

Towards a Biologically-Plausible Computational Model of Human Language Cognition

Hilton Alers-Valentín¹ and Sandiway Fong²

¹*Linguistics and Cognitive Science, University of Puerto Rico-Mayagüez*

²*Department of Linguistics, University of Arizona-Tucson*
hilton.alers@upr.edu, sandiway@arizona.edu

Keywords: Strong Minimalist Thesis, Cognitive Modeling, Computational Linguistics, Explainable Artificial Intelligence

Abstract: The biolinguistics approach aims to construct a coherent and biologically plausible model/theory of human language as a computational system coded in the brain that for each individual recursively generates an infinite array of hierarchically structured expressions interpreted at the interfaces for thought and externalization. Language is a recent development in human evolution, is acquired reflexively from impoverished data, and shares common properties through the species in spite of individual diversity. Universal Grammar, a genuine explanation of language, must meet these apparently contradictory requirements. The Strong Minimalist Thesis (SMT) proposes that all phenomena of language have a principled account rooted in efficient computation, which makes language a perfect solution to interface conditions. LLMs, albeit their remarkable performance, cannot achieve the explanatory adequacy necessary for a language competence model. We implemented a computer model assuming these challenges, only using language-specific operations, relations, and procedures satisfying SMT. As a plausible model of human language, the implementation can put to test cutting-edge syntactic theory within the generative enterprise. Successful derivations obtained through the model signal the feasibility of the minimalist framework, shed light on specific proposals on the processing of structural ambiguity, and help to explore fundamental questions about the nature of the Workspace.

1 INTRODUCTION

Recent advances in linguistics and cognitive science have contributed to a fuller understanding of human language, its biological bases, computational nature, abstract representations, mental processing, and neurological realization. We address the theoretical foundations, architectural possibilities, and limitations of a biologically plausible computational model of natural language syntax. We begin by drawing extensively from original sources to present an overview of the current state of language theory within the biolinguistic framework (§2) and the Minimalist Program’s ‘prime directive’, the Strong Minimalist Thesis (SMT), along with its three factors: computational operations, principles of efficient computation, and language-specific conditions (§3). Our model is to be contrasted with the currently popular LLMs as cognitive models: relevant aspects are briefly summarized in §4. The next sections deal with our SMT-driven implemented model, describing the basis for computation (§5), and exploring fundamental questions for parsing and the Workspace using the minimalist

model (§6). A brief conclusion in §7 summarizes the main findings.

2 LANGUAGE AS A MENTAL ORGAN: THEORETICAL FOUNDATIONS

The biological nature of human language has been pursued as an object of scientific inquiry since the 1950s [Lenneberg, 1967] and is well established in recent literature on ethology, genetics, evolution, and neurology [Di Sciullo et al., 2010, Enard et al., 2002, Musso et al., 2003, Fitch, 2010, Moro, 2015, Berwick and Chomsky, 2016, Friederici, 2018]. Language and thought seem to be a distinctive species property, “common to humans in essentials apart from severe pathology and without significant analogue in the non-human world” [Chomsky, 2021].

An important distinction must be made between the *Faculty of Language* (FL), the distinctive property shared by the human species, which enables each in-

dividual to develop (or grow in true biological fashion) a particular mind-internal system for the generation and expression of thought, an *I-language* [Chomsky, 2020], where *I* stands for *internal, individual, intensional*. The Faculty of Language constitutes the initial state of which an I-language is the steady state, “a property of the organism, a computational system coded in the brain that for each individual recursively generates an infinite array of hierarchically structured expressions, each formulating a thought, each potentially externalized in some sensory-motor (SM) medium — what we may call the *Basic Property of Language*” [Chomsky, 2021]. The combinatorial component of an I-Language is called a *grammar*, with computational procedures to form new objects, while the *lexicon* (LEX) is the set of lexical items (LI), the primitives or atoms of computation for I-Language. LIs are formatives “in the traditional sense as minimum “meaning-bearing” and functional elements.” It is conjectured that the variety of languages might be completely localized in peripheral aspects of LEX and in externalization [Chomsky, 2021].

Genomic evidence indicates that modern humans, who emerged around 200,000 years ago, began separating not long after (in evolutionary time), roughly 150,000 years ago. Since all descendants share the capacity for Language, one must conclude that Language already evolved before human populations became separated. Research suggests that FL emerged fairly suddenly in evolutionary time. If so, we would expect that FL should be simple in structure, with few elementary principles of computation, satisfying the evolvability condition [Huybregts, 2017, Chomsky, 2020, Chomsky, 2021].

A rather puzzling property of language, which appears to be its most fundamental, is *structure dependence*: “we ignore the simple computation on linear order of words [adjacency], and reflexively carry out a computation on abstract structure” [Chomsky, 2023a]. For example, the utterance

(1) the man who fixed the car carefully packed his tools

is ambiguous between ‘fixed the car carefully’ and ‘carefully packed his tools’. However,

(2) carefully, the man who fixed the car packed his tools

is unambiguously ‘carefully packed his tools’. The adverb in initial position is linearly closer to the verb phrase “fixed the car” than to “packed his tools.” But if we now assume the following simplified abstract structure

(3) carefully [[the man who fixed the car] packed his tools]

the adverb in initial position is structurally closer to the verb phrase “packed his tools” than to “fixed the car.” Experimental work has shown that this principle is available to children from the onset of syntactic acquisition at 18 months [Shi et al., 2020]. This suggests that from infancy and on through life, we reflexively ignore the linear order of words that we hear, and attend only to what we never hear but our minds construct: abstract structures generated by the mind and operations on these structures, which are non-trivial [Chomsky, 2021, Chomsky, 2023a].

