below.

Medeiros also shows how F terminal string lengths are the first to force deeper phrasal embeddings in terms of the basic molecule in (6). The issue is to determine the least embedded 'tree' possible for given terminal elements, respecting the basic (e.g., binary branching) conditions in (6). For example, with two or three terminals, only one tree is necessary (and in fact possible); namely, (9a) for two elements and (9b) for three. In contrast, for four terminals two binary trees are possible, (9c) and (9d) (labeling is irrelevant in this instance, and is thus ignored; when determinable, heads are underlined; the English examples are mere illustrations which could have been from any language):

(9)

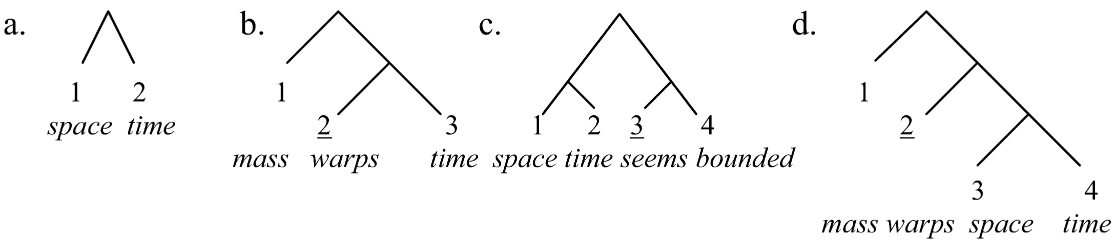

The point is that four terminals (a number not in the F sequence) does not force *a third level of embedding* as in (9d); it can sustain a representation with two such levels, as in (9c). The next number of terminals that forces a further level of embedding under these conditions is five (a number in the F sequence). The same happens with eight, thirteen, twenty one (etc.) terminal elements. Thus the following are optimal tree representations for five and eight terminals, demanding three and four levels of embedding, respectively (once again, examples in English are for mere illustration purposes):

(10) a.

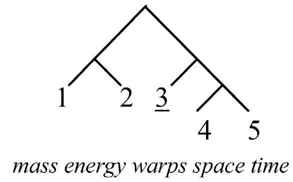

b.

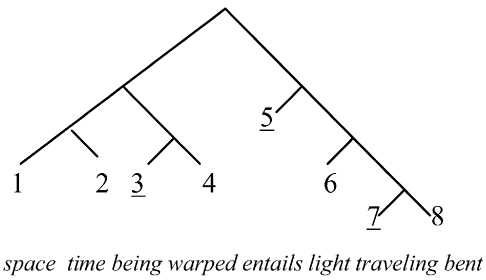

Strange symmetries like these are simple extensions of the fact that (6) is assumed as the basic F unit, and in the considerations above it is maximally expanded. *The* central property of an F pattern—as central as its overall simplicity—is that as it grows it remains the same F pattern, thus it is self-similar. To see this, consider the fact that a simple rewrite system like $0 \to 1$, $1 \to 0\ 1$, when applying maximally at each derivational line in the computation (all rewritable symbols rewrite), generates an F pattern:

(11)

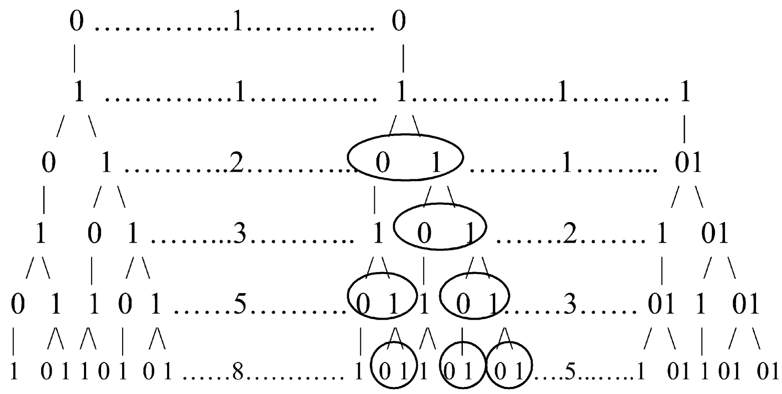

Now, observe that the first 'tree' in (11) presents an obvious regularity within it (a consequence of the rule application), which can be highlighted by circling the repeating unit ⟨0, 1⟩. If, in turn, we compile that unit as a symbol *01*, for a 'second order' graph, it is easy to see that the new object obtained is still an F pattern, the last 'tree' in (11) (generable via the rules $1 \to$ *01*, *01* $\to 1$, *01*). And within that new object we could perform the same compilation, with the unit ⟨1, *01*⟩, to form the 'third order' symbol *101*, and so on. The ensuing object is clearly self-similar, just as the maximal expansion of (6) in (7) preserves the F balance in the latter. So the real issue is why (6) should obtain in language; *if* it does, the fact that it continues to remain within the F pattern as it is maximally expands is no more (or less) surprising than the result in (11).

