
LECTURE I: WHAT IS LANGUAGE?

Author(s): NOAM CHOMSKY

Source: *The Journal of Philosophy*, DECEMBER 2013, Vol. 110, No. 12 (DECEMBER 2013), pp. 645-662

Published by: Journal of Philosophy, Inc.

Stable URL: <https://www.jstor.org/stable/43820808>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*

JSTOR

THE JOURNAL OF PHILOSOPHY

VOLUME CX, NO. 12, DECEMBER 2013

WHAT KIND OF CREATURES ARE WE?

LECTURE I: WHAT IS LANGUAGE?

The general question I would like to address in these lectures is an ancient one: What kind of creatures are we? I am not deluded enough to think I can provide a satisfactory answer, but it seems reasonable to believe that in some domains at least, particularly with regard to our cognitive nature, there are insights of some interest and significance, some new, and that it should be possible to clear away some of the obstacles that hamper further inquiry, including some widely accepted doctrines with foundations that are much less stable than often assumed.

I will consider three specific questions, increasingly obscure: What is language? What are the limits of human understanding (if any)? And what is the common good to which we should strive? I will begin today with the first, and will try to show how what may seem at first to be rather narrow and technical questions, if pursued carefully, can lead to some far-reaching conclusions that are significant in themselves, and differ sharply from what is generally believed—and often regarded as fundamental—in the relevant disciplines: cognitive science in a broad sense, including linguistics, and philosophy of language and mind.

Throughout, I will be discussing what seem to me virtual truisms, but of an odd kind. They are generally rejected. That poses a dilemma, for me at least. And perhaps you too will be interested in resolving it.

Turning to language, it has been studied intensively and productively for 2,500 years, but with no clear answer to the question of what language is. I will mention later some of the major proposals. We might ask just how important it is to fill this gap. For the study of any aspect of language the answer should be clear. Only to the extent that there is an answer to this question, at least tacit, is it

possible to proceed to investigate serious questions about language, among them acquisition and use, origin, language change, diversity and common properties, language in society, the internal mechanisms that implement the system, both the cognitive system itself and its various uses, distinct though related tasks. No biologist would propose an account of the development or evolution of the eye, for example, without telling us something fairly definite about what an eye is, and the same truisms hold of inquiries into language. Or should. Interestingly, that is not how the questions have generally been viewed, a matter to which I will return.

But there are much more fundamental reasons to try to determine clearly what language is, reasons that bear directly on the question of what kind of creatures we are. Darwin was not the first to conclude that “the lower animals differ from man solely in his almost infinitely larger power of associating together the most diversified sounds and ideas”;¹ “almost infinite” is a traditional phrase to be interpreted today as actually infinite. But Darwin was the first to have expressed this traditional concept within the framework of an incipient account of human evolution.

A contemporary version is given by one of the leading scientists studying human evolution, Ian Tattersall. In a recent review of the currently available scientific evidence, he observes that it was once believed that the evolutionary record would yield “early harbingers of our later selves. The reality, however, is otherwise, for it is becoming increasingly clear that the acquisition of the uniquely modern [human] sensibility was instead an abrupt and recent event....And the expression of this new sensibility was almost certainly crucially abetted by the invention of what is perhaps the single most remarkable thing about our modern selves: language.”² If so, then an answer to the question “What is language?” matters greatly to anyone concerned with understanding our modern selves.

Tattersall dates the abrupt and sudden event as probably lying somewhere within the very narrow window of 50–100,000 years ago. The exact dates are unclear, and not relevant to our concerns here, but the abruptness of the emergence is. I will return to the vast and burgeoning literature of speculation on the topic, which generally adopts a very different stance.

If Tattersall’s account is basically accurate, as the very limited empirical evidence indicates, then what emerged in the narrow

¹ Charles Darwin, *The Descent of Man* (London: J. Murray, 1871), chapter 3.

² Ian Tattersall, *Masters of the Planet: The Search for Our Human Origins* (New York: Palgrave Macmillan, 2012), p. xi.

window was an infinite power of “associating...the most diversified sounds and ideas,” in Darwin’s words. That infinite power evidently resides in a finite brain. The concept of finite systems with infinite power was well understood by mid-twentieth century. That made it possible to provide a clear formulation of what I think we should recognize to be the most basic property of language, which I will refer to just as the Basic Property: each language provides an unbounded array of hierarchically structured expressions that receive interpretations at two interfaces, sensorimotor for externalization and conceptual-intentional for mental processes. That allows a substantive formulation of Darwin’s infinite power, or going back much farther, of Aristotle’s classic dictum that language is sound with meaning—though work of recent years shows that “sound” is too narrow, and there is good reason, to which I will return, to think that the classic formulation is misleading in important ways.