Adequate theories of the Faculty of Language must say how acquiring one language differs from acquiring another, and how human children differ from other animals in being able to acquire either language (or both) given a suitable course of experience [Berwick et al., 2011]. It has been argued that in species-specific growth and development—in this case, of the language organ—, individual differences in outcome typically arise from interacting factors [Chomsky, 2005, Berwick et al., 2011]:

- (4)(i) innate factors or genetic endowment, apparently nearly uniform for the species, of which a distinction is made between domain-general and domain-specific.
- (ii) external stimuli or experience.
- (iii) natural law (also called third factors), like physical and developmental constraints and principles of efficient computation, data analysis, and structural architecture.

To solve the logical problem of language acquisition, it is proposed that within the human mind/brain there is a language acquisition device or Universal Grammar (UG), which is an innate factor and therefore part of the human species’ biological endowment. UG is the theory of the faculty of language, although the same term is sometimes used to refer the initial state of the human language faculty itself, i.e., the component of I-language that is shared by all human speakers which determines the class of possible (as opposed to impossible) acquired I-languages.

UG has goals that at first seem contradictory. It must meet at least three conditions:

- (5)(i) It must be rich enough to overcome the problem of poverty of stimulus.
- (ii) It must be simple enough to have evolved under the conditions of human evolution.
- (iii) It must be the same for all possible languages, given commonality of UG.

“We achieve a genuine explanation of some linguistic phenomenon only if it keeps to mechanisms that satisfy the joint conditions of learnability, evolvability,

and universality, which appear to be at odds” [Chomsky, 2021].

3 STRONG MINIMALIST THESIS

“The basic principle of language (BP) is that each language yields an infinite array of hierarchically structured expressions, each interpreted at two interfaces, conceptual-intentional (C-I) and sensorimotor (SM) – the former yielding a “language of thought” (LOT), perhaps the only such LOT; the latter in large part modality-independent, though there are preferences. The two interfaces provide external conditions that BP must satisfy, subject to crucial qualifications mentioned below. If FL is perfect, then UG should reduce to the simplest possible computational operation satisfying the external conditions, along with principles of minimal computation (MC) that are language independent. The Strong Minimalist Thesis (SMT) proposes that FL is perfect in this sense” [Chomsky, 2015].

As formulated, the SMT involves three factors: computational operations, interface or language-specific conditions, and principles that determine efficient computation [Freidin, 2021].

3.1 Computational Operations

The simplest, most economical structure-building operation (SBO) proposed is MERGE as binary set formation:

$$(6) \text{ MERGE}(X, Y) = \{X, Y\}$$

where X and Y are either lexical items or syntactic objects (SOs) already generated. MERGE allows for two subcases: EXTERNAL MERGE (EM), where X and Y are distinct and INTERNAL MERGE (IM), where one is contained in the other, i.e. X is a term of Y or Y is a term of X . This containment relation, or term-of, as is technically known, is defined recursively: Z is a term of W if Z is a member of W or of a term of W . INTERNAL MERGE yields displacement, with two copies. Thus if Y is contained in X , then $\text{MERGE}(X, Y) = \{Y, \{X, Y\}\}$ [Chomsky, 2020].

Each application of MERGE is a stage in the derivation of a SO, and there is a Workspace (WS) at each stage. A WS is “a set of already generated items that are available for carrying the derivation forward (along with LEX, which is always available). WS determines the current state of the derivation. Derivations are Markovian, in the sense that the next step does not have access to the derivational history; nevertheless, WS includes everything previously gener-

ated” [Chomsky, 2021]. At its most general formulation,

$$(7) \text{ MERGE}(X_1, \dots, X_n, \text{WS}) = \text{WS}' = \{\{X_1, \dots, X_n\}, \text{W}, \text{Y}\}$$

To satisfy SMT and LSCs, $n = 2$ (Binarity) and Y is null (nothing else is generated, by virtue of Minimal Yield, see §3.2). W is whatever is unaffected by the operation, hence carried over [Chomsky, 2021].

As the simplest SBO to satisfies both SMT and the Basic Property, (binary) MERGE counts as genuine explanation.

It is proposed that adjunction is the result of an operation PAIRMERGE [Chomsky, 2004], “which yields asymmetric (ordered) pairs rather than symmetric (unordered) sets, permitting the identification of an adjunct in a phrase-modifier configuration. PAIRMERGE may also be required for unstructured coordination” [Chomsky et al., 2019]. However, since PAIRMERGE is a formally distinct operation from simplest MERGE, it raises problems of evolvability.

Within this system, the only other permissible relation is unbounded set, which is generated by another SBO, FORMSET (FS), such that for all $X \in \text{WS}$,

$$(8) \text{ FS}(X_1, \dots, X_n) = \{X_1, \dots, X_n\}$$

It should be noted that binary FS is distinct from (External) MERGE in lacking its special θ -related properties. FS is assumed to be a costless operation available freely for all inquiry, used in constructing the workspace WS and the lexicon LEX [Chomsky, 2023b].

Agreement phenomena in languages indicate that there must be an operation AGREE relating features of syntactic objects [Chomsky, 2000, Chomsky, 2001]. AGREE seems to be a structure-dependent, asymmetric operation, that relates initially unvalued ϕ -features (grammatical person, gender, and number) on a Probe to matching, inherent ϕ -features of a Goal within the Probe’s search space (structural sister) [Chomsky et al., 2019].

3.2 Efficient Computation

Principles of efficient computation are regarded as language-independent laws of nature, “third factors” in language design. A natural condition for efficient computation is limiting search, a property of SMT. “For an operation O to apply to items it must first locate them. It must incorporate an operation Σ that searches LEX and WS and selects items to which O will apply. It is fair to take Σ to be a third factor element, [...] available for any operation” [Chomsky, 2021]. However, for the sake of computational efficiency, Σ must be limited. This condition, MINIMAL

SEARCH (MS), is another freely available “least effort” condition.

(9) *Minimal Search*

Σ searches as far as the first element it reaches and no further.

In other words, in searching WS, MS selects a member X of WS, but no term of X [Chomsky, 2021].