A related question is how (or if) the F pattern in (6) relates to the syllabic conditions discussed in Section 2 (and whether one of those conditions reduce to the other, or an even deeper unification is relevant). That such a connection should exist is explicitly argued by Carstairs-McCarthy [8], who sees phrasal structure as an 'exaptation' (in evolution) of earlier syllabification conditions. The idea for this sort of correlation is not new: it has been defended in synchronic studies as far back as Kaye et al. [27], and it is certainly sympathetic to intuitions that can be traced to Saussure [40], if not before. But it remains to be seen why such a structural translation—ultimately going from the realm of sound to that of structured meaning—is part of the fabric of language. We turn to this matter in the next two sections.

6. The cyclic character of linguistic computation

One of the oldest ideas in generative grammar is the notion that syntactic computation proceeds *in cycles*. The syntactic movement of one element in the sentence to a different (structurally higher) position in the same sentence can occur at what appears to be a rather large distance, but this never happens 'in one fell swoop'. It does, instead, via a series of strictly local movements, one after the other, each like the previous—whence the term 'cyclic movement'. For example, in a long-distance question like *who do you think Martin believes Susan likes?* the operator *who* displaces through the periphery of every single embedded clause, until it reaches its final destination (called its *scope*). The italicized (unpronounced) copy of the boldfaced **who** indicates the displacement path:

(12) [$_{\text{CP}}$ **Who** do you think [$_{\text{CP}}$ *who* Martin believes [$_{\text{CP}}$ *who* Susan likes *who*]]]

The evidence for a representation along the lines in (12) is very robust. One argument comes from child language [14]. Many toddlers acquiring English go through a stage where they utter expressions like (13) (sic):

(13) What do you think what Cookie Monster eats?

(13) may be a version of (12) where relevant 'copies' of *what*, left in the edge of every clause that this element visits in its way to its scope site, are activated—contrary to what happens in adult English. The basic state of affairs in (13) obtains in many other world languages. Take for instance the German (14), from [30]:

(14) Wen denkst du, wen Martin meint, wen Susan magt?
who-acc think you who-acc Martin believes who-acc Susan likes
'Who do you think Martin believes Susan likes?'

Many (non-standard) German dialects allow for activation of the moved Wh-word at every displacement site, but the original logical position (here, object of the verb *magt* 'likes'). This is a standard option in languages from Romani to Hindi.

Moreover, consider these English sentences:

(15) a. Which picture of herself do you think Martin believes Susan likes?
b. Which picture of himself do you think Martin believes Susan likes?
c. Which picture of yourself do you think Martin believes Susan likes?

These questions involving a reflexive are all possible, despite the fact that, normally, reflexives require an antecedent (here boldfaced) within its own clause. Observe:

(16) a. [$_{\text{CP}}$ You think [$_{\text{CP}}$ Martin believes [$_{\text{CP}}$ **Susan** likes a picture of *herself*?]]]
b. * [$_{\text{CP}}$ You think [$_{\text{CP}}$ **Martin** believes [$_{\text{CP}}$ Susan likes a picture of *himself*?]]]
c. * [$_{\text{CP}}$ **You** think [$_{\text{CP}}$ Martin believes [$_{\text{CP}}$ Susan likes a picture of *yourself*?]]]

So if reflexives must be clause-mates to their antecedents (hence the ungrammaticality of (16a, b)), how can the anaphors in (15) be grammatically licensed by their antecedents? The answer is trivial if, as (12) indicates, computational derivations proceed as in (17):

(17) a. ... [$_{\text{CP}}$ Susan likes which picture of self] ?
b. ... [$_{\text{CP}}$ Martin believes [$_{\text{CP}}$ [which picture of self] Susan likes wh-...]] ?
c. ... [$_{\text{CP}}$ you think [$_{\text{CP}}$ [which picture of self] Martin believes [$_{\text{CP}}$ wh-... Susan likes wh-...]]] ?

If the movement of the Wh-phrase containing the anaphor proceeds across each CP edge, the displacement path will put the anaphor in a situation of clause-mateness with respect to the grammatical antecedents (*Susan*, *Martin*, and *you*, respectively). But this elegant analysis presupposes a 'successive cyclic' operation, in line with the German (14).

The idea of the derivational cycle originated in phonology, when Chomsky et al. [12] showed how stress contours in English compounds are best computed in terms of rules that track constituent structure (this is easy to hear in the series: *spring-roll*, *California spring-roll*, *cuisinart California spring-roll*, where primary stress falls always on the outermost modifier of the compound). In turn, Bresnan [3] argued that the rules of the phonological cycle apply in the course of the syntactic cycle, the basis of present-day systems to track intonation patterns in the phonetic interface of the system (called PF for Phonetic Form). This very sentence clearly has intonational breaks, in particular before and after the adverb *clearly*, and before and after *in particular*. Similar considerations were raised by Jackendoff [26] and

Lasnik [31] with regards to the computation of semantic structures. For them, the syntactic cycle plays a role in the mapping to the LF interface (for Logical Form), in terms of such issues as the interpretation of anaphora, as we saw for the examples in (16).

