

Language and the Cognitive Science Revolution(s)

Noam Chomsky

Text of lecture given at the Carleton University, April 8, 2011

(Partial transcript courtesy of David P. Wilkins)

...

My own feeling is that the major cognitive science revolution was in the seventeenth century.

...

And, in some respects, what was beginning to be discovered back in the seventeenth century has been mostly lost, particularly in such crucial areas as the theory of meaning of natural language.

...

Well the title of the talk raises a more substantive question. The phrase 'Language and the Cognitive Science Revolutions' presupposes that language exists. That is, that it exists as an independent object of serious study and not as just some arbitrary collection of various phenomena and processes more or less like, say, today's weather – there's no science of today's weather, although there may be of the many factors that enter into it.

Well, that language exists in a serious sense might seem uncontentious, but it actually is very far from that. A major tendency in the cognitive sciences, probably the dominant one, holds that it does not exist and sometimes states that forcefully and quite explicitly. It's implicit in much of the other work.

It's worth remembering that there's nothing new in that stance. So, fifty years ago it was widely held by the most prominent philosophers and psychologists that language does not exist in any serious sense. It's just a matter of conditioning and some obscure notion of induction and analogy.

Well, in professional linguistics, a dominant view fifty years ago and more was that languages can differ without limits, apart from very restricted constraints like maybe choice of phonetic features which could be properties of the articulatory apparatus.

...

A further elaboration of this is that linguistic theory, to the extent that it exists at all, was just a kind of battery of procedures to reduce a body of data, called a corpus, to some organised form in one way or another, and how you do it depends on your particular purposes, but there's nothing real about it – there's no truth to the matter, there's no serious study.

Later versions of the non-existence hypothesis within philosophy were that rules of language can be justifiably postulated only if they're in principle accessible to introspection – a dogma that's largely incoherent in my opinion.

...

There are other variants, among them quite prominently the insistence, again by quite prominent philosophers and others, that language must be regarded as a socio-political entity of some kind.

...

A corollary, in the mid-twentieth century and since, is that there are no real problems to solve in the study of language – if it doesn't exist as a serious topic, can't have any problems. The only thing to do is accumulation and organisation of data, maybe what Wittgenstein called an 'assembly of particulars'.

Actually, I was a student in the nineteen-forties, and our view was – the view among students was, this is kind of fun, but what happens when we publish the structural analysis of every language, there's only a finite number of them, using the methods of analysis that were promulgated, which were assumed to be more or less settled, apart from some details, and which, in principle, could be automated.

...

The view was pretty sensible, given the non-existence assumption.

Actually similar views are quite prevalent today. Namely, that the only real object is a corpus of data and that by automated analysis, you know, complicated Bayesian statistics and so on, one can derive everything that's relevant about the language.

Well, the announced topic, Language and other cognitive systems, cognitive revolutions, that not only presupposes the existence of language but also presupposes a modular approach to the mind, taking it to be much like the rest of the organism.

That is, some complex of subsystems, often informally called organs, with enough internal integrity so that it makes sense to study each in abstraction from the others with which it is integrated and with which it interacts in the life of the organism. So, for example, the visual, immune, digestive and other organs below the neck, metaphorically speaking, and the various so-called mental organs, language, planning, various structures of memory, organisation of action, and so on, whatever the right analysis turns out to be.

The neuroscientist, Randi Gallistel, has observed that "the biological norm is modular systems with special" – I'm quoting him – "with special growth and learning mechanisms in different domains and in different species." And I think that's correct, and I think there's every reason to expect human language to keep to the biological norm in this respect. Even more so, because there are apparently crucial features of human language that appear to be quite isolated in the biological world – nothing homologous, in fact nothing closely analogous.

...

Well, my own assumption is that language does exist as a module of the mind and body, mostly the brain, but that the non-existence approach in its contemporary form in the cognitive sciences is actually raising the right questions, although pursuing them in a manner that's almost guaranteed to fail. At least as success and failure has been understood for centuries in the sciences.

...

A look at the literature, by now quite large literature, on the evolution of language reveals that it barely even addresses the topic, instead it consists almost entirely of speculations on evolution of communication, which is a totally different matter. I mean, you can use language for communication, you can use other things for communication, most language use is not for communication. So, they're just different topics.

...