MERGE also satisfies other corollaries of limiting Σ , as for example BINARITY.

Another important condition, MINIMAL YIELD (MY), limits the construction of searchable SOs:

(10) *Minimal Yield*

MERGE should construct the fewest possible new items that are accessible to further operations.

EM(P,Q) necessarily constructs one such SO: {P, Q} itself. IM(P, Q), where Q is a term of P, constructs {P, Q}, “where P contains a copy of Q, call it Q’”. The operation therefore creates two new elements: {P, Q} and the raised element Q. But Q’ is no longer accessible, thanks to MS. Q’ is protected from Σ by Q. Hence only one new accessible element is added, satisfying MY” [Chomsky, 2021].

MERGE does nothing more than take two SOs X, Y and construct a new single SO, the set {X, Y}. It otherwise leaves the combined objects unaltered. This is known as the NO TAMPERING CONDITION (NTC). So, if X, an SO, has property F before being merged with Y, another SO, X will still have property F after merging with Y.

(11) *No Tampering Condition*

MERGE does not affect the properties of the elements of computation in any way. [Hornstein, 2018].

Cyclic computation constitutes another property of computational efficiency: “A MERGE-based system will be compositional in general character: the interpretation of larger units at the interfaces will depend on the interpretation of their parts, a familiar observation in the study of every aspect of language. If the system is computationally efficient, once the interpretation of small units is determined it will not be modified by later operations –the general property of STRICT CYCLICITY that has repeatedly been found” [Chomsky, 2007]. Strict cyclicity is imposed by PHASE THEORY: “the computation will not have to look back at earlier phases as it proceeds, and cyclicity is preserved in a very strong sense” [Chomsky, 2008].

(12) *Phase Theory*

When a phase is constructed, it is dispatched to interpretation at CI and can no longer be accessed by Σ [Chomsky, 2021].

In formal languages, instances of an inscription are treated as occurrences of the same inscription, which is necessary for proper interpretation of the derivation. That convention is called STABILITY. For identical inscriptions, on the other hand, human language makes a distinction between repetitions (with different interpretations) and copies (with the same interpretation).

(13) *Stability*

Structurally identical inscriptions in the Copy relation must have exactly the same interpretation.

In fact, what is special about natural language is not the existence of copies, but rather of non-copies (repetitions) [Chomsky, 2021].

3.3 Language-Specific Conditions

Language-Specific Conditions (LSCs), which subsume the sometimes called interface or legibility conditions, are the domain of UG. That means that they are not learned from PLD nor can be reduced to or deduced from third factors (like principles of efficient computation).

First, two principles seem to be fundamental. From SMT, one guideline for inquiry is derived:

(14) *Principle S*

The computational structure of language should adhere as closely as possible to SMT [Chomsky, 2023b].

On the other hand, if I-language is basically a thought-generating system, as the Basic Property entails, it optimally should observe the following principle:

(15) *Principle T*

All relations and structure-building operations (SBO) are thought-related, with semantic properties interpreted at CI. [Chomsky, 2023b].

Language must provide argument structure at CI. Thus, predicates or θ -assigners (like verbs and prepositions, for example) assign semantic descriptions called thematic or θ -roles to constituents in θ -positions (arguments of such predicate) This is known as Θ -Theory, a module in the Principles and Parameters framework [Chomsky, 1981, Adger, 2003].

(16) *Θ -Theory*

- (i) A θ -assigner assigns θ -roles to θ -positions.
- (ii) Every θ -role must be assigned.

Simplest MERGE is the one SBO that satisfies both language-specific and computational efficiency conditions. Furthermore, it follows from Principle T that MERGE is thought-related. An LSC, DUALITY

OF SEMANTICS, relates each subcase of MERGE to a category of thought:

(17) *Duality of Semantics*

EM is associated with θ -roles (propositions) and IM with force-/discourse-related functions (clauses). [Chomsky, 2021, Chomsky, 2023b].

The LSCs so far formulated seem to be concerned mainly with legibility conditions at CI. It may turn out that LSCs are restricted to the core function of language in generating thought.

4 COGNITIVE MODELS AND LLMs

A genuine model or theory of language must aim at *descriptive* and *explanatory adequacy*. A descriptively adequate model “is concerned to give a correct account of the linguistic intuition of the native speaker; [...] with the output of the [language] device [...] and specifies the observed data (in particular) in terms of significant generalizations that express underlying regularities in the language.” Furthermore, to achieve explanatory adequacy, the model must be “concerned with the internal structure of the device; that is, it aims to provide a principled basis, independent of any particular language, for the selection of the descriptively adequate grammar of each language” [Chomsky, 1964]. In other words, descriptive adequacy deals with the issue of *strong generation* of linguistic structures, as opposed to mere observational adequacy, which is only concerned with the *weak generation* of strings. Explanatory adequacy, on the other hand, deals with the problem of language acquisition. And beyond explanatory adequacy lies the deeper question of *why* language is the way it is.

Current approaches to artificial intelligence (AI), based almost exclusively on Deep Learning, show promising results in domains involving pattern recognition. In fact, Large Language Models (LLMs), a technological achievement of generative AI, have been proposed as theories of human language because of their impressive text generation when prompted by a query. LLMs are characterized by (i) enormous requirements of training data and energy consumption, (ii) attention mechanisms that “allow the next word in sequence to be predicted from some previous far in the past”, (iii) embeddings, with words stored as vectors whose locations in a multi-dimensional vector space are supposed to “include not just some aspects of meaning but also properties that determine how words can occur in sequence”, and (iv) “massive over-parameterization” that should provide “space for