Ideas along these lines have undergone various amendments and re-interpretations over the decades, and are back in fashion in terms of systems of so-called multiple Spell-out, such as the one defended in [10], involving what he calls 'phases' in a syntactic derivation—the modern equivalent of the traditional cycle. The intuition behind the notion 'phase' is that the syntactic derivations that generate a complex sentence have specific and mandatory points of closure. These partially structured chunks are 'sent' to the interfaces (PF and LF) before the sentence is completed—and once sent, they become impenetrable to further computations.

Chomsky suggests that systemic phases, aside from their various phonetic and semantic consequences, are necessary for the grammar to make its computations workable. Otherwise, inasmuch as sentences are arbitrarily complex, dealing with their information would be unfeasible. This is best seen if we consider a difficulty emphasized by Townsend and Bever [41]. Syntactic conditions are thought to be bottom-up; for example, (6) is constructed by first merging the head X to its complement YP, as ample linguistic evidence suggests—starting with the fact that, as we have seen above, Wh-movement proceeds *outwards* in the syntactic derivation. But linguistic behavior in sentence parsing proceeds serially and incrementally from before to after (the reader, for instance, started parsing this very sentence by first processing the words *while syntactic conditions*..., not the parenthetical that we are now dealing with). 'Commensurability' between these two orthogonal computations becomes possible only if we break down the structures for which the problem arises to a size that makes it not just computationally solvable, but *effectively* so, as argued by Berwick and Weinberg [1].

But although the 'Two Orthogonal Computations' model directly argues for systemic cycles, it does not tell us much about what Boeckx [6] calls the cycle's *periodicity*. Intuitively, one may need to take a break once in a while from writing, reading, or otherwise using a computer—but *how often* should that be? Empirically, there is a good understanding of what the syntactic periods are. Recasting much current work on the topic, Richards [38] suggests that the right periodicity among phases is as indicated in (18) (for P = phase and N = non-phase):

(18) $[_{P} [_{N}[_{P} [_{N}[_{P} [_{N} \ldots]]]]]]$

The idea is that successive categories in the syntactic skeleton (whatever they turn out to be: TP, CP and so on) stand in a phase/not-phase 'rhythm' with regard to one another. When Richards speaks of this phasal periodicity, he provides factual evidence, of the sort we discussed in this section, that each and every phase edge is visited in a long-distance displacement (see Boeckx [4] for detailed discussion of this sort of evidence).

Gallego [21] admits that (18) is indeed common, but suggests that in some languages things are more intriguing (see also Fortuny [20]). In particular, he suggests that just as not all languages present the same syllabic structure (e.g., all have CV syllables, but only some have CVC syllables)—and this does not mean that all languages do not participate on the same underlying conditions—so too languages may differ on the periodicity of their phases. Concretely, in some languages there are more projected materials between, for example, the left-peripheral CP and the core TP (he thinks of this extra material as the FP region in [42]). So his alternative to (18) is (19):

(19) $[_{P} [_{N} [_{N} \ldots [_{P} [_{N} \ldots [_{P} [_{N} [_{N} \ldots]]]]]]]]$

In effect Gallego suggests that the 'rhythm' Richards alluded to may be more intriguing than implied in (18): not only does a phase immediately dominate a non-phase and vice-versa (Richard's scenario), but under certain predictable conditions a non-phase can immediately dominate yet another non-phase (which entails that, in those languages, there are more intermediate categories, like FP between CP and TP).

7. Two opposing forces in language

Uriagereka [43] shows how, if we assign a + representation to phase edges and a − representation to core phase spaces (separating the N instance that Richards postulated, as in (18), from the double N instance that Gallego added to the discussion, as in (19), in terms of a − vs. double − representation), then we obtain the rhythmic units in (20a) and (20b). Moreover, Uriagereka shows how three more phasal conditions can be argued for when right edges are studied, concretely (20c), (20d) and (20e):

(20) a. +− b. +−− c. +−+ d. +−−+ e. − f. −+

This, although an observation about the conditions under which phases dominate other phases within a clausal skeleton, is of course identical to the series of syllabic patterns that we sketched in (4) in Section 3. Moreover, Uriagereka shows how (20a) is by far the most familiar pattern, with the characteristic rhythm that Richards described. (20b) is Gallego's variant, which corresponds only to what Chomsky [11] calls non-defective phases (which is to say that, in languages of the relevant sort, (20a) exists as well, albeit for defective phases associated to infinitival and similar conditions). (20e) is rare, and (20f) is the rarest. The important point is that we already know how this sort of abstract ontology can emerge as an F pattern by the rules in (2) in Section 3, under conditions of the form in (3). But what does all of this mean for the linguistic system at large?

