

The language capacity: architecture and evolution

Noam Chomsky¹

Published online: 1 July 2016
© Psychonomic Society, Inc. 2016

Abstract There is substantial evidence that the human language capacity (LC) is a species-specific biological property, essentially unique to humans, invariant among human groups, and dissociated from other cognitive systems. Each language, an instantiation of LC, consists of a generative procedure that yields a discrete infinity of hierarchically structured expressions with semantic interpretations, hence a kind of “language of thought” (LOT), along with an operation of externalization (EXT) to some sensory-motor system, typically sound. There is mounting evidence that generation of LOT observes language-independent principles of computational efficiency and is based on the simplest computational operations, and that EXT is an ancillary process not entering into the core semantic properties of LOT and is the primary locus of the apparent complexity, diversity, and mutability of language. These conclusions are not surprising, since the internal system is acquired virtually without evidence in fundamental respects, and EXT relates it to sensory-motor systems that are unrelated to it. Even such properties as the linear order of words appear to be reflexes of the sensory motor system, not available to generation of LOT. The limited evidence from the evolutionary record lends support to these conclusions, suggesting that LC emerged with *Homo sapiens* or not long after, and has not evolved since human groups dispersed.

Keywords Language comprehension · Language capacity

Evidently, we can study the evolution of some system only to the extent that we know what it is. In the present case, what

has evolved is not languages, which do not evolve, in the technical sense of the term, any more than states of the visual system evolve. Rather, what has evolved is the capacity for language (LC), analogous to the genetic basis for a mammalian, not insect, visual system.¹

The study of the evolution of LC was opened, and given a solid basis, by Eric Lenneberg in his classic *Biological Foundations of Language* (1967). Among many other contributions, Lenneberg reviewed the sharp divergence between human language and the symbolic systems of other animals and the dissociations of language from other cognitive processes. These discoveries and insights have since been considerably extended and deepened. He also discussed the biological plausibility of qualitative discontinuity, a conclusion that is also much better grounded today—and a logical necessity for accounting for the emergence of systems of discrete infinity such as human language, which may indeed be unique in this regard, and the root of others, such as knowledge of arithmetic, a matter that greatly concerned Darwin and Wallace.

LC appears to be a true species property, unique to humans in essentials and invariant among human groups, indicating that it has undergone little if any evolutionary change at least since our ancestors left Africa some 60,000 ya—possibly about twice that long, as some very recent genomic studies have indicated. Because of the apparent uniformity and stability of LC, the theory of theory of LC has in recent decades been called “universal grammar” (UG), adapting a traditional notion to a new biological context. UG is not to be confused with generalizations about surface properties, a topic that has been studied in highly informative ways, notably by Joseph Greenberg, but an entirely different one.

✉ Noam Chomsky
chomsky@mit.edu

¹ Massachusetts Institute of Technology, Cambridge, MA, USA

¹ For further detail, throughout, see Berwick and Chomsky (2016), and sources cited.

The most salient property of LC is that languages consist of a discrete infinity of structured expressions that are interpretable in a definite way by the conceptual–intentional system (CI) of thought and action and by a sensory-motor system (SM) for externalization, thus yielding a sound-meaning correlation over an infinite range, though the sound system, while convenient, is only one option. LC is thus based on a generative computational system (GEN) consisting of combinatorial operations that operate on atomic elements of the lexicon to yield the two interface representations. At the CI interface, GEN yields a kind of “language of thought” (LOT); in different terminology, a system of “conceptual structures” (CS). The lexicon includes substantive terms (word-like elements, though not words) and others that play a role in GEN.

The study of evolution of LC therefore seeks to establish the nature of the Lexicon and of the combinatorial operations of GEN. As to the former, there is solid evidence that the terms of human languages are radically different from those of animal symbolic systems. The latter appear to be associated with mind-independent events, while even the simplest human terms violate this condition, an insight tracing back to classical Greece. There is, in short, no “word–object/thing” relationship for human language, though terms and more complex expressions of language may be used to refer to mind-external objects, a different matter. The evolutionary origin of even the most elementary human concepts/terms is a mystery, particularly those used to refer, and hence to relate the internal language to the external world. Even the task of giving accurate descriptions of these fundamental elements has barely been undertaken.²

Turning to GEN, one of its crucial properties, which came to light as soon as the first efforts were undertaken to construct explicit accounts of language (generative grammars), is *structure-dependence*: the operations of GEN apply to structures, not strings, even ignoring such elementary properties of externalized expressions as linear order. Thus, quite generally, operations keep to minimal distance, but in such expressions as “the man and the woman *copula* angry at *reflexive*,” the copula and reflexive are plural, agreeing with the closest *phrase* (*the man and the woman*), not singular, agreeing with the linearly closest possible antecedent (*the woman*).

