

PAPER • OPEN ACCESS

The Galilean Challenge: Architecture and Evolution of Language

To cite this article: Noam Chomsky 2017 *J. Phys.: Conf. Ser.* **880** 012015

View the [article online](#) for updates and enhancements.

You may also like

- [The Molecular Gas Reservoirs of \$z \sim 2\$ Galaxies: A Comparison of CO\(10\) and Dust-based Molecular Gas Masses](#)
M. Kaasinen, N. Scoville, F. Walter et al.
- [Chemical and Kinematic Properties of the Galactic Disk from the LAMOST and Gaia Sample Stars](#)
Yepeng Yan, Cuihua Du, Shuai Liu et al.
- [Probing the Pulsar Population of Terzan 5 via Spectral Modeling](#)
H. Ndiyavala, C. Venter, T. J. Johnson et al.



ECS
The
Electrochemical
Society
Advancing solid state &
electrochemical science & technology

DISCOVER
how sustainability
intersects with
electrochemistry & solid
state science research

The Galilean Challenge: Architecture and Evolution of Language

Noam Chomsky

Linguistics & Philosophy, MIT, Cambridge MA 02139, USA

E-mail: nchomsky3@gmail.com

In the early days of the modern scientific revolution, Galileo and his contemporaries issued a crucial challenge to those concerned with the nature of human language, a challenge that was scarcely recognized until it was taken up in the mid-20th century and became the primary concern of much of the study of language. For short, I will refer to it as the Galilean challenge. These great founders of modern science expressed their awe and wonder at the fact that language permits us to construct “from 25 or 30 sounds an infinite variety of expressions, which although not having any resemblance in themselves to that which passes through our minds, nevertheless do not fail to reveal all of the secrets of the mind, and to make intelligible to others who cannot penetrate into the mind all that we conceive and all of the diverse movements of our souls.”¹

We can now see that the Galilean challenge requires some qualifications, but it is very real, and should I think be recognized as one of the deepest insights in the rich history of inquiry into language and mind in the past 2500 years.

The challenge was not entirely ignored. For Descartes, the human capacity for unbounded and appropriate use of language was a primary basis for his postulation of mind as a new creative principle. In later years, there is occasional recognition that language is a creative activity that involves “infinite use of finite means,” in Wilhelm von Humboldt’s formulation, that it provides “audible signs for thought,” in the words of linguist William Dwight Whitney a century ago. There has also been awareness that these capacities are a species-property, shared by humans and unique to them, the most striking feature of this curious organism and a foundation for its remarkable achievements. But there was never much to say beyond a few phrases.

There is a good reason why the insights languished until mid-20th century: intellectual tools were not available for even formulating the problem in a clear enough way to address it seriously. That changed thanks to the work of Alan Turing and other great mathematicians who established the general theory of computability on a firm basis, showing in particular how a finite object like the brain can generate an infinite variety of expressions. It then became possible, for the first time, to address at least part of the Galilean challenge directly – although, regrettably, the earlier history was entirely unknown at the time.

¹ The quote is from Antoine Arnauld and Claude Lancelot, *Grammaire générale et raisonnée* (1660). Galileo’s earlier version, in his *Dialogo sopra i due massimi sistemi del mondo* (1632) was similar, except for referring to the alphabet rather than the sounds of language.



With these intellectual tools available, it becomes possible to formulate what we may call the Basic Property of human language: the language faculty of the human brain provides the means to construct a digitally infinite array of structured expressions, each of which has a semantic interpretation expressing a thought, and each of which can be externalized by means of some sensory modality. The infinite set of semantically interpreted objects constitutes what has sometimes been called a language of “thought”: the system of thoughts that receive linguistic expression and that enter into reflection, inference, planning, and other mental processes, and when externalized can be used for communication and other social interactions. We may fairly assume that the language faculty is shared among humans. There are no known group differences in language capacity, and individual variation is at the margins.

We now know that although speech is the usual form of sensorimotor externalization, it can just as well be sign or even touch, discoveries that require a slight reformulation of the Galilean challenge. A more fundamental qualification has to do with the way the challenge is formulated: in terms of production of expressions. So formulated, the challenge overlooks some basic issues. Production, like perception, accesses the internal language but cannot be identified with it. We must distinguish the internalized system of knowledge from the actions that access it. The theory of computability enables us to establish the distinction, which is an important one, familiar in other domains.