What the previous argument suggests is that the conditions studied in Section 3, which Medeiros unearthed for syntactic representations, are possibly a direct reflection of a deeper periodicity, arising for phases more generally. We return shortly to the 'opposing forces' in this instance, but it should be clear that conditions of the form in (3) are not peculiar to phonology, and are to be understood as totally abstract in nature, broadly seeking to maximize something like stable signaling spaces and boundaries thereof. This is quite direct for the Nucleus Constraint in (3i). In turn, the other two constraints seek to determine boundaries to maximal signal spaces, in an asymmetric fashion: first the onset (or 'left') of the signaling space, then the coda (or 'right') of that very space, arguably an adjustment on an abstract computational space of a lower-level information-driven requirement. For conditions of this type to emerge, it seems immaterial whether we are dealing with phonological spaces (be they consonant/vowels or hand position/movements), syntactic spaces (phasal periodicity), or some interactive situation (cyclic stress and prosodic conditions). Whether it is at the level of phonemes, words, phrases, sentences or discourses, speech is an affair that goes from before to after, mediated through articulatory/perceptual mechanisms with little computational room for maneuver.

That said, though, what are the opposing forces, implicit in the F game in (2)? As it turns out, we have an answer there too if we take Townsend and Bever's [41] 'Two Orthogonal Computations' model at face value: as we saw already, the disparate computations of language can only hope to be measured against one another by breaking each unit into cycles that are manageable to the other. We can think of the bottom-up, syntactico-semantic computation as the equivalent of the 'repulsion' forces that we studied above for syllables: this is the associative procedure that makes the system grow, endlessly as it turns out. Now, if that were all, there would not be 'gluing' forces that break things into manageable units. Perhaps inner thought does work that way, at lightning speed; perhaps if we had telepathy we could communicate all of that directly. But alas, we do not—and then we are left with the need of a 'mini-' part for this 'mini-max' problem. This is what discretizes associative thought into squeezable and then parseable speech, yielding phasal periodicity.

If this is the right approach, even the elegant molecule in (6) is, itself, a reflex of such a mini-max compromise, as Boeckx [6] shows within more technical presuppositions. In standard bottom-up terms, the head-complement relation is what the system naturally associates without any difficulty. In contrast, the head-specifier relation forces the system to go into a separate derivational workspace, as is easily seen when considering the derivation of *the man saw a woman*, as in (21):

(21)

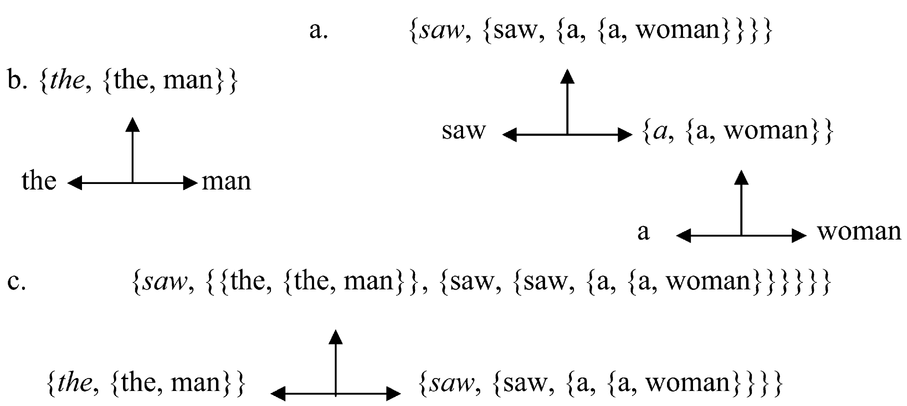

There is no bottom-up way to merge *the man* directly to *saw a woman*: we must assemble the former in a separate derivational work-space (21b), place it on a 'derivational buffer', and then assemble the results to structure still active in (21a), as in (21c). Resorting to this 'derivational buffer' is a virtual definition of 'specifier'. Then the question is how the system treats material within this buffer, vis-à-vis the rest of the structure.

Various special grammatical conditions hold of specifiers, as discussed by Uriagereka [43], summarizing a long tradition that goes back to Huang [24]. One of them is the fact that, as Chomsky argued in [11] and elsewhere, specifiers are generally not transferred to interpretation when the complement of the Probe that hosts them does. In other words, phase-wise, the edge does not go hand-in-hand with the rest of the computation, as is expected if the derivational buffer alluded to above is real: these elements occupy their own derivational dimension. In any case, the point is that, if these asymmetries do hold—and empirically, there is little doubt that they do—then the layers in the molecule in (6) follow. First-merge yields the complement space, which works in harmony with the Probing head. In contrast, the specifier is a boundary, which acts as counterpoint, if at all present, to the phase space; this is what the elsewhere-merge condition yields; the molecule is a simple reflex of the asymmetry.