To take a more interesting case, consider (1):

- (1) Which girls and boys did the men expect to like each other

Like other anaphors, *each other* seeks the closest potential antecedent. However, it skips *the men*, instead choosing *which girls and boys*, ignoring linear (and even structural) order, raising questions to which we will return.

Similar cases abound. Consider (2), (3):

- (2) birds that fly instinctively swim
(3) instinctively, birds that fly swim

Sentence (2) is ambiguous: “fly instinctively” or “instinctively swim”; (3) is unambiguous: “instinctively swim.” Again, the construal rule that relates the adverb to the verb ignores linear distance and observes minimal structural distance.

Structure-dependence holds for all relevant constructions in all languages. That universal property would be paradoxical if linear order were available to GEN, since linear distance is far more easily computed than structural distance. The only plausible conclusion is that linear order is simply not available to generation of the core semantic properties at CI. Other suggestions have been made, but those that are clear enough to investigate quickly collapse on inspection.³

We conclude, then, that there is a fundamental asymmetry between the two interfaces: GEN yields an infinite array of structured expressions at CI (LOT/CS). Ancillary operations of externalization map the structures produced by GEN to some sensory modality, introducing linear order and other properties that are required by SM but are irrelevant to core semantic/conceptual properties of language.

Much other evidence supports this conclusion, including simple semantic facts. Thus, languages can have head-complement or complement-head order (say, VO or OV), the *head parameter*, but the semantic relationships are identical, a fact that generalizes widely (with apparent exceptions that go beyond the discussion here). Or consider again (1). Its interpretation at CI is, loosely, (4):

- (4) for which girls and boys, the men expected those girls and boys to like each other

That would follow directly if what is generated at CI is (5), with two copies of the phrase *which girls and boys*, in which case *each other*, as expected, selects the closest potential antecedent (as always ignoring linear order, hence ignoring *boys*):

- (5) [which girls and boys] did the men expect [which girls and boys] to like each other

Accordingly, we expect UG to determine that (5) is generated at CI, while the externalized form (1) is derived by deletion of a copy, thus minimizing internal and SM computation.

What reaches the mind, then, is (5), while what reaches the mouth/ear is (1), violating minimal distance, a matter of no concern given the basic architecture of LC, with its fundamental asymmetry of interfaces. Far more intricate cases of semantic interpretation fall under the same mechanisms.

² For review and sources, see Chomsky (2013).

³ See Berwick et al. (2011).

A wide variety of linguistic evidence supports the conclusion that externalization is an ancillary process, with properties that are in part a reflex of SM. There is independent support from psycholinguistics and the neurosciences. One productive paradigm has been to present subjects with two kinds of artificial languages, one modeled on a human language, thus conforming to UG, the other violating UG, for example, with rules that use linear order—say, negating a sentence by placing the negative particle after the third word, a rule far simpler to compute than actual linguistic rules. It turns out that, in the case of conformity to UG, there is normal activation in the language areas, though not when linear order is used. In that case, the task is interpreted as a non-linguistic puzzle, or so brain activity indicates. Work by Neil Smith and Ianthi-Maria Tsimpli with a cognitively impaired but linguistically fluent subject reached similar conclusions. They also found that normals were unable to deal with the violations of UG using linear order if the task was presented to them as linguistic, though they could handle the problem if it was presented as a puzzle.⁴

Though many questions remain, the basic architectural asymmetry appears to be reasonably well established. It has many consequences. One is that uses of language that depend on externalization, notably communication, are even more peripheral to the nature of language, contrary to widely held beliefs. The asymmetry also illustrates again the sharp dissociation of language from animal symbolic systems.

A problem that arises at once is the apparent variety, diversity, and easy mutability of language, properties that appear to be inconsistent with the conclusion that UG is a species property, not having evolved throughout detectable human history. The problem would be resolved if these properties of language are confined, perhaps completely, to the lexicon and to externalization (hence to the SM interface). There is substantial evidence supporting this conclusion. In particular, investigation of languages of wide typological variety has repeatedly shown that apparent sharp differences in underlying structures—for example, “flat” vs highly structured expressions—dissolve under deeper inquiry,⁵ lending support to the conclusion that the core system generating CI is close to uniform, as we should expect simply from the fact that it is acquired with little direct evidence, in many cases none at all.

Still, a question arises about the evolution of the options of variation—“parameters” as they have been called in recent work. The optimal conclusion would be that they did not evolve at all. Some might represent alternative solutions to the cognitive problem of relating a virtually invariant system that may satisfy conditions of minimal computation to SM systems that had long been in place and are unrelated to it. Another possibility is that parameters, at least many of them, do not exist. They are options left open by GEN and the

overarching principles of minimal computation. Consider again the head parameter. It is simply a mismatch between GEN, which assigns no order, and the SM systems that require it. Languages have to make a choice for externalization, and do it one way or another. English and Japanese, for example, are virtual mirror images. Note that the options resulting from mismatch lend further support to the conclusion that the internal system lacks order and other surface arrangements.⁶

Pursuing further the nature of what has evolved, UG, consider again (1). It illustrates the ubiquitous property of *dislocation*: phrases that are heard in one position are interpreted both there and somewhere else, in a position where similar phrases can appear. Like structure-dependence, that property has been regarded as highly problematic from the earliest days of generative grammar.