There happens to be one recent essay which very neatly encapsulates clearly the main assumptions of the contemporary non-existence approach to language and its evolution, in a prestigious journal Science magazine,

journal of the AAAS. So, I'll use that as an example, but it's just one illustration of many. This is a review article discussing books on evolution of language by a scholar at the Max Planck Institute, N.J. Enfield.

He finds essentially nothing of value in the books that are reviewed. And the reason, he says, is that they are all lethally tainted by the existence assumption – that is, the belief that there is something to study, the belief that there are rule systems of some kind that determine what's a sentence, form-meaning relations, conditions of language use, and so on. To illustrate the fallacy of the existence approach, the article is accompanied by a photograph of three infants, suitably interracial and apparently noticing one another if you look at the picture, and the caption reads 'communication without syntax'. The point is to show that rule systems of the kind studied under the existence assumption aren't necessary for communication. Actually, they could have given a picture of three bacteria, that would have shown the point just as well – this is both the author and the editors, that's their idea.

The title of the article is 'Without Social Context?' with a question mark. The question mark expresses the surprise at the fallacy of the existence approach because it ignores social context when it seeks to determine the functioning rules that allegedly, say, relate form and meaning, determine conditions of language use and so on.

Well, again, to try to make the matter concrete, let's take an actual example, it's simple enough so we don't need a blackboard. So, just take the sentence:

'He wondered whether the mechanics fixed the cars.'

And just consider two questions that you could ask about that.

The two questions are:

'How many cars did he wonder whether the mechanics fixed?'

Answer: 'Three cars'

'How many mechanics did he wonder whether fixed the cars?'

Answer: 'Three mechanics'

But the trouble is, you can't say that for some reason. You can't ask: 'How many mechanics did he wonder whether fixed the cars?'

In technical terms it's called an ECP violation, but it doesn't matter what that is.

The thought is fine – fine thought – but you have to express it in some kind of paraphrase. There's something about the language design which poses a barrier to communication. You just can't express a simple thought like that, you need a circumlocution. Many such examples.

Well, to investigate such questions as these, according to, say, the editors of Science and the author of the article, you have to consider the social context of the use of these expressions, the normal use. That's very quickly finished, there is none, so you're done. But it's a mistake to raise the question anyway, from this point of view, because the sentences are constructed. They're constructed as an experiment, they're not drawn from a massive corpus of data, where you'd never find anything like this. That is, the inquiry proceeds by the methods of the sciences as they've been understood since antiquity, but that's another serious flaw, you're not allowed to do that.

The observation about social context is actually uncontroversial with regard to communication – you can't study communication without social context, that's virtual tautology. So, of course, the study of communication takes it

into account. But it's also uncontroversial that the study of the mechanisms that we put to use in action, whatever it is, that study typically ignores social context and quite rightly so. For example, for those of you who know this work, the classic work on neurophysiology of vision, say, Hubel and Wiesel's work, for which they got the Nobel Prize.

...

Or in fact virtually all of the fundamental work that aims to determine the properties of the modules of cognition at whatever level it's conducted whether it's neurophysiological, behavioural, perceptual, whatever, - it ignores social context totally, just following the normal methods of the sciences. However, we're instructed that the study of mechanisms used, say, in the examples I mentioned, these ECP examples, or the study of, for example, vowel harmony in Turkish, or of the relative scope of operators, or, in fact, everything else about language has to depart from the scientific norm. That's a principle. It cannot follow the methods of the sciences.

Well, this kind of critique, which is quite widespread, is, in fact, accompanied by a novel concept of science that has emerged in the computational cognitive sciences and related areas of linguistics. With this new notion of science, which is all over the literature, an account of some phenomena is taken to be successful to the extent that it approximates unanalysed data.

...

The major cognitive science journals, and general journals like *Science*, regularly publish articles triumphantly listing dramatic failures which are called successes because they accord with this new concept, which is unique in the history of the sciences and very radically restricted, in fact, almost specifically to language. So, nobody would suggest it for physics or bee communication or almost anything else, because it's so obviously absurd that people would just laugh. In fact, it's not even suggested for systems as close to a language as arithmetical capacity. So, you don't study arithmetical capacity by constructing models based on a statistical analysis of masses of observations of what happens when, say, people try to multiply numbers in their heads without external memory. At least, I hope nobody does that.

Enfield, in the same article, he also puts forth a far-reaching thesis which is quite standard in the cognitive sciences and a very clear expression of the non-existence hypothesis, I'll quote him. He says: "Language is entirely grounded in a constellation of cognitive capacities that each, taken separately, has other functions as well." Notice, that's kind of an updating of the nineteen-fifties position that I quoted. Well, that means language exists only in the sense that there exists such a thing as today's weather, which is also a constellation of many factors that operate independently.