These are very old ideas. As Townsend and Bever [41] note, they are revamping considerations in [13,22], the Analysis by Synthesis 'hypothesize-and-test' method in parsing. Parsing is important not just for the processing of speech as it is produced, normally by other speakers; it is crucial also for language acquisition, as a child must parse out the nuances of her language from incoming speech. As Poeppel et al. [37] argue, this general approach to the matter is extremely plausible, from both psychological and biological perspectives:

> Based on minimal sensory information, the perceptual system generates knowledge-based 'guesses' (hypotheses) about possible targets and internally synthesizes these targets. Matching procedures between the synthesized candidate targets and the input signal ultimately select the best match; in other words, the analysis is guided by internally synthesized candidate representations. [p. 1072]

Which is to say, to put it in Marantz's [33] apt terms, that 'there is no escape from syntax.' In effect, in this view of things syntax is literally *the physics of language* (and see both Hinzen [23] and Boeckx [5] for quite literal articulations of this idea), and thus guiding cognitive psychology through its pathways is pretty much the same sort of situation that life finds itself into when its emergence is guided through physical constraints. Growth, in both instances, is a matter of channeling.

If 'Analysis by Synthesis' is central to linguistic architecture, it is all we need to understand the structuring conditions in language. From the 'hypothesize-and-test' perspective, grammatically possible forms are creatively generated (plausibly in language acquisition, i.e. linguistic growth) from a first-pass analysis, stemming from the local symbolic combinations identified upon lexical analysis, and fed-back into the system for comparative accuracy. That is what effectively produces two systemic 'forces': the associative ones constructing structure, and the looping ones, which in the process of feeding information back into the system must determine characteristic periodicities. Symbolic units, whether phases, syllables or anything else, can be seen as *conservation conditions* at various time-resolution levels; their symbolic character is thus an emergent property, given system dynamics. Language architecture is based on language physics.

Structuralists famously called the characteristic pattern duality that has fascinated linguists of all times the 'double articulation' of language. In this piece, the two instances we have observed in this regard are at two apparently quite distinct levels: syllables and phases. What is important, however, is that, different though these 'articulators' surely are, in a deep sense they are also abstractly very similar. In other words, in this regard at least the double articulation turns out to be a single affair, albeit with different degrees of coarseness, possibly determined by that mysterious unit in between: the word. We call the word mysterious as did Saussure before us [40], and for reasons that relate to Fodor's [18] views on atomic lexical units. Nothing we have said here touches on what a word is, or why it exists, not to mention how it is that a human child can acquire about a dozen a day during the explosive acquisition period. However, words are, and they are in between the two F articulators studied here, syllables and phases—to the point that they may be wedges to separate these articulators: syllables go inside words, phases outside.

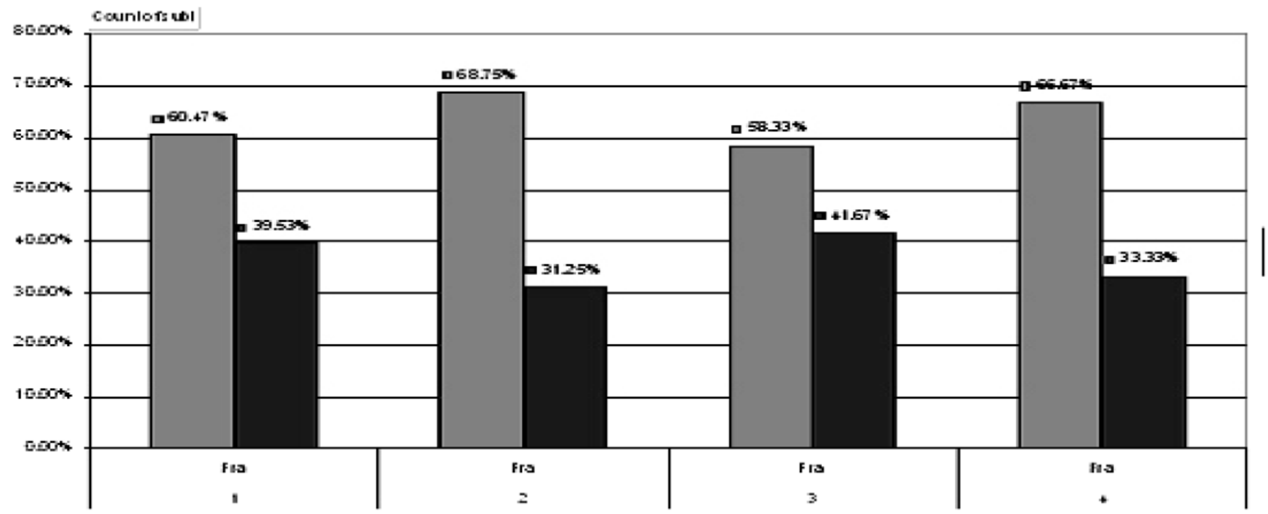

Fig. 3. For each substitution version, bars show accuracy of identifying a 10 s sample taken from the Fibonacci training set (gray) in a forced choice decision with a 10 s sample taken from a random string created with that substitution set.