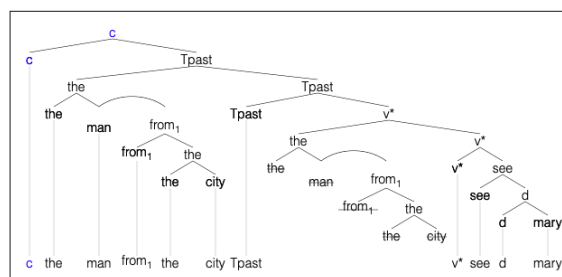
inferring hidden variables and relationships” [Piantadosi, 2023]. However, these characteristics that are intrinsic to their architecture automatically disqualify LLMs as genuine explanatory models of human cognition. The bigger LLMs are, the more prone they are to overparameterization (having more parameters than data points), which tends to overfitting (memorizing the data rather than generalizing). GPT-4 is rumored to have around 100 trillion parameters, is estimated to be trained with data in an order of magnitude of petabytes (1024 TB or approximately 10^{15} bytes), and would take almost a million megawatt-hour of training, which would be the equivalent to running the human brain for around five million years at the oft-cited figure of 20 watts [Fong, 2023]. LLMs are very susceptible to perturbations in the training data, and even fail to produce some commonsense inferences and generalizations that are natural to humans [Fong, 2022]. The big data approach to LLMs is not only unsuitable for domains where massive amounts of data are not available, but is also in stark contrast with human language development, which thrives with impoverished data and produces correct generalizations from almost non-existent direct evidence (zero-shot learning) [Alers-Valentín et al., 2023]. Experiments have shown that LLMs (i) do not exhibit the same linguistic biases and representations as humans in acceptability judgements and language universals, (ii) do not align with humans in the competence-performance distinction, (iii) lack a distinction between likelihood and grammaticality and (iv) lack the capacity for generalizations common in humans [Katzir, 2023, Moro et al., 2023]. LLM failures in inference, generalization, and trustworthiness are due to the absence of explicit internal representations and a dynamic world model [Lenat and Marcus, 2023]. Lastly, a most fundamental difference between humans and LLMs is “the fact that there is no comparable state for the machine to the “Impossible language state” characterizing human brains” [Moro et al., 2023]. When human brains compute impossible languages (e.g., violations of structure dependence), the canonical networks selectively associated to language computation are progressively inhibited [Musso et al., 2003]. “LLMs do not have intrinsic limits nor any similar hardware correspondence [nor] any embodied syntax which is in fact the fingerprint of human language. [...] [D]espite their (potential) utility for language tasks, [LLMs] can by no means be considered as isomorphic to human language faculty as resulting from brain activity” [Moro et al., 2023].

5 THE BASIS FOR COMPUTATION

The SMT sets forth strict and austere guidelines for the simplest possible generative theory of language. In particular, computational devices often taken for granted (for algorithmic reasons) are not permitted in the case of a SMT-based computational engine. For example, although a so-called *covering* phrase structure grammar (PSG), such as that employed in [Fong, 1991], to generate candidate parses (in accordance with X-bar theory and phrasal movement) would admit a variety of efficient and well-understood PSG algorithms, e.g. top-down Earley or bottom-up LR(*k*) methods, such a device would fail the test of evolutionary plausibility. Under the hypothesis that modern humans have only recently arrived (on the scene) via a small change that unlocked (Simplest) MERGE as the recursive basis of thought expression computation, there has not been enough time (on the evolutionary timescale) to evolve a multitude of other mechanisms. As PSG parsing algorithms compute hierarchical expressions (from linear word order), unless it is already in use by the brain (for other purposes), we cannot adopt such an approach.

Following Chomsky's lead, even slightly elaborated versions of MERGE, e.g. parallel or sideways MERGE (proposed in the linguistics literature), are not permitted as possible operations. The challenge therefore is to make use of existing resources only, ideally without modification. Assuming only MERGE can build hierarchical expressions, this is what we must utilize (rather than a separate structure-building parsing primitive). Contra what we think of as left-to-right (online) parsing, MERGE is categorically bottom-up in nature (and right-to-left for head-initial languages such as English), i.e. MERGE takes two pre-existing objects and create a larger expression composed of the original two (without modifying either one, i.e. we must respect the Non-Tampering Condition (NTC)). MERGE operates on a scratchpad, termed a Workspace (WS), and initially applies to linguistic heads sourced from LEX, the lexicon. The result of MERGE is dumped back into the WS for possible input for further Merges. Prior inputs to MERGE are not available for further computation. Therefore the WS cannot increase in size (and complexity) during the course of a derivation, a desirable result with respect to computational complexity. Further limiting operative complexity, the WS must be structured, i.e. divided into sub-WSs. For example, the subject of a sentence, as in Figure 1, is constructed independently from the main spline of the sentence. In Figure 1, the adjunct phrase *from the city* is also computed indepen-

dently and linked to the subject *man* by PAIRMERGE (shown as a curve).¹ Finally, computation is localized into Phases, this results in staging of recursively embedded clauses, as in Figure 2, sentence taken from ([Chomsky, 2000]: 110), in which three Phases are identified, viz. $P_1 = \textit{that global warming is taken seriously}$, $P_2 = \textit{that glaciers are receding}$, and $P_3 = \textit{the demonstration } P_2 \textit{ showed } P_1$.



Spell-out:
the man from the city -ed(sg) see -acc Mary (after morpheme realization)
the man from the city see -ed(sg) Mary -acc (after affix-hop)
the man from the city see -ed(sg) Mary -acc (after morpheme realization, stage 2)
the man from the city saw Mary

Figure 1: The man from the city saw Mary.

Structured WS must be constructed so that MERGE can fire appropriately. For the examples in Figures 1 and 2, the appropriate list of initial heads are (18) and (19), respectively. We use single bracketing to indicate the sub-lists that must be computed in a sub-WS, and double bracketing to indicate the sub-Phases.