...

There's another influential version of the idea that language doesn't exist. It's sort of highly dominant in language acquisition studies and the leading figure is Michael Tomasello. So, in a recent handbook of child development he explains that there aren't any linguistic rules and there's nothing to say about descriptive regularities, say, like those ECP examples. Rather, there's nothing at all except a structured inventory of meaningful linguistic constructions, all of them meaningful linguistic symbols that are used in communication. That's his topic, there being no such thing as language. The inventory is structured only in the sense that its elements – words, idioms, sentences like the one I'm now speaking – they're all acquired by processes of pattern finding, schematization and abstraction that are common to all primates. A few other processes, all left quite obscure. So, in other words, these ECP examples that I mentioned, according to this story, are learned just the way a child learns 'horse', or an idiom like 'how do you do' or, say, 'kick the bucket' meaning 'die', and so on, or this sentence, they're all learned exactly the same way. And the child somehow learns that the ECP violation is not usable for

communication, even though the thought is fine, although the other expression somehow is. And presumably, the expressions could have virtually any other properties in the next language you look at. In fact, the inventory, as in the 1950's versions is essentially an arbitrary collection of unanalysed linguistic symbols and it's also finite, just like Quine's pattern, apart from some hand-waving. In fact, I can think of no rational interpretation for any of this, but it's overwhelmingly dominant in the fields, you might think about it.

Enfield also presents a closely related thesis, that's also very widely held, I'll quote it: "There are well-developed gradualist evolutionary arguments to support the conclusion that there's no such thing as language, except as an arbitrary complex of independent cognitive processes." Again, no relevant sources cited, and none exist.

...

Well, the fact is, we know very little about evolution of language, maybe we never will, but there are a few fairly clear conclusions, and they're kind of suggestive. So, for example, there's very good evidence that the language capacity is the same for all human groups, which is kind of surprising if you think about it.

...

There are individual differences, but there's no known group differences. So, it follows that there's been no known meaningful evolutionary change with regard to language ever since the time when our common ancestors, who are generally assumed to be quite a small group, maybe even in the hundreds, left Africa roughly about fifty thousand years ago. And if you go back, say, fifty thousand years before that, there's no evidence in the archaeological record of the existence of language at all. So, somewhere in that very narrow window, from an evolutionary point of view a bare flick of an eye, somewhere in that window there seems to have been a sudden explosion – and you see it in the archaeological record – a sudden explosion of creative activity, complex social organisations, symbolic behaviour of various kinds, recording of astronomical events, and so on. It's what Jared Diamond called a great leap forward. It's generally assumed, and plausibly assumed, by palaeo-anthropologists to be associated with the emergence of language. And then, since that time, there appears to be no relevant evolutionary change. If you want to double or triple the numbers, nothing much changes.

Well, these simple observations, which are reasonably well-supported, provide some plausible suggestions about language design. I'll return to them briefly.

Well, suppose we now assume that language exists – that is, that this meeting is not a waste of time – and suppose we want to look at language as a normal biological system – some module of the organism – then the standard questions for any module arise at once. Roughly speaking, what-, how- and why-questions: What are the properties of the acquired system? – each one regarded as an internal system of the individual; How are these properties acquired?; and Why do they have THESE properties and not some other properties?

Well, perhaps the most elementary property of human language is that it consists of a discrete infinity of interpretable expressions – so there's five-word sentences, and six-word sentences, no five-and-a-half words sentence, so it goes on indefinitely like the integers. That's kind of unusual, there's nothing like that known in the biological world. So, that's one isolated property. I mean, maybe, down to the level of DNA. What it means is that each speaker, each child, has somehow internalised a computational process, what's called a generative process, that yields an infinite array of hierarchically structured expressions. And each of these is interpreted – in fact, dually interpreted – at what are called interfaces with other systems. The sensory-motor interface for externalisation, for, you know, making it available to others, and what's called the conceptual intentional interface, that means thought and planning of action and so on. It's a way of saying that language is sound with meaning, Aristotle's common sense dictum.

The unbounded range is a very elementary fact about language that appears to have been barely noticed, if at all, until the seventeenth century scientific revolution. It was noticed by Galileo, Descartes and others, and they were quite mystified by it. And it played quite a significant role, especially in Cartesian thought if you look at it.