At the very least, then, each language incorporates a computational procedure satisfying the Basic Property. Therefore a theory of the language is by definition a generative grammar, and each language is what is called in technical terms an I-language—“I” standing for internal, individual, and intensional: we are interested in discovering the actual computational procedure, not some set of objects it enumerates, what it “strongly generates,” in technical terms, loosely analogous to the proofs generated by an axiom system.

There is also a notion “weak generation”—the set of expressions generated, analogous to the set of theorems generated. There is also a notion “E-language,” standing for external language, which many—not me—identify with a corpus of data, or with some infinite set that is weakly generated.³ Philosophers, linguists, and cognitive and computer scientists have often understood language to be what is weakly generated. It is not clear that the notion weak generation is even definable for human language. At best, it is derivative from the more fundamental notion of I-language. These are matters extensively discussed in the 1950s, though not properly assimilated, I believe.⁴

I will restrict attention here to I-language, a biological property of humans, some subcomponent of (mostly) the brain, an organ of the mind/brain in the loose sense in which the term “organ” is used in

³ The term is mine. See Noam Chomsky, *Knowledge of Language: Its Nature, Origin, and Use* (New York: Praeger, 1986). But I defined it almost vacuously, as any concept of language other than I-language.

⁴ A source of misunderstanding may be that in early work, “language” is sometimes defined in introductory expository passages in terms of weak generation, though the usage was quickly qualified, for reasons explained.

biology. I take the mind here to be the brain viewed at a certain level of abstraction. The approach is sometimes called the biolinguistic framework. It is regarded as controversial, but without grounds, in my opinion.

In earlier years, the Basic Property resisted clear formulation. Taking some of the classics, for de Saussure, language (in the relevant sense) is a storehouse of word images in the minds of members of a community, which “exists only by virtue of a sort of contract signed by the members of a community.” For Leonard Bloomfield, language is an array of habits to respond to situations with conventional speech sounds and to respond to these sounds with actions. Elsewhere, Bloomfield defined language as “The totality of utterances made in a speech community”—something like William Dwight Whitney’s earlier conception of language as “the body of uttered and audible signs by which in human society thought is principally expressed,” thus “audible signs for thought”—though this is a somewhat different conception in ways to which I will return. Edward Sapir defined language as “a purely human and non-instinctive method of communicating ideas, emotions, and desires by means of a system of voluntarily produced symbols.”⁵

With such conceptions it is not unnatural to follow what Martin Joos called the Boasian tradition, holding that languages can differ arbitrarily and that each new one must be studied without preconceptions.⁶ Accordingly, linguistic theory consists of analytic procedures to reduce a corpus to organized form, basically techniques of segmentation and classification. The most sophisticated development of this conception was Zellig Harris’s *Methods*.⁷ A contemporary version is that linguistic theory is a system of methods for processing expressions.⁸

In earlier years it was understandable that the question “What is language?” received only such indefinite answers as the ones

⁵ Ferdinand de Saussure, *Course in General Linguistics* (1916; repr., New York: Philosophical Library, 1959), pp. 13–14. Leonard Bloomfield, “Philosophical Aspects of Language” (1942), in Charles F. Hockett, ed., *A Leonard Bloomfield Anthology* (Bloomington: Indiana, 1970), pp. 267–70; Bloomfield, *A Set of Postulates for the Science of Language* (Indianapolis: Bobbs-Merrill, 1926); Bloomfield, “A Set of Postulates for the Science of Language,” *Language*, II, 3 (September 1926): 153–64; William Dwight Whitney, *The Life and Growth of Language* (London: H. S. King, 1875); Edward Sapir, *Language* (New York: Harcourt, Brace, 1921), p. 8.

⁶ Martin Joos, comments in Joos, ed., *Readings in Linguistics* (Washington, D.C.: American Council of Learned Societies, 1957).

⁷ Zellig Harris, *Methods in Structural Linguistics* (Chicago: University Press, 1951).

⁸ A regression, I think, since it confuses the fundamentally different notions of competence and performance—roughly, what we know and what we do—unlike Harris’s system, which does not.

mentioned, ignoring the Basic Property. It is, however, surprising to find that similar answers remain current in contemporary cognitive science. Not untypical is a current study on evolution of language, where the authors open by writing that “we understand language as the full suite of abilities to map sound to meaning, including the infrastructure that supports it,”⁹ basically a reiteration of Aristotle’s dictum, and too vague to ground further inquiry. Again, no biologist would study evolution of the visual system assuming no more about the phenotype than that it provides the full suite of abilities to map stimuli to percepts along with whatever supports it.