8. A simple experiment

The vast, now fashionable, statistical learning literature takes a more conservative approach to these matters, uncritically assuming that behavioral systems ought to simply involve more or less clever generalizations over encountered data, which may be structured for reasons that are rarely, if ever, specified. In such a view, systematic properties of symbol strings are taken to be directly delivered by transition probabilities (technically, first order Markov processes), a hypothesis that Noam Chomsky forcefully argued against over half a century ago, at least for language. Taking (yet again) that sort of question at face value, in a recent study Fink and Saddy [17] investigated what humans do in front of pseudo-random strings of syllables, for which there are no simple statistical regularities at the level of their neighboring occurrences, one after the other. Interestingly for our purposes, these researches generated such strings via the recursive procedures that a Fibonacci grammar deploys ($0 \rightarrow 1$, $1 \rightarrow 1\ 0$, using syllables [bi] and [ba] to represent 0 and 1), which generates deterministic pseudo-random strings.

Subjects of the experiment would hear synthetically generated sequences like *ba-bi-ba-ba-bi-ba-bi-ba* (corresponding to the sequence of *1*'s and *0*'s in the bottom line of the tree in (11)). Importantly, the local transition probabilities between symbols in strings generated by these grammars are close to random, so success in learning the regular properties of these systems cannot be attributed to transition probabilities. In the experiment, these sorts of pseudo-random strings were compared to others generated as follows. First a random string was chosen (making sure it could not have been generated by the Fibonacci grammar). Then four different substitutions of symbols for string bits if the Fibonacci sort were executed, to make the strings comparable, yielding four versions of the experiment:

(i) *1* is replaced with *1 0* and *0* is replaced with *1*;
(ii) *1* is replaced with *1 0 1 1 0 1 0 1* and *0* is replaced with *1 0 1 1 0*;
(iii) *1* is replaced with *1 0 1 1 0* and *0* is replaced with *1 0 1*; and
(iv) *1* is replaced with *1 0 1* and *0* is replaced with *1 0*.

Subjects were asked to listen to 3 minute long strings of these syllables. After the training phase, they were presented auditorily with 64 pairs of candidate strings lasting ten seconds each, and were asked to indicate which of the pair was most similar to the training set. As Fig. 3 shows, subjects were rather good at making the discrimination.

From the perspective presented here, that human ability is not all that surprising. What people are clearly doing is identifying *constituents* in the string of syllables, obviously at a higher level of abstraction than the mere list of signals heard. To be sure, the 'brackets' implied by these constituents are not pronounced (this is speech generated by a synthesizer), intonation is flat, there are absolutely no clues for subjects other than what their own minds cannot help but project over the perceived structure. In fact, as was mentioned above already, one of the properties of the expression of Fibonacci grammars is that they are self-similar: any given generation of output can be seen to be composed of combinations of previous generations. This simply means that the strings have large constituents composed of sub-constituents, indeed of the same sort. There are no reliable transition probabilities of the first order (Markov) type in the recursive stimuli in this experiment, and since the stimuli are all pseudo-random, in order to discriminate one sample from another it must be the case that subjects hold an abstract representation of the systematic properties of the grammars. This seems to us like an interesting result.

Perhaps it is worth noting that the Fibonacci grammar generates deterministic chaos, which entails that, in some sense at least, the computational complexity of the relevant system generated is arbitrarily large. If human language 'rides on a Fibonacci wave', then the potential computational complexity of human language is, in the relevant sense, also arbitrarily large. Of course, we do not see this in human languages: if they were that complex, they would probably be unusable. So somehow capacity limitations—attention, memory, sampling rates, timing properties and more—interact with the, as it were, Fibonacci space of possibilities, to yield a viable physiological and, ultimately, psychological system. This too may, at least in part, underlay the range of variability found across human languages. In a sense to be understood, any human language is a kind of approximation, of a computational system that embodies arbitrary complexity but is nevertheless deterministic; or so it seems at the relevant level of abstraction.

9. Towards a conclusion

Just as we are willing to exhibit our limitations, we also want to be honest in declaring what we do not take to be a serious difficulty for this overall approach—and in so doing, we want to wrap up the present piece. Evidently, the 'forces' we have talked about here are not forces, without the scare quotes, in any standard physical sense. Certainly, only in the latter instance does present-day science know how to relate the dynamics thereof to the emergent logarithmic eddies. We cannot presently calculate how it is that *any* living entity (be it a flower or a peacock, a virus coating or a musical thought) emerges in F terms from systemic dynamics. In all treatments a metaphor is invoked, since nobody knows how to translate physical conditions into computational interactions. Knowing that would be akin to resolving the mind/body problem, or at least the form/body problem.