Well, an approach to language that doesn't capture at least this property, can't possibly be taken seriously.

For language and any other module internal system, whether it's the visual system or anything else, its growth and development in the individual involves at least three factors. One factor is external data, which can have a big effect on any system. The second is a genetic endowment that converts - for a cognitive module - converts data to experience and guides the general course of development. And third, principles of broader scope, presumably laws of nature which are invariably a major factor in evolution and development. It's much better understood now than it was fifty years ago.

Well, for language, we could analyse the genetic endowment into a component specific to human language - the standard term for that is Universal Grammar, UG - and other components of the genetic endowment that somehow are relevant to language development, so, other cognitive systems, neurophysiological structures and so on. It would seem, prima facie that UG is critically involved in determining such phenomena as the one I gave as an example, this ECP example. Though, if you look closely it's likely that the third factor, principles general to laws of nature, are also crucial for that. But we can explore the role of UG in the very early stages of language development, not with complex examples like that. For example, it's now known that newborn infants instantly and reflexively extract language relevant data from the famous blooming, buzzing confusion around them. That's no trivial task, no one knows how that's done, and it's human specific. So, the infants pet kitten or chimpanzee or whatever can't do it. It just gets the data. But somehow an infant picks out of all of this data, instantly and reflexively, the parts that are somehow language-related.

...

Although real results remain very sparse, contrary to many claims, the role of statistical reasoning and other cognitive processes in language acquisition, potentially at least, should be a significant area of research.

...

There are also presumably conditions imposed on language by the structure of the brain, although so little is known about that that you just can't draw any conclusions.

...

Well, let's turn to UG. The question whether language exists is basically the question whether UG exists. Is there some genetic element that's specific to language? As I mentioned, this is commonly denied, but I'm simply unaware of any coherent alternative.

In the early work on this topic in the 1950's, it seemed as though UG must be extremely rich to allow any degree of descriptive adequacy in dealing with the complexities of language. And one major goal of theoretical linguistics, in my view THE major goal, since that time, has been to reduce the postulated complexity of UG in accounting for the phenomena of language. The reasons for this pursuit are quite straightforward. One of them is just normal rational inquiry, so that is trying to achieve greater explanatory depth. Another reason is the hope that maybe someday there will be an eventual serious study of evolution of language. Evidently, this task, to the extent that it's feasible at all, is rendered more difficult to the extent that the postulated target UG is more complex. Well, actually, the non-existence approach shares the same goal - to reduce UG, actually to zero in that conception. But there are several salient differences to these two approaches to the question. The first is with regard to results. I think, I can't show it here, but I think it's fair to say there are virtually no results in the

non-existence literature, except in terms of this curious notion of science which has been contrived, departing from the entire history of science. The one I mentioned. In contrast, there are quite substantial results in the existence literature. Overwhelmingly, they result from investigating what I call third-factor considerations, considerations of computational complexity. Actually, the inquiry hasn't been phrased in those terms, but that's what it comes down to if you look at it closely.

So, these steps, [uh, for those of you who know what I'm talking about, it won't mean anything to others] include:

dissolving constructions into more general components, the elimination of phrase structure grammars, which have extremely rich stipulations;

radical reduction of the complexity of transformational grammars – these were designed to accommodate non-contiguous relations, which are ubiquitous phenomena and in displacement and morphological discontinuity;

and, finally, in more recent years, unification of these two approaches in what appears to be the simplest possible computational operation – one which has to operate in any generative system, and seems to accommodate both of these, or so it's hoped.

In recent years, the inquiry into these topics is often called the minimalist program, but that term has been extremely misleading. I'm sorry I ever used it. It's just ordinary science, normal ordinary science, extending the main thrust of theoretical linguistics since the earliest days of contemporary work in the nineteen-fifties. And these developments are pretty natural if you think about it. Language is very clearly a computational system, so it makes very good sense to seek the role of general principles of computational efficiency that apply far more broadly and may indeed be rooted in organism-independent natural law. And that's particularly natural given the little we know about the evolution of language.