Much earlier, at the origins of modern science, there were hints at a picture somewhat similar to Darwin’s and Whitney’s. Galileo wondered at the “sublimity of mind” of the person who “dreamed of finding means to communicate his deepest thoughts to any other person...by the different arrangements of twenty characters upon a page,” an achievement “surpassing all stupendous inventions,” even those of “a Michelangelo, a Raphael, or a Titian.”¹⁰ The same recognition, and the deeper concern for the creative character of the normal use of language, was soon to become a core element of Cartesian science-philosophy, in fact a primary criterion for the existence of mind as a separate substance. Quite reasonably, that led to efforts to devise tests to determine whether another creature has a mind like ours, notably by Géraud de Cordemoy.¹¹ These were somewhat similar to the “Turing test,” though quite differently conceived. De Cordemoy’s experiments were like a litmus test for acidity, an attempt to draw conclusions about the real world. Turing’s imitation game, as he made clear, had no such ambitions.

These important questions aside, there is no reason today to doubt the fundamental Cartesian insight that use of language has a creative character: it is typically innovative without bounds, appropriate to circumstances but not caused by them—a crucial distinction—and can engender thoughts in others that they recognize they could have expressed themselves. We may be “incited or inclined” by circumstances and internal conditions to speak in certain ways, not others, but we are not “compelled” to do so, as Descartes’s successors put it. We should also bear in mind that Wilhelm von Humboldt’s now

⁹ Dan Dediu and Stephen C. Levinson, “On the Antiquity of Language: The Reinterpretation of Neandertal Linguistic Capacities and Its Consequences,” *Frontiers in Psychology*, iv, 397 (Jul. 5, 2013): 1–17, doi:10.3389/fpsyg.2013.00397.

¹⁰ Galileo Galilei, *Dialogue Concerning the Two Chief World Systems* (1632), end of First Day.

¹¹ For references and discussion, see Chomsky, *Cartesian Linguistics*, 3rd ed. (New York: Cambridge, 2009), with extended (and improved) translations, and introduction by James McGilvray.

oft-quoted aphorism that language involves infinite use of finite means refers to *use*. More fully, he wrote that “language is quite peculiarly confronted by an unending and truly boundless domain, the essence of all that can be thought. It must therefore make infinite employment of finite means, and is able to do so through the power which produces identity of language and thought.”¹² He thus placed himself in the tradition of Galileo and others who associated language closely with thought, though going well beyond, while formulating one version of a traditional conception of language as “the single most remarkable thing about our modern selves,” in Tattersall’s recent phrase.

There has been great progress in understanding the finite means that make possible infinite use of language, but the latter remains largely a mystery despite significant progress in understanding conventions that guide appropriate use, a much narrower question. How deep a mystery is a good question, to which I will return in the next lecture.

A century ago Otto Jespersen raised the question of how the structures of language “come into existence in the mind of a speaker” on the basis of finite experience, yielding a “notion of structure” that is “definite enough to guide him in framing sentences of his own,” crucially “free expressions” that are typically new to speaker and hearer.¹³ The task of the linguist, then, is to discover these mechanisms and how they arise in the mind, and to go beyond to unearth “the great principles underlying the grammars of all languages,” and by unearthing them to gain “a deeper insight into the innermost nature of human language and of human thought”—ideas that sound much less strange today than they did during the structuralist/behavioral science era that came to dominate much of the field, marginalizing Jespersen’s concerns and the tradition from which they derived.

Reformulating Jespersen’s program, the primary task is to investigate the true nature of the interfaces and the generative procedures that relate them in various I-languages, and to determine how they arise in the mind and are used, the primary focus of concern naturally being “free expressions.” And to go beyond to unearth the shared biological properties that determine the nature of I-languages accessible to humans, the topic of UG, universal grammar, in the contemporary version of Jespersen’s “great principles underlying the grammars of all languages,” now reframed as a question of the

¹² Wilhelm von Humboldt, *On Language* (1836; repr., New York: Cambridge, 1988), p. 91.

¹³ Otto Jespersen, *The Philosophy of Grammar* (New York: H. Holt and Co., 1924).

genetic endowment that yields the unique human language capacity and its specific instantiations in I-languages.

The mid-twentieth-century shift of perspective to generative grammar within the biolinguistic framework opened the way to much more far-reaching inquiry into language itself and language-related topics. The range of empirical materials available from languages of the widest typological variety has enormously expanded, and they are studied at a level of depth that could not have been imagined sixty years ago. The shift also greatly enriched the variety of evidence that bears on the study of each individual language to include acquisition, neuroscience, dissociations, and much else, and also what is learned from the study of other languages, on the well-confirmed assumption that the capacity for language relies on shared biological endowment.