Consider for example human arithmetical competence. In studying it, we routinely distinguish the internal system of knowledge from the actions that access it, like multiplying numbers in our head, an action that involves many factors beyond intrinsic knowledge; memory constraints, for example. The same is true of language. Production and perception access the internal language, but involve other factors as well, including again short-term memory, matters that began to be studied with some care as soon as the Galilean challenge was addressed in the 1950s.

There has been considerable progress in understanding the nature of the internal language, but its free creative use remains a mystery. That comes as no surprise. In a recent review of the state of the art concerning far simpler cases of voluntary action, two leading researchers, neuroscientists Emilio Bizzi and Robert Ajemian, write that we are beginning to learn something about the puppet and the strings, but the puppeteer remains shrouded in mystery.² That is even more dramatically true for such creative acts as the normal use of language, the unique human capacity that so impressed the founders of modern science.

The fundamental task of inquiry into language is to determine the nature of the Basic Property. To the extent that its properties are understood, we can seek to investigate particular internal languages, each an instantiation of the Basic Property, much as each individual visual system is an instantiation of the human faculty of vision. We can investigate how the internal languages are acquired and used, how the language faculty itself evolved, its basis in human genetics and the ways it functions in the human brain. This general program of research has been called the Biolinguistic Program. The theory of the genetically based language faculty is called Universal Grammar; the theory of each individual language is called its Generative Grammar.

Languages appear to be extremely complex, varying radically from one another. And indeed, a standard belief among professional linguists 60 years ago was that languages can vary in arbitrary ways and each must be studied without preconceptions. Similar views were held at the time about organisms generally. Many biologists would have agreed with molecular biologist

² Emilio Bizzi and Robert Ajemian, “A Hard Scientific Quest: Understanding Voluntary Movements,” *Daedalus* 144.1 (2015).

Gunther Stent's conclusion that the variability of organisms is so free as to constitute "a near infinitude of particulars which have to be sorted out case by case."³ When understanding is thin, we expect to see extreme variety and complexity.

A great deal has been learned since then. Within biology, it is now recognized that the variety of life forms is very limited, so much so that the hypothesis of a "universal genome" has been seriously advanced. My own feeling is that linguistics has undergone a similar development, along lines that I will outline.

The Basic Property takes language to be a computational system, which we therefore expect to observe general conditions on computational efficiency. A computational system consists of a set of atomic elements and rules to construct more complex ones. For generation of the language of thought, the atomic elements are word-like, though not words; for each language, the set of these elements is its lexicon. The lexical items are commonly regarded as cultural products, varying widely with experience and linked to extra-mental entities – an assumption expressed in the titles of standard works, such as W.V. O. Quine's influential study *Word and Object*. Closer examination reveals a very different picture, one that poses many mysteries. Let's put that aside for now, turning to the computational procedure.

Clearly, we will seek the simplest computational procedure consistent with the data of language, for reasons that are implicit in the basic goals of scientific inquiry. It has long been recognized that simplicity of theory translates directly to explanatory depth. A more specific version of this quest for understanding was provided by a famous dictum of Galileo's, which has guided the sciences since their modern origins: nature is simple, and it is the task of the scientist to demonstrate this, from the motion of the planets, to an eagle's flight, to the inner workings of a cell, to the growth of language in the mind of a child. Linguistics has an additional motive of its own for seeking the simplest theory: it must face the problem of evolvability. Not a great deal is known about evolution of modern humans, but the few facts that are well established, and others that have recently been coming to light, are rather suggestive, and conform well to the conclusion that the language faculty is near optimal for a computational system, the goal we should seek on purely methodological grounds.

One fact that does appear to be well established I have already mentioned: that the faculty of language is a true species property, invariant among human groups – and, furthermore, unique to humans in its essential properties. It follows that there has been little or no evolution of the faculty since human groups separated from one another. Recent genomic studies place this date not very long after the appearance of anatomically modern humans about 200,000 years ago, perhaps some 50,000 years later, when the San group in Africa separated from other humans. There is no evidence of anything like human language, or symbolic activities altogether, before the emergence of modern humans. That leads us to expect that the faculty of language emerged along with modern humans or not long after, a very brief moment in evolutionary time. It follows, then, that the Basic Property should indeed be very simple. The conclusion conforms to what has been discovered in recent years about the nature of language, a welcome convergence.