So, go back to that. It appears overwhelmingly clear that a generative process suddenly emerged at some pretty recent point – maybe that window I mentioned, the great leap forward, right about then. Well it emerged in an individual, mutations don't take place in groups, so some individual was fortunate or unfortunate enough to get this generative capacity. Well, maybe that's some slight rewiring of the brain. Furthermore, there was no selectional pressure at that time. There couldn't be. It's just something that happened to an individual. So, what you'd presumably expect what appeared at that point to be just determined by natural law, there's no other pressure, something kind of like a snowflake. And the same would be true as this capacity of this property is transmitted to offspring. Notice that the capacity itself HAS selective advantage, the person who had this capacity could think, it could plan, it could interpret, you know, could construct, internally of course, complex thoughts. That, you would expect, would have advantage transmitted to offspring in some small hunter gatherer group, maybe a couple of hundred people. It could take over most of the group after some period, and at that point there would be a reason to externalise it, to make it available to others so that it's not just in your own head. And that seems to be the way language works, with externalisation being an ancillary process.

Well, consider again the what-questions. As soon as the question was addressed within the framework – this general framework – about fifty years ago, many puzzles came to light which had previously been unnoticed. Actually, some of them are still on the agenda. Actually, that moment is somewhat reminiscent of the very early modern scientific revolution. So, for millennia, the scientists had been satisfied with traditional answers to such simple questions as to why a rock falls to the ground and steam rises to the sky – they're seeking their natural place, Aristotle's answer, so that ends that problem. Seventeenth century scientists were the first ones actually to have allowed themselves to be puzzled by that and to investigate it and they very quickly discovered that conventional beliefs were completely false. Well, at that point, science entered a new phase, its modern phase.

And the capacity to be puzzled is very well-worth cultivating. That's essentially the Galilean Revolution. I don't think it's happened in other fields, but it begins with a capacity to be puzzled about very simple phenomena, not just to say well it's natural, so nothing to look at. There were myriad puzzles that emerged as soon as the first efforts were made to go beyond hand-waving about analogy and training and pattern formation and so on.

I'll mention two, at least one, maybe two if there's time, because these are among the very few cases for which answers have been sought, relying on non-linguistic cognitive processes and typically adopting the non-existence assumption. One very hoary example has to do with what's called auxiliary inversion. So, take the sentence 'Can eagles that fly swim?'. A simple sentence. We understand right away that the question is whether eagles can swim, not whether they can fly. That is, the auxiliary element 'can' is associated with 'swim' not with 'fly'. That's obvious from the interpretation. It's also obvious from just plain morphology. You can say things like 'Are eagles that fly swimming?' where the 'are' is going with 'swim' obviously. 'Have eagles that fly been swimming?' You can't however say 'Are eagles that flying swim?'. It has a meaning. It means, is it the case that eagles that are flying swim, but you can't say it that way. There's something about the design of language that hampers communication and compels a paraphrase, much like those ECP examples that I mentioned. It's a perfectly good thought, you just can't express it, except in some roundabout complex way. Some result of language design. So the thoughts are fine, it's just that you can't articulate them. Well this simple question about auxiliary inversion, which is quite trivial. It had never arisen in millennia of careful study of language, because the point was just taken to be obvious, like objects seeking their natural place. But why should it be the case? So, take the sentence again 'Can eagles that fly swim?'. Just on the basis of ease of computation you'd expect that the fronted auxiliary 'can' should be associated with the closest verb. That's the easiest computation, hence with 'fly' not with 'swim'. And, actually, communication would be facilitated if both options were available. That would avoid the need for a complex paraphrase in one case.

So, if you want to put the matter a little bit differently, there are two concepts of minimal distance that are available, one is minimal linear distance, let's get the two things that are closest, that would relate the sentence 'Eagles that can fly swim' to 'Can eagles that fly swim?'. That's minimal linear distance, wrong result. Minimum structural distance relates 'Eagles that fly can swim' to 'Can eagles that fly swim?', the right result. The reason is that 'eagles that fly' is a phrase and you look at structural distance and can is the closest one if it's near 'swim'.

Well, the question then is, why does the language learner reflexively minimise the property of structural distance rather than adopting the computationally far simpler property of linear distance? Or, for that matter, adopting both of them, which would facilitate communication, as I said. Well, notice that there is a good answer to the what-question: minimal structural distance remains unchallenged, not only in these constructions but in plenty of comparable constructions and in every language where such things exist. So, it's an answer to the what-question: somehow minimal structural distance is built-in, and the question doesn't arise for the child, it's the only choice.