What we do not find useful is to deny the sorts of facts discussed here, for whichever organism one cares to investigate, simply because we are missing a step in the reasoning. It is a fact, as plain as it gets, that F growth patterns centrally exist in nature—including natural language. We can explain these patterns very nicely (which we have not attempted here in any detail) when standard forces interact in well-understood conditions. We can describe the patterns computationally (which we have done already) in a variety of domains, and it is descriptively the case that, in such instances, there is a sense, also, in which interactive 'forces' yield the relevant results. Herein the metaphor, however, and the obvious question for the future: How do these computational 'forces' emerge, in nature? We do not know, in either general living systems or mental systems. But to insist: we are in the same darkness in mental systems as in other living systems more generally. This is not a particular problem for cognitive scientists, which means that, aside from being justified in pursuing this research, we should all keep trying to understand what's going on, from various angles. Our information seems useful in that regard at least.

Do other F patterns in the human mind relate to the sorts of linguistic patterns we have discussed here? It is hard to know. As we have shown, it is not unreasonable to unify apparently unrelated phonologically and syntactically-based F patterns, in terms of considerations of the Analysis by Synthesis sort, applying at various levels of analytical coarseness, plausibly in the acquisition process (thus counting as an instance of growth). If our approach is correct, the apparent reflex of syntactic structuring on phonological patterns, or vice-versa for that matter, are best seen as deeper conditions emerging because of the broad dynamics of the language faculty, existing as a min-max solution to a very interesting conflict of cognitive interests. If so, the question is whether similar considerations apply to other mental (musical, aesthetic, even moral) patterns, and these too emerge as the result of dynamic interactions, or rather they are to be seen as parasitic on a linguistic solution, much as we think that the classic structure in (6) is probably not a primitive in itself. Only serious research will determine this.

We do not want to finish without emphasizing, yet again, that the considerations we have raised here are very old, at least within linguistics. The Analysis by Synthesis method that we have assumed, via Townsend and Bever [41], is the first serious model of linguistic performance in the twentieth century—and it remains the dominant paradigm after almost half a century of research. The idea that language somehow emerges at the cross-roads between two orthogonal computations was first explicitly explored, as far as we know, by Lucien Tesnière in the thirties (and published posthumously in 1959). The notion, aptly expressed by Jakobson, that there is beauty in language, in the sense this notion had for Plato, is probably as old as humanity. The only thing that modern science allows us to do with all these insights is to place them in an always evolving perspective, which new technologies (e.g., in molecular biology) and paradigms (e.g., the return of the Laws of Form via studies in Complex Dynamic Systems) undoubtedly help us focus.

What does it mean for linguistic forms to obey those conditions of growth, including such nuanced asymmetrical organizations as we have studied here? Why has natural law carved this particular niche among logically imaginable ones, thus yielding the attested sub-cases, throughout the world languages? What in the evolution of the species directed the emergence of these forms, and How was that evolutionary path even possible and, in the end, successful? The biolinguistics take that we assume attempts to address these matters from the perspective of coupling results within contemporary linguistic theorizing with machinery from systems biology and bio-physics more generally. Again, we do not fully understand the 'embodiment' of any of the F patterns in living creatures. But the program seems clear: proceeding with hypothetical models based on various theoretical angles, from physical and bio-molecular ones, to grammatical studies isolating abstract patterns in the linguistic phenotype. A synthesis proceeding this way seems viable in the not so distant future, at the rate that new discoveries march.

Acknowledgements

We wish to thank Leonid Perlovsky for his interest in our work and useful commentary. We also thank the following individuals for various suggestions: Tom Bever, Cedric Boeckx, Andrew Carney, Bill Idsardi, David Medeiros and Doug Saddy. The usual disclaimers apply about our errors.