As soon as the earliest attempts were made to construct explicit generative grammars sixty years ago, many puzzling phenomena were discovered that had not been noticed as long as the Basic Property was not clearly formulated and addressed and syntax was just considered “use of words” determined by convention and analogy. This is somewhat reminiscent of the early stages of modern science. For millennia scientists had been satisfied with simple explanations for familiar phenomena: rocks fall and steam rises because they are seeking their natural place; objects interact because of sympathies and antipathies; we perceive a triangle because its shape flits through the air and implants itself in our brains; and so on. When Galileo and others allowed themselves to be puzzled about the phenomena of nature, modern science began—and it was quickly discovered that many of our beliefs are senseless and our intuitions often wrong. Willingness to be puzzled is a valuable trait to cultivate, from childhood to advanced inquiry.

One puzzle about language that came to light sixty years ago, and remains alive and I think highly significant in its import, has to do with a simple but curious fact. Consider the sentence “instinctively, eagles that fly swim.” The adverb “instinctively” is associated with a verb, but it is *swim*, not *fly*. There is no problem with the thought that eagles that instinctively fly swim, but it cannot be expressed this way. Similarly the question “can eagles that fly swim” is about ability to swim, not to fly.

What is puzzling is that the association of the clause-initial elements “instinctively” or “can” to the verb is remote and based on structural properties, rather than proximal and based solely on linear properties, a far simpler computational operation and one that would be optimal for processing language. Language makes use of a property

of minimal structural distance, never using the much simpler operation of minimal linear distance; in this and numerous other cases, ease of processing is ignored in the design of language. In technical terms, the rules are invariably *structure-dependent*, ignoring linear order. The puzzle is why this should be so—not just for English but for every language, not just for these constructions but for all others as well, over a wide range.

There is a simple and plausible explanation for the fact that the child reflexively knows the right answer in such cases as these, even though evidence is slight or non-existent: linear order is simply not available to the language learner confronted with such examples, who is guided by a deep principle that restricts search to minimal structural distance, barring the far simpler operation of minimal linear distance. I know of no other explanation. And this proposal of course at once calls for further explanation: Why is this so? What is it about the genetically determined character of language—UG—that imposes this particular condition?

The principle of minimal distance is extensively employed in language design, presumably one case of a more general principle, call it Minimal Computation, which is in turn presumably an instance of a far more general property of the organic world or even beyond. There must however be some special property of language design that restricts Minimal Computation to structural rather than linear distance, despite the far greater simplicity of the latter for computation and processing.

There is independent evidence from other sources, including the neurosciences, supporting the same conclusion. A research group in Milan studied brain activity of subjects presented with two types of stimuli: invented languages satisfying UG and others not conforming to UG; in the latter case, for example, a rule for negation that places the negative element after the third word, a far simpler computational operation than the rules for negation in human language. They found that in the case of conformity to UG there is normal activation in the language areas, though not when linear order is used.¹⁴ In that case the task is interpreted as a non-linguistic puzzle, so brain activity indicates. Work by Neil Smith and Ianthi-Maria Tsimpli with a cognitively impaired but linguistically gifted subject reached similar conclusions—but, interestingly, found that normals as well were unable to deal with the violations of UG

¹⁴ Mariacristina Musso, Andrea Moro, et al., "Broca's Area and the Language Instinct," *Nature Neuroscience*, VI (Jun. 22, 2003): 774–81, doi:10.1038/nn1077.

using linear order. As Smith concludes: “the linguistic format of the experiment appeared to inhibit them from making the appropriate structure-independent generalization, even though they could work out comparable problems in a non-linguistic environment with ease.”¹⁵

There is a small industry in computational cognitive science attempting to show that these properties of language can be learned by statistical analysis of Big Data. This is, in fact, one of the very few significant properties of language that has been seriously addressed at all in these terms. Every attempt that is clear enough to be investigated has been shown to fail, irremediably.¹⁶ But more significantly, the efforts are beside the point in the first place. If they were to succeed, which is a virtual impossibility, they would leave untouched the original and only serious question: *Why* does language invariably use the complex computational property of minimal structural distance in the relevant cases, while always disregarding the far simpler option of minimal linear distance? Failure to grasp this point is an illustration of the lack of willingness to be puzzled that I mentioned earlier, the first step in serious scientific inquiry, as recognized in the hard sciences at least since Galileo.