- (18) Mary, d, see, v*, [man, [city, the, from₁], the],
 T_{past}, c
- (19) [[warming, global, d, take, seriously, prt, v~, must, T, c_e]], show, v_{unerg}, [[glaciers, d, recede, v_{unacc} prog, v~, T, c_e]], demonstration, the], T_{past}, c

In addition to Set and Pair MERGE, the system also implements FORMSET (introduced earlier in §3.1), which handles general coordination and so-called *stacked* relatives as in (20)(i), the two relative clauses, CP₁ and CP₂ in (20)(ii) and (20)(iii), respectively, forms a set {CP₁, CP₂}. The parse is shown in Figure 3 (with a horizontal line connecting the two set members).

- (20)(i) The student who lives here who studies English

¹Generally, for any pair {XP, YP}, XP and YP non-head phrases, XP and YP must be independently computed, both for (Set) MERGE and FORMSET (as will be discussed for example (20)(i)). The same applies for cases of PAIRMERGE in the case of adjunction <X/XP, YP>, an example is *the man from the city* in Figure 1. with X = *man* and XP = *from the city*.




Figure 2: The demonstration that glaciers are receding showed that global warming must be taken seriously.

- (ii) $CP_1 = \{ \text{who}_{rel} \text{ student}, \{ C_{rel}, \{ \text{who}_{rel} \text{ student lives here} \} \} \}$
- (iii) $CP_2 = \{ \text{who}_{rel} \text{ student}, \{ C_{rel}, \{ \text{who}_{rel} \text{ student studies English} \} \} \}$

The parses shown previously are automatically constructed by our computer program from these structured WSs. Implementation details aside, these lists of heads are pre-ordered so that constituents are formed in the correct positions, e.g. internal arguments such as objects MERGE earlier to verbs than subjects.² All MERGE operations follow precisely the theory described earlier in §3.1. A successful derivation is obtained when all heads in a WS are used in a sequence of (valid) Merges that lead to a single syntactic object, e.g. the parses shown above, with the WS emptied.

6 FUNDAMENTAL QUESTIONS FOR PARSING

In the previous section, an extremely important (and rather fundamental) question has been omitted, viz. how do we (magically) come up with a correct WS that results a convergent derivation? With respect to thought generation, we can ask: which heads from LEX get placed into the initial Lexical Array (LA)? The literature is largely silent on this question. In our computer program that computes the syntactic structures shown above, initial WS's have been hand-constructed. It seems there is an apparent circularity in the logic: how can we know which initial WS

²There is good reason to assume that subjects and objects are structurally asymmetric, as in the thematic configuration $\{ \text{Subject}, \{ v^*, \{ R, \text{Object} \} \} \}$. For example, in Figure 1, the verbal root R is *see*, Subject is *the man from the city* and Object is *Mary*. This asymmetry affects the timing of MERGE; in particular, the Object must be scheduled for MERGE before the Subject.

will correctly drive MERGE and converge without actually testing/doing MERGE? Moreover, we also need to know the connection between arguments and verbs to structure the LA into a structured WS so that *John saw Mary* does not spell out as *Mary saw John*. (See also note 2.)

In the case of parsing, the question is perhaps better posed. We ask, given the signal, e.g. speech, sign or written language: which heads in LEX activate and populate the initial WS? We can tentatively assume recognition of a word (and its morphemes) activate appropriate lexical heads and trigger transfer from LEX into the WS. Note there may be more than one appropriate WS for the signal. For example, a structurally ambiguous sentence such as (21) may have two distinct parses, beginning with two distinct initial WS's shown in (22)(i)–(ii).

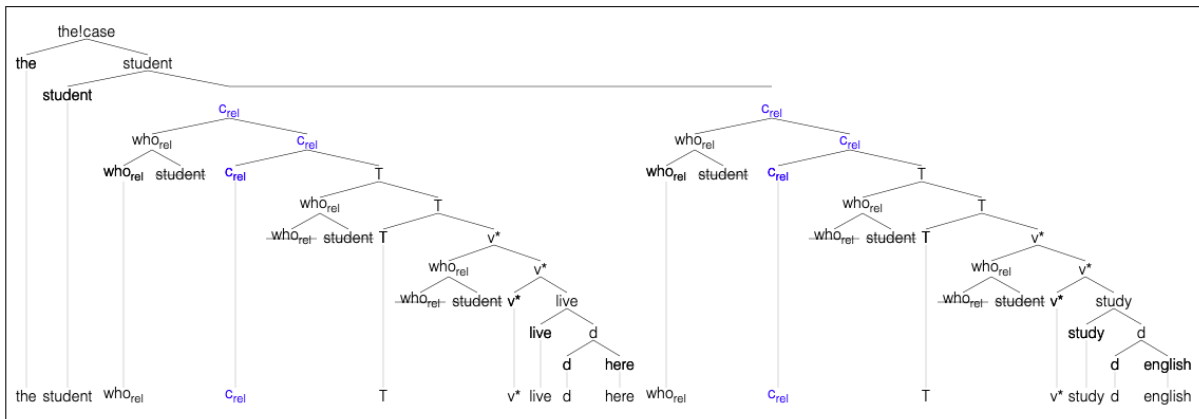
(21) The chicken is ready to eat

- (22)(i) chicken, the, eat, v_{unerg} , [PRO, d0], T_{inf} , c, ready, v_{be} , T, c
- (ii) eat, v_{unerg} , [chicken, the], T_{inf} , c, ready, v_{be} , T, c

We propose both WS's must be initialized when we hear (21). The fact that humans readily spot this kind of ambiguity in the absence of contextual cues means that both parses shown in Figure 4 are computed.³ (Of course, given sufficient disambiguating context, we may strongly prefer one analysis over the other.)