That still leaves the how-question and the why-question. There's a very simple and quite likely correct answer, and that has to do with the general architecture of the language faculty. So, go back to Aristotle's dictum, language is sound with a meaning. It's probably the case that that's not quite accurate, language is not sound with a meaning, but it's meaning with sound, which is quite different. That means that the core of the language faculty is a generative procedure that yields structured expressions that are interpreted at the thought system, the conceptual-intentional system, while externalization to the sensory-motor system is a secondary process. Well, linear order is clearly required for externalization but not, it appears, for the core systems of thought, from syntax and semantics. It seems then, that it's just a reflex of the sensory-motor system, and, as you would expect, it varies depending on the sensory modality that's used for externalization. So, it's somewhat different for sign and speech.

Well, though it's not uncontroversial, and there are empirical issues, there is growing evidence that this is indeed the case, and, as I mentioned, it fits very well with what we know about evolution and what we know about language more generally – I won't go into that, but it's true. It would also follow that communication is an even more peripheral aspect of language than I mentioned before. Because communication is contingent on externalisation, and that the substantial literature of speculation on language evolution is simply on the wrong track.

Well, if that's correct, you immediately get answers to the how- and why-questions for auxiliary inversion. Linear order is simply not available to the language learner. The language learner is guided by a very simple principle of Universal Grammar, look for minimal structural distance, and that's the only one available – because you don't have linear distance, except at some peripheral part of language – and that determines the general architecture of language and it's computationally quite natural.

Well, auxiliary-inversion has been the topic of a considerable inquiry industry in computational cognitive science. In fact, it's the main topic that's been studied to try to show that the child acquires this knowledge on the basis of statistical analysis of a corpus of data, in accordance with the non-existence thesis. And there are new papers that come out regularly on this, even in the past few weeks. They have very curious properties. One is, every one of them fails, and fails dramatically. Second is, they're all cited as examples of successes in the literature. And they are examples of successes in terms of this new notion of success which departs from the entire history of science, more or less matching unanalysed data. Another curiosity is that they all ignore the simple explanation, which is pretty well conceptually and empirically based – minimal structural distance, for the reasons I mentioned. Another reason is, it wouldn't matter if these approaches succeeded, although the people carrying them out don't seem to perceive that. If these methods succeeded, it would mean nothing, because if you had a non-language which used linear distance rather than structural distance the methods would work for that too. But that's not the case, every language uses structural distance. So even if they worked, then they're total failures cause it wouldn't mean anything. But that doesn't seem to matter.

Well, kind of a background question which you might think about is: How does the child even know what the intended interpretation is in such a sentence as 'Can eagles that fly swim?' unless it's already relying on the minimum structural distance principle without any data at all?

Well, I was going to talk about some other cases but I think I'll skip em. But, there are one or two other cases which have been extensively discussed in the computational science literature and you find the same thing: total failures, regarded as successes, wouldn't matter if they worked, and it's considered fine. It's a strange thing cognitive science, cognitive scientists ought to be interested in thinking about it.

Well, I don't want to end without at least mentioning another extremely serious problem which has been barely addressed. A computational procedure requires certain atoms of computation, something to start with. For language that's called a lexicon, a lexicon of minimal elements, maybe roughly word like, but not exactly. Even the simplest of these elements raise fundamental puzzles, fundamental problems about how they relate to the mind-external world.

Well, there are two aspects to the question – one is meaning, the other is sound. The sound aspect is secondary, ancillary, if the reasoning I reviewed earlier proves accurate. For sound it's been investigated. The answers lie in articulatory and acoustic phonetics, and the problems are hard. They've been studied intensively for sixty years using sophisticated technology. It's given some answers, but it leaves many outstanding problems.

Well, what about meaning? Now that hasn't been investigated because of a failure to reach the level of a Galilean revolution. There's a standard answer, like 'objects go to their natural place'. The standard answer is provided by

what is called referentialist doctrine. So, the word 'cow' picks out cows, maybe by some causal relation. Actually something like that does seem to be true for animal communication. So, symbols in animal communication do appear to relate to physically identifiable external or internal states. So, say, for a vervet monkey, motion of leaves elicits a warning cry, which we interpret as saying 'an eagle's coming', I don't know how the vervet interprets it. Or, say, 'I'm hungry' goes back to an internal hormonal state which is identifiable. Nothing remotely like that is true for even the simplest elements of human language - 'cow', 'river', 'person', 'tree', pick any one you want.