The discoveries about early separation of the San people are highly suggestive. They appear to share the general human language capacity, but have significantly different externalized languages. With irrelevant exceptions, their languages are all and only the languages with phonetic clicks, with corresponding adaptations in the vocal tract. The most likely explanation for these facts, developed in detail in current work by Dutch linguist Riny Huijbregts,⁴ is that

³ Gunther Stent, "From Probability to Molecular Biology," *Cell* 36 (1984).

⁴ Riny Huijbregts, "Phonemic Clicks and the Mapping Asymmetry: How Language Emerged and Speech Developed" <<http://www.sciencedirect.com/science/article/pii/S0149763416305450>>, *Neuroscience & Biobehavioral Reviews* (2017).

possession of internal language preceded separation, which in turn preceded externalization, the latter in somewhat different ways in separated groups. Externalization seems to be associated with the first signs of symbolic behavior in the archaeological record, after the separation. Putting these observations together, it seems that we are reaching a stage in understanding where the account of evolution of language can perhaps be fleshed out in ways that were unimaginable until quite recently.

Returning to the Basic Property, as we have seen we have reason to believe that it may be quite simple. The challenge for research, then, is to show how the facts of language are accounted for in terms of the Basic Property: more fully, by the interaction of the Basic Property, specific experience, and language-independent principles, including principles of computational efficiency. The challenge is of particular interest and significance when it is clear that the child's experience provides little or no relevant evidence – a situation far more prevalent than commonly realized, so careful examination reveals, from acquisition of word meaning on to syntactic structures and the semantic properties of the generated language of thought.

Also of particular interest are the universal properties of the language faculty that began to come to light as soon as serious efforts were undertaken to construct generative grammars, including quite simple principles that had never been noticed, and that are quite puzzling. One crucial and puzzling principle is structure-dependence: the rules that yield the language of thought attend solely to structural properties, ignoring properties of the externalized signal, even such simple properties as linear order.

The property is illustrated by elementary examples. Consider the sentence “John and his father are tall” – not *is tall*, though the bigram frequency of *father-is* is much greater than of *father-are*, and the computation using linear order (adjacency) is far simpler than the computation that has to analyze the sentence into phrases and use phrasal locality.

To take an example of semantic construal, consider the sentence *birds that fly instinctively swim*. It is ambiguous: the adverb “instinctively” can be associated with the preceding verb (*fly instinctively*) or the following one (*instinctively swim*). Suppose now that we extract the adverb from the sentence, forming *instinctively, birds that fly swim*. Now the ambiguity is resolved: the adverb is construed only with the linearly more remote but structurally closer verb *swim*, not the linearly closer but structurally more remote verb *fly*. The only possible interpretation – *birds swim* – is the unnatural one, but that doesn't matter: the rules apply rigidly, independent of meaning and fact. What is puzzling, again, is that the rules ignore the simple computation of linear distance and keep to the far more complex computation of structural distance.

The principle of structure dependence holds for all constructions in all languages, and it is indeed puzzling. Furthermore, it is known without relevant evidence, as is clear in cases like the one I just gave and innumerable others. Experiment shows that children understand that rules are structure-dependent as early as they can be tested, by about age 3 – and are of course not instructed. We can be quite confident, then, that structure-dependence follows from principles of universal grammar that are deeply rooted in the human language faculty.

There is evidence from other sources that supports the conclusion that structure-dependence is a true linguistic universal, deeply rooted in language design. Research conducted in Milan a decade ago, initiated by Andrea Moro, showed that invented languages keeping to the principle of structure-dependence elicit normal activation in the language areas of the brain, but much simpler systems using linear order in violation of these principles yield diffuse activation, implying that experimental subjects are treating them as a puzzle, not a language.⁵ Similar

⁵ Mariacristina Musso et al., “Broca's Area and the Language Instinct” <<http://www.nature.com/neuro/journal/v6/n7/abs/n1077.html>>, *Nature Neuroscience* 6.7 (2003).

results were found in work by Neil Smith and Ianthi Tsimpli in their investigation of a cognitively deficient but linguistically gifted subject. They also made the interesting observation that normals can solve the problem if it is presented to them as a puzzle, but not if it is presented as a language, presumably activating the language faculty.