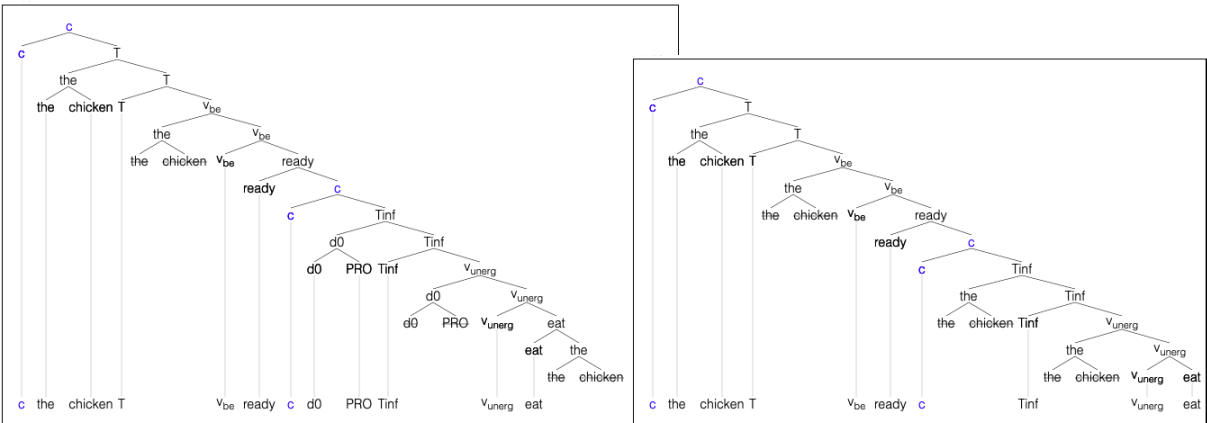
In the extended verbal projection in English, rather than a single verb, a sequence of heads must enter the WS, viz. the verbal root itself plus a choice of verbal categorizer (the so-called 'little v') and a tense morpheme typically. For example, the verb *break* is compatible with different alternations as in (23)(i)

³In Figure 4, in the left parse *the chicken* is the object of the verb *eat*, and the subject of *eat* in the right parse. Also in the left parse, arbitrary PRO, a pronominal meaning (*for anyone*, is the subject of *eat*.



Spell-out:
 the student who -s live -acc here who -s study -acc English (after morpheme realization)
 the student who live -s here -acc who study -s English -acc (after affix-hop)
 the student who live -s here -acc who study -s English -acc (after morpheme realization, stage 2)
 the student who lives here who studies English

Figure 3: The student who lives here who studies English.



Spell-out:
 the chicken -s be ready to eat (after morpheme realization)
 the chicken be -s ready to eat (after affix-hop)
 the chicken be -s ready to eat (after morpheme realization, stage 2)
 the chicken is ready to eat

Spell-out:
 the chicken -s be ready to eat (after morpheme realization)
 the chicken be -s ready to eat (after affix-hop)
 the chicken be -s ready to eat (after morpheme realization, stage 2)
 the chicken is ready to eat

Figure 4: The chicken is ready to eat.

and (23)(ii) distinguished in the WS by choice of v^* or v_{unacc} heads (that have different syntactic properties). *Broke* in (23)(i) and (23)(ii) require (24)(i) and (24)(ii), respectively. (Table 1 provides a sample inventory of fundamental heads our system implements.)

- (23)(i) John broke the vase
- (ii) The vase broke

- (24)(i) *broke* = spellout of $v^* + break + T_{past}$
- (ii) *broke* = spellout of $v_{unacc} + break + T_{past}$

Therefore both little v 's have to be activated and populate distinct initial WS's.⁴ Choice of little v results in different syntactic structures. For parsing, we

⁴The lack of space prevents us from going into details on the theoretical possibilities for functional heads. But

propose a simple answer to this apparent conundrum (and to the *John saw Mary* vs. *Mary saw John* question discussed earlier): spellout of convergent parses must match the initial signal.

Even when there is only one LA for the input signal, it is possible that it can be structured differently as WS's. For example, (25) has a single (unordered)

generally, multiple heads (with different syntactic properties) will be available for selection, see Table 1. But not all heads are available. For example, both v^* and v_{unacc} are available for the verb *break* given (23)(i)-(ii). However, the same is not true for *crack*, as *John cracked an unknown code* is grammatical but *the unknown code cracked* is not. The association between appropriate v and verbal roots is both explicitly acquired (from primary linguistic evidence) and computed based on meaning, i.e. with consideration for Lexical Semantics.

Table 1: A selection of functional heads.

Functional head	uFeatures	Other	Spell-Out (English)
<i>Little v</i>			
v^* (transitive)	phi:Person,Number	ef(theta); value acc Case	
v_{unerg} (unergative)		ef(theta)	
v_{unacc} (unaccusative)		ef check theta	<i>be</i>
v_{\sim} (<i>be</i>)		ef check theta	<i>be</i>
<i>Auxiliaries</i>			
prt (participle)	phi:Number; Case	ef	<i>-ed</i>
prog (progressive)		ef	<i>-ing</i>
perf (perfective)			<i>-en</i>
<i>Tense</i>			
T (non-past)	phi:Person,Number	ef; value nom Case	[1,sg]:-m, [2,sg]:-re, [3,sg]:-s, [-,pl]:-re
T_{past} (past)	phi:Person,Number	ef; value nom Case	[1,sg]:-ed, [1,pl]:-ed, [2,-]:-ed,
			[3,sg]:-ed, [3,pl]:-ed
T_{inf} (non-finite)	phi:Person,Number	ef; value null Case	<i>to</i>
<i>Complementizer</i>			
C (declarative)		Local Extent (LE) head	
C_e (decl embedded)	T	ef; LE head	
C_Q (interrogative)	Wh; T	ef; LE head	<i>do</i>
C_eQ (int embedded)	Wh; T	ef; LE head	<i>do</i>
C_{rel} (relative)	Wh; T	ef(wh); LE head	

LA, as in (26), but two different possibilities as structured WS's in (27)(i)–(ii) leading to different parses (and therefore interpretations at INT). The two parses are shown in Figure 5.