There are inklings of that understanding in classical philosophy back to Aristotle, and it was considerably enriched during the first cognitive revolution in the work of British neo-Platonists and classical empiricists. They recognised that there is no direct link between the elementary elements of language and thought and some mind-dependent external entity. There's just no link. There's no relation of reference in human language and thought, though there is an act of referring. That's a very different matter. Rather, these minimal elements, like these elementary words, they provide rich perspectives for interpreting the mind-independent world. In the seventeenth century they were essentially studying what we call, they didn't call them this, Gestalt properties, cause-and-effect, sympathy of parts, concerns directed to a common end, psychic continuity, and other mentally imposed properties. Properties that the mind imposes on the world, physicists can't find them.

In this respect, meaning is rather similar to sound. So, for example, every act of articulating some item, say, articulating the internal syllable 'ta', let's say, every act of articulating it yields a physical event, but no one seeks some category of physical events associated with 'ta', there isn't any. Similarly, some, but by no means all, utterances, uses, of the word 'cow' or 'river' or whatever you like, some of them refer to physically identifiable entities, but there's no category of such entities that's identifiable in principle by a physicist investigating the mind-external world. David Hume had a simple phrase summarizing a century of inquiry into this, he said that "the identity that we ascribe to vegetables, animal bodies, artefacts, persons and their minds – the array of properties – is a fictitious one established by" – what were called then – "our cognoscitive powers" – our cognitive powers we would call now.

Well, most of this has been forgotten, unfortunately, but there's quite strong evidence that it's basically correct. And, once again, failure to be puzzled by these things is a very serious error. It seems that these elements, so fundamental to human language and thought, reveal another vast chasm between humans and other animals, and they pose a huge problem for evolutionary biology and a comparably huge problem for acquisition of language – what's sometime called a poverty of stimulus problem, similar to those I mentioned. Now, there's interesting experimental evidence that a child acquires these words – say a word like 'cow' – on a single exposure. That means on a single exposure to the word 'cow' in some situation, the child conjectures what it means, and that's almost invariably right – confirming evidence doesn't seem to have any effect, I mean disconfirming evidence would have some effect, then maybe we'd abandon it and pick up something else. But it seems to be one-trial learning, in other words, no learning at all, something just clicks and you get the association with the meaning. Of course, the sound symbol association is acquired, but the meaning is not acquired, that's somehow coming from the inside, without any direct evidence, much as in the cases I discussed. And if you investigate the meanings of these words, it turns out what's coming from the inside is quite rich.

The what- how- and why-questions raised by these systems have only barely been explored, because just of the unwillingness to be puzzled, despite their fundamental significance for thought and action. The origins of these basic elements and their organisation, now that remains entirely unknown, and it Richard Lewontin is correct, then they may remain a permanent mystery for human inquiry, but they certainly will be if they're not even investigated.

So, in this domain, a basic one, we really are in the pre-Galilean stage.

Okay, I'll stop here.

+++++

Q & A Session

Q1: I wonder Mr. Chomsky if you might briefly discuss the difference between interpreting language literally and figuratively or ironically. Sometimes we tend to take a singular interpretation versus multiple meanings from any given snippet of text.

NC: That's certainly true. I mean, these are all very important complexities. But, what's striking, is we don't even begin to understand the simplest cases. Let's go back to Galileo and motion. I mean, Galileo didn't try to find out how leaves fall. I mean, it's true leaves fall too, but, you know, that's so overwhelmingly complex that even modern physicists don't even try to study it. First, you want to see if you can find whatever works for the simplest cases, and maybe on the basis of that you can extend to other cases, and I think that's true here too. So, yeah, these are important questions. Doesn't mean they shouldn't be investigated. You know, you can learn things about them, but one should recognise, we're not going to understand much about them if we can't even understand how the child comes to grasp the concept 'cow' or 'person' or 'tree' or 'river' or anything else. And we're very far from that.

Q2: One of the things you argued in support of is Language's productivity, and Fodor uses the same argument for a language of thought. I wondered if you think language plays a greater role, like Fodor does, and is there such a thing as a Universal Grammar of thought?