References

[1] Berwick R, Weinberg A. The grammatical basis of linguistic performance. Cambridge: MIT Press; 1984.
[2] Blevins J. The syllable in phonological theory. In: Goldsmith J, editor. The handbook of phonological theory. Oxford: Blackwell; 1995. p. 206–44.
[3] Bresnan J. Sentence stress and syntactic transformations. Language 1971;47(2):257–81.
[4] Boeckx C. Understanding minimalist syntax: lessons from locality in long-distance dependencies. Oxford: Blackwell; 2007.
[5] Boeckx C. Bare syntax. Oxford: Oxford University Press; 2008.
[6] Boeckx C. Elementary syntactic structures. Ms. Harvard University; in progress.
[7] Boeckx C, Carnie A, Medeiros D. Some consequences of natural law in syntactic structure. Ms. University of Arizona and Harvard University; 2005.
[8] Carstairs-McCarthy A. The origins of complex language: an inquiry into the evolutionary beginnings of sentences, syllables, and truth. Oxford: Oxford University Press; 2000.
[9] Chomsky N. The minimalist program. Cambridge: MIT Press; 1995.
[10] Chomsky N. Minimalist inquiries. In: Martin R, Michaels D, Uriagereka J, editors. Step by step. Cambridge: MIT Press; 2000. p. 89–155.
[11] Chomsky N. On phases. In: Freidin R, Otero C, Zubizaretta ML, editors. Foundational issues in linguistic theory. Cambridge: MIT Press; 2008. p. 133–66.
[12] Chomsky N, Halle M, Lukoff F. On accent and juncture in English. In: Halle M, Lunt HG, McLean H, van Schooneveld CH, editors. For Roman Jakobson. The Hague: Mouton; 1956. p. 65–80.
[13] Chomsky N, Miller G. Introduction to the formal analysis of natural language. In: Luce R, Bush R, Gallanter E, editors. Handbook of mathematical psychology, vol. II. New York: Wiley; 1963. p. 269–307.
[14] Crain S, Thornton R. Investigations in universal grammar: a guide to experiments in the acquisition of syntax and semantics. Cambridge: MIT Press; 1998.
[15] Douady S, Couder Y. Phyllotaxis as a physical self-organized growth process. Phys Rev Lett 1992;68:2098–101.
[16] Duckworth GE. Structural patterns and proportions in Virgil's Aeneid: a study in mathematical composition. Ann Arbor: University of Michigan Press; 1962.
[17] Fink A, Saddy D. Ms. University of Reading; in progress.
[18] Fodor J. Concepts: where cognitive science went wrong. Oxford: Oxford University Press; 1998.
[19] Fodor J, Piattelli-Palmarini M. Evolution without adaptation. Cambridge: MIT Press; forthcoming.
[20] Fortuny J. The emergence of order in syntax. PhD dissertation, Universitat de Barcelona; 2007.
[21] Gallego A. Phase theory and parametric variation. PhD dissertation, Universitat Autònoma de Barcelona; 2007.
[22] Halle M, Stevens K. Speech recognition: a model and a program for research. IEEE Transactions on Information Theory 1962;8(2):155–60.
[23] Hinzen W. Mind design and minimal syntax. Oxford: Oxford University Press; 2006.
[24] Huang J. Logical relations in Chinese and the theory of grammar. PhD dissertation, MIT; 1982.
[25] Idsardi W. Combinatorics for metrical feet. Ms. University of Maryland; 2005.
[26] Jackendoff R. Semantic interpretation in generative grammar. Cambridge: MIT Press; 1972.

[27] Kaye J, Lowenstamm J, Vergnaud J-R. The internal structure of phonological elements: a theory of charm and government. Phonology Yearbook 1985;2:305–28.
[28] Kayne R. Connectedness and binary branching. Dordrecht: Foris; 1984.
[29] Kayne R. The antisymmetry of syntax. Cambridge: MIT Press; 1994.
[30] Kroch A, Santorini B. The syntax of natural language: an online introduction using the Trees program. ling.upenn.edu/~beatrice/syntax-textbook, 2007.
[31] Lasnik H. Analyses of negation in English. PhD dissertation, MIT; 1972.
[32] Lewontin RC. The triple helix: gene, organism and environment. Cambridge, MA: Harvard University Press; 2000.
[33] Marantz A. Generative linguistics within the cognitive neuroscience of language. The Linguistic Review 2005;22:429–45.
[34] Medeiros D. PhD dissertation. University of Arizona; forthcoming.
[35] Perlmutter D. Sonority and syllable structure in American Sign Language. Linguistic Inquiry 1992;23:407–42.
[36] Perlovsky LI. Toward physics of the mind: Concepts, emotions, consciousness, and symbols. Phys Life Rev 2006;3(1):22–55.
[37] Poeppel D, Idsardi W, van Wassenhove V. Speech perception at the interface of neurobiology and linguistics. Philosophical Transactions of the Royal Society, B: Biological Sciences 2008;363:1071–86.
[38] Richards M. Deriving the edge: What's in a phase? Ms., University of Cambridge; 2006.
[39] Soschen A. Natural law: the dynamics of syntactic representations in the minimalist program. DEAL. Linguistics in Potsdam 2006;25:43–76.
[40] Saussure F. Cours de linguistique générale. In: Bally C, Sechehaye A, editors. Lausanne: Payot; 1916.
[41] Townsend DJ, Bever TG. Sentence comprehension: the integration of Habits and Rules. Cambridge, MA: MIT Press; 2001.
[42] Uriagereka J. Rhyme and reason: an introduction to minimalist syntax. Cambridge: MIT Press; 1998.
[43] Uriagereka J. Where to look for Fibonacci patterns in syntax. In: Fortuny F, Gallego A, editors. Spelling out universal grammar. Barcelona: Universitat de Barcelona; forthcoming.