A broader thesis is that linear order is never available for computation in the core parts of language involving syntax-semantics. Linear order, then, is a peripheral part of language, a reflex of properties of the sensorimotor system, which requires it: we cannot speak in parallel, or produce structures, but only strings of words. The sensorimotor system is not specifically adapted to language in fundamental respects: the parts essential for externalization and perception appear to have been in place long before language emerged. There is evidence that the auditory system of chimpanzees might be fairly well adapted for human speech,¹⁷ though apes cannot even take the first step in language acquisition, extracting language-relevant data from the “blooming, buzzing confusion” surrounding them, as human infants do at once, reflexively, not a slight achievement. And though capacity to control the vocal tract for speech appears to be human-specific, that fact cannot bear too much weight given that production of human language is modality-independent,

¹⁵ Neil Smith, *Chomsky: Ideas and Ideals*, 2nd ed. (New York: Cambridge, 2004), p. 136. See also Neil Smith and Ianthi-Maria Tsimpli, *The Mind of a Savant: Language Learning and Modularity* (Cambridge: Blackwell, 1995).

¹⁶ Robert C. Berwick, Paul Pietroski, Beracah Yankama, and Noam Chomsky, “Poverty of the Stimulus Revisited,” *Cognitive Science*, xxxv, 7 (September/October 2011): 1207–42, doi:10.1111/j.1551-6709.2011.01189.x.

¹⁷ W. Tecumseh Fitch, “Speech Perception: A Language-Trained Chimpanzee Weighs In,” *Current Biology*, xxi, 14 (Jul. 26, 2011): R543–46, doi:10.1016/j.cub.2011.06.035.

as recent work on sign language has established, and there is little reason to doubt that apes have adequate gestural capacities. Evidently much deeper cognitive properties are involved in language acquisition and design.

Though the matter is not settled, there is considerable evidence that the broader thesis may in fact be correct: fundamental language design ignores order and other external arrangements. In particular, semantic interpretation in core cases depends on hierarchy, not the order found in the externalized forms. If so, then the Basic Property is not exactly as I formulated it before, and as it is formulated in recent literature—papers of mine too. Rather, the Basic Property is generation of an unbounded array of hierarchically structured expressions mapping to the conceptual-intentional interface, providing a kind of “language of thought”—and quite possibly the only such LOT, though interesting questions arise here. Interesting and important questions also arise about the status and character of this mapping, which I will put aside.

If this line of reasoning is generally correct, then there is good reason to return to a traditional conception of language as “an instrument of thought,” and to revise Aristotle’s dictum accordingly; language is not sound with meaning but meaning with sound—more generally, with some form of externalization, typically sound though other modalities are readily available: work of the past generation on sign language has shown remarkable similarities to spoken language in structure, acquisition, and neural representation, though of course the mode of externalization is quite different.

It is worth noting that externalization is rarely used. Most language use by far is never externalized. It is a kind of internal dialogue, and the limited research on the topic, going back to some observations of Vygotsky’s,¹⁸ conforms to what introspection suggests—at least mine: what reaches consciousness is scattered fragments. Sometimes, full-formed expressions instantly appear internally, too quickly for articulators to be involved, and probably even instructions to them. This is an interesting topic that has been barely explored, but could be subjected to inquiry, and has many ramifications.

The latter issue aside, investigation of the design of language gives good reason to take seriously a traditional conception of language as essentially an instrument of thought. Externalization then would be an ancillary process, its properties a reflex of the largely or completely

¹⁸ Charles Fernyhough, “The Voices Within: The Power of Talking to Yourself,” *New Scientist*, 2919 (Jun. 3, 2013): 32–35.

independent sensorimotor system. Further investigation supports this conclusion. It follows that processing is a peripheral aspect of language, and that particular uses of language that depend on externalization, among them communication, are even more peripheral, contrary to virtual dogma that has no serious support. It would also follow that the extensive speculation about language evolution in recent years is on the wrong track, with its focus on communication.

It is, indeed, virtual dogma that the function of language is communication. A typical formulation of the idea is the following: "It is important that in a community of language users that words be used with the same meaning. If this condition is met it facilitates the chief end of language which is communication. If one fails to use words with the meaning that most people attach to them, one will fail to communicate effectively with others. Thus one would defeat the main purpose of language."¹⁹

It is, in the first place, odd to think that language has a purpose. Languages are not tools that humans design, but biological objects, like the visual or immune or digestive systems. Such organs are sometimes said to have functions, to be *for* some purpose. But that notion too is far from clear. Take the spine. Is its function to hold us up, to protect nerves, to produce blood cells, to store calcium, or all of the above? Similar questions arise when we ask about the function and design of language. Here evolutionary considerations are commonly introduced, but these are far from trivial; for the spine as well. For language, the various speculations about evolution typically turn to the kinds of communication systems found throughout the animal kingdom, but that it is just again a reflection of the modern dogma, and is likely to be a blind alley, for reasons already mentioned and to which I will return.