The challenge is to explain why UG should have such properties as the ones reviewed here. There is a very straightforward answer, which holds over a substantial range: GEN observes language-independent principles of computational efficiency, and UG itself is based on the simplest computational operation, call it Merge, an operation that is embedded somehow in every more complex computational procedure: take two objects X and Y already constructed and form a new object Z without modifying X or Y, or imposing any further structure on them: thus $\text{Merge}(X, Y) = \{X, Y\}$.

This elementary assumption suffices to yield the observations discussed above: the asymmetry of the interfaces and the locus of surface complexity and variety; the ubiquity of dislocation; structure-dependence; minimal structural distance for anaphoric and other construal; and the difference between what reaches the mind for semantic interpretation and what reaches the mouth and ear. I will not run through the explanation, presented elsewhere (see note 1). Instead, a few words on the possible evolution of LC (UG).

Empirical evidence on evolution of LC is very slight. One assumption, already mentioned, can be put forth with fair confidence: there has been little or no evolution of language (that is, of LC) since our ancestors left Africa and quickly spread all over the world. A second fact is that evidence for non-trivial forms of symbolic behavior date to not long before that time, less than 100,000 ya. And there is no convincing evidence of LC in any species other than *Homo sapiens* (presumably with date of origin about 200,000 ya). So we are considering very brief periods of evolutionary time, even if the dates are somewhat extended.

Recall that there are two basic problems: the origins of the lexical/conceptual atoms and of GEN. On the former, there is very little to say, particularly for the terms used to refer and to relate internal language to the world. On GEN, the simplest assumption consistent with the limited evidence about evolution

⁴ Musso et al. (2003). Smith (2004); Smith and Tsimpli (1996).

⁵ For one very important example (of many), see Legate (2002).

⁶ See Moro (2008), 107. For detailed development of these ideas, see Obata, Epstein and Baptista (2015).

and what is known about what evolved is that some, perhaps relatively small, rewiring of the brain produced the simplest computational operation, accessing the lexicon and thus yielding LOT/CS and the capacity for thought, planning, reflection, and creativity, over an (in principle) unbounded range. Yielding selectional advantages, the trait might have proliferated through a small community, over time providing the motive for devising a mode of externalization of the internally generated linguistic expressions that express thoughts.⁷ That poses the cognitive problem of relating an internal system that might conform closely to principles of computational efficiency to SM systems that have no special relationship to it, a problem that can be solved in many ways, possibly by just solving the mismatch problems, perhaps with little if any further evolution of any significance. These are, it seems, the simplest assumptions that would yield an outcome satisfying the observations about LC discussed here. And they appear to be consistent with the very limited evidence available about the evolution of LC.

Needless to say, these remarks only scratch the surface of what is at all understood, let alone the vast array of questions yet to be addressed.

References

- Berwick, R., Pietroski, P., Yankama, B., & Chomsky, N. (2011). Poverty of the stimulus revisited. *Cognitive Science*, *35*, 1–36.
- Berwick, R., & Chomsky, N. (2016). *Why Only US? Language and Evolution*. Cambridge MA: MIT press.
- Chomsky, N. (2013). Notes on Denotation and Denoting. In I. Caponigro & C. Cecchetto (Eds.), *From Grammar to Meaning: the Spontaneous Logicality of Language*. New York: Cambridge U. Press.
- Legate, J. A. (2002). *Warlpiri. Theoretical Implications*. Cambridge MA: MIT press.
- Lenneberg, E. (1967). *The Biology of Language*. New York: Wiley.
- Moro, A. (2008). *The Boundaries of Babel: The Brain and the Enigma of Impossible Languages*. Cambridge MA: MIT press.
- Musso, M., Moro, A., Glauche, V., Rijntjes, M., Reichenbach, J., & Buechel, C. (2003). Broca's area and the language instinct. *Nature Neurosci.*, *6*, 774.
- Obata, M., S. Epstein, M. Baptista (2015). Can crosslinguistically variant grammars be formally identical? Third factor underspecification and the possible elimination of parameters of UG. *Lingua* *156* (2015) 1–16. Online at www.ScienceDirect.com.
- Smith, N., & Tsimpli, I. (1996). *Mind of a Savant*. Oxford: Blackwell.
- Smith, N. (2004). *Chomsky: Ideas and Ideals*, 2nd ed. (New York: Cambridge, 2004), p. 136.

⁷ The system of motor control presumably has much deeper evolutionary origins, and was exapted for externalization of internal language. The same for sign, of course. See reference of note 1.