The only plausible conclusion, then is that structure-dependence is an innate property of the language faculty. That raises the question why this should be so? There is only one known answer, and fortunately, it is the answer we seek for general reasons: the computational operations of language are the simplest possible ones. Again, that is the outcome that we hope to reach on methodological grounds, and that is to be expected in the light of the evidence about evolution of language already mentioned.

The simplest recursive operation, embedded in one or another way in all others, takes two objects already constructed, say X and Y , and forms a new object Z , without modifying either X or Y or adding any further structure. Accordingly, Z can be taken to be just the set $\{X, Y\}$. In current work, the operation is called Merge. Since Merge imposes no order, the objects constructed, however complex, will be hierarchically structured but unordered, and operations on them will necessarily keep to structural distance, ignoring linear distance. It follows that the linguistic operations yielding the language of thought will be structure-dependent, as indeed is the case, resolving the puzzle.

Externalization of language maps internal structures into some sensorimotor modality, usually speech. The sensorimotor system requires linear order; we cannot speak in parallel. But none of this enters the generation of the language of thought, which keeps to structural relations. More generally, externalization of language seems to be a peripheral aspect of language, not entering into its core function of providing a language of thought, contrary to a long tradition, including the formulation of the Galilean challenge.

Perception yields further evidence in support of this conclusion. The auditory systems of apes are quite similar to those of humans, even attuned to the phonetic features that are used in language. But the shared auditory-perceptual systems leave apes without anything remotely like the human faculty of language. Some have attributed this lack to deficiency of articulatory apparatus and vocal learning, but apes can gesture quite easily and as has been known for many years, sign language is virtually identical to spoken language in its basic properties and acquisition, and human groups (including children), under the right conditions, even invent normal sign languages with no linguistic input at all.

These results support the conclusion that internal language is independent of externalization, and that internal language evolved quite independently of the process of externalization.

The most powerful evidence, however, is what we learn from investigation of language design. One crucial example is the one just mentioned: the explanation for the puzzle posed by the linguistic universal of structure-dependence, which follows from the null hypothesis: that the computational system is optimal, and hence ignores linear order, the most elementary feature of externalization.

Not long ago it would have seemed absurd to propose that the operations of human language could be reduced to Merge, along with language-independent principles of computational efficiency. But work of the past few years has shown that quite intricate properties of language follow directly from this assumption, along with a few other quite simple ones. One important result has to do with the property of displacement, a ubiquitous and also quite puzzling property of language: phrases are heard in one position but interpreted both there and in some other position that is silent but where they could have occurred – a puzzling property, which is never built into artificial symbolic systems for metamathematics, programming, or other purposes.

For example, the sentence “which book will you read?” is interpreted as meaning roughly: “for which book x , you will read the book x ,” with the nominal phrase *book* heard in one position but interpreted in two positions.

I will not go into the details, but it is easy to show that Merge-based computation automatically yields displacement with copies; in this case, two copies of *which book*, yielding the correct semantic interpretation directly. The same process yields quite intricate semantic interpretations, and also has significant implications about the nature of language. To see why, consider for example the sentence “the boys expect to see each other” and the same sentence preceded by “which girls”: “which girls do the boys expect to see each other.” In the latter sentence, “each other” does not refer back to the closest antecedent, “the boys,” as such phrases universally do, but rather to the more remote antecedent “which girls.” The sentence means “for which girls the boys expected those girls to see each other.” That is what reaches the mind, under Merge-based computation with automatic copies, though not what reaches the ear. What reaches the ear violates the locality condition of referential dependency. Deletion of the copy in externalization causes processing problems: such filler-gap problems, as they are called, can be become quite severe, and are among the major problems of automatic parsing and perception. If the copy were not deleted, the problem would not arise. Why then is it deleted? Again, because of principles of efficient computation that reduce what is computed to the minimum: at least one copy must appear or there is no evidence that displacement took place at all, so only the structurally most prominent one remains (with important qualifications strengthening the conclusion, which I will put aside), leaving a gap that must be filled by the hearer – a matter that can become quite intricate.

These examples illustrate a general phenomenon of some significance. Language design appears to maximize computational efficiency but disregards communicative efficiency. In fact, in every known case in which computational and communicative efficiency conflict, communicative efficiency is ignored. These facts argue against the common belief that communication is the basic function of language. They also further undermine continuity assumptions about language evolving from animal communication. And they provide further evidence that externalization, which is necessary for communication, is a peripheral aspect of language.