Furthermore, even insofar as language is used for communication, there is no need for meanings to be shared (or sounds, or structures). Communication is not a yes-or-no but rather a more-or-less affair. If similarities are not sufficient, communication fails to some degree, as in normal life.

Even if the term "communication" is largely deprived of substantive meaning and used as a cover term for social interaction of various kinds, it remains a minor part of actual language use, for whatever that observation is worth.

¹⁹ William Uzgalis, "John Locke," *The Stanford Encyclopedia of Philosophy* (Fall 2012 Edition), Edward N. Zalta, ed., URL = <<http://plato.stanford.edu/archives/fall2012/entries/locke/>>.

In brief, there is no basis for the standard dogma, and there is by now quite significant evidence that it is simply false. Doubtless language is sometimes used for communication, as is style of dress, facial expression and stance, and much else. But fundamental properties of language design indicate that a rich tradition is correct in regarding language as essentially an instrument of thought, even if we do not go as far as Humboldt in identifying the two.

The conclusion becomes even more solidly entrenched if we consider the Basic Property more closely. Naturally we seek the simplest account of the Basic Property, the theory with fewest arbitrary stipulations—each of which is, furthermore, a barrier to some eventual account of the origin of language. And we ask how far this resort to standard scientific method will carry us.

The simplest computational operation, embedded in some manner in every relevant computational procedure, takes objects X and Y already constructed and forms a new object Z . Call it *Merge*. The principle of Minimal Computation dictates that neither X nor Y is modified by Merge, and that they appear in Z unordered. Hence $\text{Merge}(X, Y) = \{X, Y\}$. That does not of course mean that the brain contains sets, as some current misinterpretations claim, but rather that whatever is going on in the brain has properties that can properly be characterized in these terms—just as we do not expect to find the Kekulé diagram for benzene in a test tube.

Note that if language really does conform to the principle of Minimal Computation in this respect, we have a far-reaching answer to the puzzle of why linear order is only an ancillary property of language, apparently not available for core syntactic and semantic computations: language design is perfect in this regard (and again we may ask why). Looking further, evidence mounts in support of this conclusion.

Suppose X and Y are merged, and neither is part of the other, as in combining *read* and *that book* to form the syntactic object corresponding to “read that book.” Call that case *External Merge*. Suppose that one is part of the other, as in combining $Y = \textit{which book}$ and $X = \textit{John read which book}$ to form *which book John read which book*, which surfaces as “which book did John read” by further operations to which I will return. That is an example of the ubiquitous phenomenon of displacement in natural language: phrases are heard in one place but interpreted both there and in another place, so that the sentence is understood as “for which book x , John read the book x .” In this case, the result of Merge of X and Y is again $\{X, Y\}$, but with two *copies* of Y ($= \textit{which book}$), one the original one remaining in X , the other the displaced copy merged with X . Call that *Internal Merge*.

It is important to avoid a common misinterpretation, found in the professional literature as well. There is no operation *Copy* or *Remerge*. Internal Merge happens to generate two copies, but that is the outcome of Merge under the principle of Minimal Computation, which keeps Merge in its simplest form, not tampering with either of the elements merged. New notions of Copy or Remerge are not only superfluous; they also cause considerable difficulties unless sharply constrained to apply under the highly specific conditions of Internal Merge, which are met automatically under the simplest notion of Merge.

External and Internal Merge are the only two possible cases of binary Merge. Both come free if we formulate Merge in the optimal way, applying to any two syntactic objects that have already been constructed, with no further conditions. It would require stipulation to bar either of the two cases of Merge, or to complicate either of them. That is an important fact. For many years it was assumed—by me too—that displacement is a kind of “imperfection” of language, a strange property that has to be explained away by some more complex devices and assumptions about UG. But that turns out to be incorrect. Displacement is what we should expect on the simplest assumptions. It would be an imperfection if it were lacking. It is sometimes suggested that External Merge is somehow simpler, and should have priority in design or evolution. There is no basis for that belief. If anything, one could argue that Internal Merge is simpler since it involves vastly less search of the workspace for computation—not that one should pay much attention to that.

Another important fact is that Internal Merge in its simplest form—satisfying the overarching principle of Minimal Computation—commonly yields the structure appropriate for semantic interpretation, as just illustrated in the simple case of “which book did John read.” However, these are the wrong structures for the sensorimotor system: universally in language, only the structurally most prominent copy is pronounced, as in this case; the lower copy is deleted. There is a revealing class of exceptions that in fact support the general thesis, but I will put that aside.²⁰

Deletion of copies follows from another uncontroversial application of Minimal Computation: compute and articulate as little as possible. The result is that the articulated sentences have *gaps*. The

²⁰ Tue Trinh, “A Constraint on Copy Deletion,” *Theoretical Linguistics*, xxxv, 2–3 (October 2009): 183–227. I also put aside here several topics that raise a variety of further questions, among them “covert operations” in which only the first-merged copy is externalized.

hearer has to figure out where the missing element is. As is well known in the study of perception and parsing, that yields difficult problems for language processing, so called *filler-gap* problems. In this very broad class of cases too, language design favors minimal computation, disregarding the complications in the processing and use of language.