- (25) The man saw the boy with a telescope
- (26) {boy, telescope, a, with₁, the, see, v^* , man, the, T_{past} , c}
- (27)(i) boy, [telescope, a, with₁], the, see, v^* , [man, the], T_{past} , c
- (ii) boy, the, see, v^* , [man, the], [telescope, a, with₁], ' T_{past} , c

Finally, it is also possible that two distinct sentences have the same initial WS. For example, both (28)(i)–(ii) can be generated from WS (29), assuming *that* is the optional spellout of Tense at the complementizer position.⁵

- (28)(i) Mary thinks Sue will buy the book
- (ii) Mary thinks that Sue will buy the book
- (29) book, the, buy, v^* , [Sue, d], will, T, c_e , think, v_{unerg} , [Mary, d], T, c

The brain is largely chemical in nature, rather than electrical ([Gallistel and King, 2009] notwithstanding), one that is slow in operation and rather demanding of computational efficiency, but with limited possibilities for parallelism. The challenge for parsing

⁵Along the same lines, the WS shown earlier in (19) for *the demonstration (that) glaciers are receding showed (that) global warming must be taken seriously* generates four different sentences as both complementizers, *that*, are optional.

is to limit the population of WS candidates to those that are combinatorially plausible in a biological setting. If the SMT is on the right track, not only is the locus of variation in language shifted to Externalization, but also constraints on processing should be induced from primary linguistic evidence, i.e. what the child hears and internalizes must help limit computational complexity so that, ultimately, comprehension becomes not only possible (given enough resources), but be readily made efficient (over time). Much work remains to be done, particularly with respect to how memory and the lexicon must be organized, but we believe that the SMT has both simplified the theoretical landscape and severely limited the biologically plausible options for parsing (that we must now explore).

7 CONCLUSIONS

In this paper, we have sketched how a practical sentence parser can be designed (and constructed) while adhering to the austere conditions imposed by evolutionary considerations. Language is a computational system coded in the mind/brain that for each individual recursively generates an infinite array of hierarchically structured expressions interpreted at the interfaces for thought and externalization. As a cognitive organ, language is subject to constraints from domain-general and domain-specific innate factors, external stimuli, and natural laws like principles of efficient computation. Universal Grammar, a genuine explanation of language, must satisfy these apparently

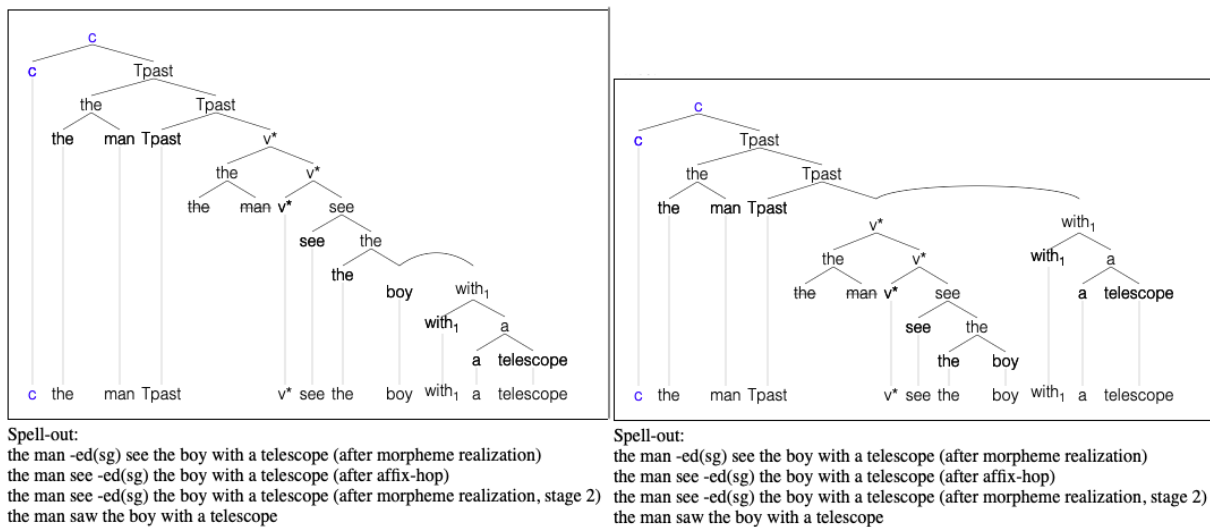


Figure 5: The man saw the boy with a telescope.

contradictory conditions. The Strong Minimalist Thesis (SMT) proposes that all phenomena of language have a principled account rooted in efficient computation, which makes language a perfect solution to interface conditions. LLMs, in spite of their performance achievements, do not satisfy the conditions of learnability, evolvability and universality, necessary for a biologically-plausible competence model, as their data and energy requirements vastly exceed the capacities of organic systems. Our proposed system constitutes a model of UG, as it only implements operations, relations, and procedures that satisfy SMT, like computational operations external and internal MERGE (simplest SBO), PAIRMERGE (for adjuncts), FORMSET (for stacked relative clauses) and AGREE (to relate Probe's and Goal's matching features). Computational devices often taken for granted (for algorithmic reasons) are not permitted in the case of a SMT-based computational engine, not even efficient and well-understood PSG algorithms that would fail the test of evolutionary plausibility. The system implements derivation by phases, following the strict cyclicity condition. It also satisfies principles of efficient computation, restricting all operations to comply by NTC, MS, and MY. Our minimalist language model automatically construct parses from structured WSs. This posits a question about how do humans come up with a correct WS that results a convergent derivation, a fundamental problem in understanding human cognition and thought generation. In processing structural ambiguity, it is proposed that different WS's must be initialized when hearing an ambiguous utterance. Since humans reflexively detect this kind of ambiguity in the absence of contextual cues, it suggests that several parses must be computed. On

the other hand, it is possible that structurally different sentences can be derived from the same initial WS. Within the brain's biological limitations, MERGE must operate in parallel and linguistic stimuli must induce constraints on processing, which still needs to be investigated. Other SMT devices should be implemented to develop a more complete model of this promising, cutting-edge framework.