As I mentioned, there are methodological reasons and also some evolutionary reasons to expect that the basic design of language will be quite simple, perhaps even close to optimal. With regard to externalization of language, the same methodological arguments hold, but the evolutionary arguments do not apply. In fact, externalization of language may not involve evolution at all. The sensorimotor systems were in place long before the appearance of language. Mapping the internal language to some sensorimotor system for externalization is a hard cognitive problem, relating two systems that are unrelated: an internal system that may be highly efficient computationally, and a sensory modality unrelated to it. That would lead us to expect that the variety, complexity, and easy mutability of observed languages might lie primarily in externalization. Increasingly, it seems clear that that is the case. And in fact it should be expected, since the principles of the internal language are largely known by children without evidence, as, indeed, is a great deal more about language, including almost all semantic and most syntactic properties – a matter of contention, but solidly established, I think.

Let us return finally to the second component of a computational system, the atomic elements: for language, the lexical items. As I mentioned, the conventional view is that these are cultural products, and that the basic ones – those used for referring to the world – are associated with extra-mental entities. This representationalist doctrine has been almost universally adopted in the modern period. The doctrine does appear to hold for animal communication: a monkey’s calls, for example, are associated with specific physical events. But the doctrine is radically false for human language, as was recognized as far back as classical Greece.

To illustrate, let's take the first case that was discussed in pre-Socratic philosophy, the problem posed by Heraclitus: how can we cross the same river twice? To put it differently, why are two appearances understood to be two stages of the same river? Contemporary philosophers have suggested that the problem is solved by taking a river to be a four-dimensional object, but that simply restates the problem: why this object and not some different one, or none at all?

When we look into the question, puzzles abound. Suppose that the flow of the river has been reversed. It is still the same river. Suppose that what is flowing becomes 95% arsenic because of discharges from an upstream plant. It is still the same river. The same is true of other quite radical changes in the physical object. On the other hand, with very slight changes it will no longer be a river at all. If its sides are lined with fixed barriers and it is used for oil tankers, it is a canal, not a river. If its surface undergoes a slight phase change and is hardened, a line is painted down the middle, and it is used to commute to town, then it is a highway, no longer a river. Exploring the matter further, we discover that what counts as a river depends on mental acts and constructions. The same is true quite generally of even the most elementary concepts: *tree, water, house, person, London*, or in fact any of the basic words of human language. Unlike animals, the items of human language and thought uniformly violate the representationalist doctrine.

For such reasons, incidentally, deep-learning approaches to object-recognition, whatever their interest, cannot in principle discover the meanings of words.

Furthermore, the intricate knowledge of the means of even the simplest words, let alone others, is acquired virtually without experience. At peak periods of language acquisition, children are acquiring about a word an hour, that is, often on one presentation. It must be, then, that the rich meaning of even the most elementary words is substantially innate. The evolutionary origin of such concepts is a complete mystery, one that may not be resolvable by means available to us.

Returning to the Galilean challenge, it has to be reformulated to distinguish language from speech, and to distinguish production from internal knowledge, the latter an internal computational system that yields a language of thought, a system that might be remarkably simple, conforming to what the evolutionary record suggests. Secondary processes map the structures of language to one or another sensorimotor system for externalization. These processes appear to be the locus of the complexity and variety of linguistic behavior, and its mutability over time.

There are suggestive recent ideas about the neural basis for the operations of the computational system, and about its possible evolutionary origins. The origin of the atoms of computation, however, remains a complete mystery, as does a major question that concerned those who formulated the Galilean challenge, the Cartesian question of how language can be used in the normal creative way, in a manner appropriate to situations but not caused by them, in ways that are "incited and inclined" but not "compelled," in Cartesian terms. The mystery holds for even the simplest forms of voluntary motion, as discussed earlier.

A great deal has been learned about language since the Biolinguistic Program was initiated. It is fair to say, I think, that during these years more has been learned about the nature of language, and about a very wide variety of typologically different language, than in the entire 2500 year history of inquiry into language. New questions have arisen, some quite puzzling. And there are some surprising answers, which lead us to revise what has long been believed about language and mental processes generally.

But as is familiar in the sciences, the more we learn, the more we discover what we do not know. And the more puzzling it often seems.