Notice that any linguistic theory that replaces Internal Merge by other mechanisms has a double burden of proof to meet: it is necessary to justify the stipulation barring Internal Merge and also the new mechanisms intended to account for displacement—in fact, displacement with copies, generally the right forms for semantic interpretation.

The same conclusions hold in more complex cases. Consider for example the sentence “[which of his pictures] did they persuade the museum that [[every painter] likes best]?” It is derived by Internal Merge from the underlying structure “[which of his pictures] did they persuade the museum that [[every painter] likes [which of his pictures] best]?” formed directly by Internal Merge, with displacement and two copies. The pronounced phrase “which of his pictures” is understood to be the object of “likes,” in the position of the gap, analogous to “one of his pictures” in “they persuaded the museum that [[every painter] likes [one of his pictures] best].” And that is just the interpretation that the underlying structure with the two copies provides.

Furthermore, the quantifier-variable relationship between *every* and *his* carries over in “[which of his pictures] did they persuade the museum that [[every painter] likes best]?” The answer can be “his first one”—different for every painter, as in one interpretation of “they persuaded the museum that [[every painter] likes [one of his pictures] best].” In contrast, no such answer is possible for the structurally similar expression “[which of his pictures] persuaded the museum that [[every painter] likes flowers]?” in which case “his pictures” does not fall within the scope of “every painter.” Evidently, it is the unpronounced copy that provides the structure required for quantifier-variable binding as well as for the verb-object interpretation. The results once again follow straightforwardly from Internal Merge and copy deletion under externalization. There are many similar examples—along with interesting problems as complexity mounts.

Just as in the simpler cases, like “instinctively, eagles that fly swim,” it is inconceivable that some form of data processing yields these outcomes. Relevant data are not available to the language learner. The results must therefore derive “from the original hand of nature,” in Hume’s phrase—in our terms, from genetic endowment, specifically

the architecture of language as determined by UG in interaction with such general principles as Minimal Computation. In ways like these we can derive quite far-reaching and firm conclusions about the nature of UG.

One commonly reads claims in the literature that UG has been refuted, or does not exist. But this must be a misunderstanding. To deny the existence of UG—that is, of a biological endowment underlying the capacity for language—would be to hold that it is a miracle that humans have language but other organisms do not. The reference in these claims is presumably not to UG, however; rather, to descriptive generalizations—Joseph Greenberg’s very important proposals on language universals, for example. For example, in a foreword to the new edition of Quine’s *Word and Object*,²¹ Patricia Churchland, with an irrelevant citation, writes that “linguistic universals, long the darlings of theorists, took a drubbing as one by one they fell to the disconfirming data of field linguists.” Presumably she takes this to be confirmation of Quine’s view that “Timely reflection on method and evidence should tend to stifle much of the talk of linguistic universals,” meaning generalizations about language. In reality, it is field linguists who have discovered and confirmed not only the generally valid and quite important generalizations, but also the invariant properties of UG. The term “field linguists” means linguists concerned with data, whether they are working in the Amazon jungle, or in their offices in Belem, or in New York.

The fragment of truth in such observations is that generalizations are likely to have exceptions, which can be quite valuable as a stimulus to inquiry: for example, the exceptions to deletion of copies, which I just mentioned. That is a common experience in the sciences. The discovery of perturbations in the orbit of Uranus did not lead to the abandonment of Newton’s principles and Kepler’s laws, or to the broader conclusion that there are no physical laws, but to the postulation—later discovery—of another planet, Neptune. Exceptions to largely valid descriptive generalizations play a similar role quite generally in the sciences, and have done so repeatedly in the study of language.

There is, then, persuasive and quite far-reaching evidence that if language is optimally designed, it will provide structures appropriate for semantic interpretation but that yield difficulties for perception and language processing (hence communication). There are many

²¹ Patricia S. Churchland, foreword, W. V. O. Quine, *Word and Object* (1960; repr., Cambridge: MIT, 2013), p. xiii.

other illustrations. Take, say, passivization. It has been argued that passivization supports the belief that language is well-designed for communication. Thus in the sentence “the boys took the books,” if we wish to foreground “the books,” the passive operation allows us to do so by saying “the books were taken by the boys.” In fact, the conclusion is the opposite. The design of language, following from Minimal Computation, regularly bars this option. Suppose in the sentence “the boys took the books from the library” we wish to foreground “the library,” yielding “the library was taken the books from by the boys.” That is barred by language design, yet another barrier to communication.