Acknowledgements

This material is based upon work supported by the National Science Foundation (NSF) under Grant No. 2219712 and 2219713. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

- Adger, D. (2003). *Core Syntax: A Minimalist Approach*. Oxford University Press, Oxford.
- Alers-Valentín, H., Fong, S., and Vega-Riveros, J. F. (2023). Modeling syntactic knowledge with neuro-symbolic computation. In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence*, volume 3, pages 608–616.
- Berwick, R. C. and Chomsky, N. (2016). *Why only us: Language and evolution*. MIT Press.
- Berwick, R. C., Pietroski, P., Yankama, B., and Chomsky, N. (2011). Poverty of the stimulus revisited. *Cognitive Science*, 35(7):1207–1242.

- Chomsky, N. (1964). *Current Issues in Linguistic Theory*, volume 38. Mouton, The Hague.
- Chomsky, N. (1981). *Lectures on Government and Binding. The Pisa Lectures*. Number 9 in Studies in Generative Grammar. Foris, Dordrecht.
- Chomsky, N. (2000). Minimalist Inquiries: The Framework. In *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, pages 89–155. MIT Press.
- Chomsky, N. (2001). Derivation by phase (mitopl 18). In *Ken Hale: A Life in Language*, pages 1–52. MIT Press.
- Chomsky, N. (2004). Beyond explanatory adequacy. In *Structures and Beyond*, pages 104–131. Oxford University Press.
- Chomsky, N. (2005). Three factors in language design. *Linguistic inquiry*, 36(1):1–22.
- Chomsky, N. (2007). *Approaching UG from below*, volume 89. Mouton de Gruyter Berlin.
- Chomsky, N. (2008). On phases. In *Foundational Issues in Linguistic Theory: Essays in Honor of Jean-Roger Vergnaud*, pages 133–166. MIT Press.
- Chomsky, N. (2015). *The Minimalist Program: 20th Anniversary Edition*. MIT Press.
- Chomsky, N. (2020). Fundamental operations of language: Reflections on optimal design. *Cadernos de Linguística*, 1(1):1–13.
- Chomsky, N. (2021). Minimalism: Where are we now, and where can we hope to go. *Gengo Kenkyu*, 160:1–41.
- Chomsky, N. (2023a). Genuine explanation and the Strong Minimalist Thesis. *Cognitive Semantics*, 8(3):347–365.
- Chomsky, N. (2023b). The Miracle Creed and SMT. <http://www.icl.keio.ac.jp/news/2023/04/2023-theoretical-linguistics-at-keio-emu.html>.
- Chomsky, N., Gallego, Á. J., and Ott, D. (2019). Generative grammar and the faculty of language: Insights, questions, and challenges. *Catalan Journal of Linguistics*, pages 229–261.
- Di Sciullo, A. M., Piattelli-Palmarini, M., Wexler, K., Berwick, R. C., Boeckx, C., Jenkins, L., Uriagereka, J., Stromswold, K., Cheng, L. L.-S., Harley, H., Wedel, A., McGilvray, J., van Gelderen, E., and Bever, T. G. (2010). The biological nature of human language. *Biolinguistics*, 4(1):004–034.
- Enard, W., Przeworski, M., Fisher, S. E., Lai, C. S., Wiebe, V., Kitano, T., Monaco, A. P., and Pääbo, S. (2002). Molecular evolution of FOXP2, a gene involved in speech and language. *Nature*, 418(6900):869–872.
- Fitch, W. T. (2010). *The evolution of language*. Cambridge University Press.
- Fong, S. (1991). *Computational properties of principle-based grammatical theories*. PhD thesis, Massachusetts Institute of Technology.
- Fong, S. (2022). Simple models: Computational and linguistic perspectives. *Journal of the Institute for Research in English Language and Literature*, 46:1–48.
- Fong, S. (2023). SMT and Parsing. Lecture slides.
- Freidin, R. (2021). The Strong Minimalist Thesis. *Philosophies*, 6(4):97–115.
- Friederici, A. D. (2018). *Language in our Brain: The Origins of a Uniquely Human Capacity*. The MIT Press, Cambridge, MA.
- Gallistel, C. and King, A. (2009). *Memory and the Computational Brain: Why Cognitive Science will Transform Neuroscience*. Wiley-Blackwell.
- Hornstein, N. (2018). The minimalist program after 25 years. *Annual Review of Linguistics*, 4:49–65.
- Huybregts, M. R. (2017). Phonemic clicks and the mapping asymmetry: How language emerged and speech developed. *Neuroscience & Biobehavioral Reviews*, 81:279–294.
- Katzir, R. (2023). Why large language models are poor theories of human linguistic cognition. A reply to Piantadosi (2023). <https://lingbuzz.net/lingbuzz/007190>.
- Lenat, D. and Marcus, G. (2023). Getting from Generative AI to Trustworthy AI: What LLMs might learn from Cyc. <https://arxiv.org/abs/2308.04445>.
- Lenneberg, E. H. (1967). *Biological Foundations of Language*. John Wiley, New York.
- Moro, A. (2015). *The boundaries of Babel: The brain and the enigma of impossible languages*. MIT Press.
- Moro, A., Greco, M., and Cappa, S. F. (2023). Large languages, impossible languages and human brains. *Cortex*, 167:82–85.
- Musso, M., Moro, A., Glauche, V., Rijntjes, M., Reichenbach, J., Büchel, C., and Weiller, C. (2003). Broca’s area and the language instinct. *Nature Neuroscience*, 6(7):774–781.
- Piantadosi, S. (2023). Modern language models refute Chomsky’s approach to language. <https://lingbuzz.net/lingbuzz/007180/v1.pdf>.
- Shi, R., Legrand, C., and Brandenberger, A. (2020). Toddlers track hierarchical structure dependence. *Language Acquisition*, 27(4):397–409.