The interesting cases are those in which there is a direct conflict between computational and communicative efficiency. In every known case, the former prevails; ease of communication is sacrificed. Many such cases are familiar, among them structural ambiguities and “garden path sentences” such as “the horse raced past the barn fell,” interpreted as ungrammatical on first presentation. Another case of particular interest is so-called *islands*—constructions in which extraction (Internal Merge) is barred—insofar as these can be given principled explanations invoking computational efficiency. An illustration is the questions associated with the expression: “they asked if the mechanics fixed the cars.” We can ask “how many cars,” yielding “how many cars did they ask if the mechanics fixed?” Or we can ask “how many mechanics,” yielding “how many mechanics did they ask if fixed the cars?” The two interrogatives differ sharply in status: asking “how many mechanics” is a fine thought, but it has to be expressed by some circumlocution, again impeding communication; technically an ECP (Empty Category Principle) violation. Here too there appear to be counterexamples, in Italian for example. Recognition of these led to discoveries about the nature of null subject languages by Luigi Rizzi,²² reinforcing the ECP principle, again illustrating the value of proposed generalizations and apparent exceptions.

There are many similar cases. Insofar as they are understood, the structures result from free functioning of the simplest rules, yielding difficulties for perception and language processing. Again, where ease of processing and communicative efficiency conflict with computational efficiency in language design, in every known case the former are sacrificed. That lends further support to the view of language as an instrument of thought, in interesting respects perfectly designed, with externalization an ancillary process, hence a fortiori communication and other uses of externalized language. As is often

²² Luigi Rizzi, *Issues in Italian Syntax* (Dordrecht: Foris, 1982).

the case, what is actually observed gives quite a misleading picture of the principles that underlie it. The essential art of science is reduction of “complex visibles to simple invisibles,” as Nobel laureate in chemistry Jean Baptiste Perrin put the matter.

To bring out more clearly just what is at stake, let us reverse the argument outlined here, putting it in a more principled way. We begin with the Basic Property of language, and ask what the optimal computational system would be that captures it, adopting normal scientific method. The answer is *Merge* in its simplest form, with its two variants, External and Internal Merge, the latter yielding the “copy theory of movement.” In a wide and important range of cases, that yields forms appropriate for semantic interpretation at the conceptual-intentional interface, forms which lack order or other arrangements. An ancillary process of externalization then converts the internally generated objects to a form adapted to the sensorimotor system, with arrangements that vary depending on the sensory modality for externalization. Externalization too is subject to Minimal Computation, so that copies are erased, yielding difficulties for language processing and use (including the special case of communication). A fallout of the optimal assumptions is that rules are invariably structure-dependent, resolving the puzzle discussed at the outset and others like it.

A broader research project—in recent years called *the minimalist program*—is to begin with the optimal assumption—the so-called *strong minimalist thesis* (SMT)—and to ask how far it can be sustained in the face of the observed complexities and variety of the languages of the world. Where a gap is found, the task will be to see whether the data can be reinterpreted, or principles of optimal computation can be revised, so as to solve the puzzles within the framework of SMT, thus producing some support, in an interesting and unexpected domain, for Galileo’s precept that nature is simple, and it is the task of the scientist to prove it. The task is of course a challenging one. It is fair to say, I think, that it seems a good deal more realistic today than it did only a few years ago, though enormous problems of course remain.

All of this raises at once a further question: Why should language be optimally designed, insofar as the SMT holds? This question leads us to consideration of the origin of language. The SMT hypothesis fits well with the very limited evidence we have about the emergence of language, apparently quite recently and suddenly in the evolutionary time scale, as Tattersall discussed. A fair guess today—and one that opens rich avenues of research and inquiry—is that some slight rewiring of the brain yielded Merge, naturally in its

simplest form, providing the basis for unbounded and creative thought, the “great leap forward” revealed in the archaeological record, and the remarkable differences separating modern humans from their predecessors and the rest of the animal kingdom. Insofar as the surmise is sustainable, we would have an answer to questions about apparent optimal design of language: that is what would be expected under the postulated circumstances, with no selectional or other pressures operating, so the emerging system should just follow laws of nature, in this case the principles of Minimal Computation—rather the way a snowflake forms.

These remarks only scratch the surface. Perhaps they can serve to illustrate why the answer to the question “What is language?” matters a lot, and also to illustrate how close attention to this fundamental question can yield conclusions with many ramifications for the study of what kind of creatures humans are.