NC: Well, I DO think so, and, in fact, what I was describing was a language of thought. What actually seems to have happened, as far as we can piece anything together, and as far as the empirical evidence shows, is that at some point – maybe seventy-five thousand years ago – some small neural rewiring took place, of course, in some individual, because it's the only possibility, and that individual had a computational process which was somehow linked to pre-existing conceptual structures. Now, what those pre-existing conceptual structures are, we haven't a clue about that. That's the problem I ended with. Nobody has a clue what that could be or where it could have come from. But it's there, and it's totally different from anything in the animal world. And if this generative procedure could link to it, you do have thought. So, that's a language of thought. And then somewhere down the line it got externalised and you get interactions among individuals. So, yeah, I think it is a language of thought. But I don't see any reason to think it's separate from language, I think it just is language. In fact, if you look at Fodor's work, and you ask the question 'What do we know about the language of thought?', well it turns out to be English.

[Audience murmurs]

Q3: You are saying that everybody thinks in English.

NC: I'm not saying that we think in English, but the reason it turns out to be English is that's the language everybody's using. Whatever the language of thought – this internal, our own internal language, yours and mine – whatever it is, is inaccessible to introspection. Okay, if you introspect, you can't go one minute without talking to yourself. It takes a tremendous act of will not to talk to yourself. In fact you do it all night, it keeps you up all night. It's just impossible not to do it, but what you're introspecting is the externalised language. So, you can tell when you're talking to yourself, you can tell whether two sentences rhyme, okay. Or you can tell how long they are, or something like that. And actually, if you really pay attention, you're not really talking to yourself in sentences, just kind of odd little fragments. Something is going on deeper which we CAN'T introspect into,

anymore than you introspect into the mechanism of vision, and that's the language of thought. And it's probably universal. It's hard to imagine how it could be anything else. There's no evidence for acquiring it.

...

Q4: Do you think that despite the complications and lack of knowledge we have that we'll develop some sort of machine with language capability which is a) functionally equivalent to our own, or b) similar to our own, in the next 10, 20, 50, 100 years, ever?

NC: It's hard to know what that means actually. I mean, we're a machine, okay? We're not made out of ectoplasm. So, we're some kind of biological system which means, you know, that at least in principle somebody could construct that system. You could ask that question about the visual system. Can you construct a mammalian visual system? Well, you know, everyone presumes that probably you can, but that question is the same as can we understand what the visual system is. And I think it's the same with language. To the extent that you understand what the system is, you can at least, in principle, duplicate it in some other physical mechanism, but it doesn't tell you anything. And in fact we can't get very far in it, because so little is understood. Look, what I was talking about is artificial intelligence. It's not called that, but this is artificial intelligence, trying to figure out what the mechanisms are that enter into one capacity of intelligent behaviour. That's presumably what artificial intelligence is. What people usually mean by artificial intelligence is, you know, a little robot that does amusing things, you know. Well, okay, that's ...

There was a famous debate – this is kind of unfair, but – there was a famous debate at MIT in, well, long before your day, back around nineteen-sixty I guess, between Marv Minsky, who's the kind of guru of Artificial Intelligence, and Jerry Lettvin, who was a very smart biologist. And in the course of the debate, Jerry famously said he thought it was a great topic – artificial intelligence – but he thought it was misnamed, it should be called 'Natural Stupidity'. A little unfair, but ...

Q5. Dr Chomsky, I was wondering, when you were speaking about the evolutionary biology part of the question, i.e., what moment in time when a member of the species developed the capacity to have language. First of all, without a written language, what kind of evidence can you find in the archaeological record, and secondly, is it not possible that at the moment that the primate that is us developed language we became a new species entirely, such that language is such an integral part of homo-sapiens that you can't really distinguish between people without language and people that have language.

NC: Well, as to becoming a new species, that's a matter of definition. If some rewiring took place, then it's in a sense, a new species. You know, how much change has to take place to call it a new species, there no answer to that. So, yes it was a different creature in a minimal sense. What you were just saying is kind of similar to Richard Lewontin's argument. He says by the kind of inquiry that anyone knows how to carry out, we're never going to find anything out about this. Not only is there no written language, there's no tape recordings, there's no direct evidence of what must have happened. But that's true of palaeo-anthropology altogether. All you have are artefacts, you don't even have soft tissue, cause that degenerates. What you have is, you know, skulls, hard tissue and you have artefacts, like tools, you know, drawings, recordings of, which you can pretty well determine are recordings of natural events. Well it turns out that if you look at the palaeo-anthropological literature – there are good summaries of this, like by Ian Tattersall, for example, one of the main scientists – it turns out that roughly in the window that I mentioned, you know you can't give a date to it, but, say, seventy-five thousand years ago, plus or minus a couple of tens of thousands of years, it doesn't matter, you do get this so-called great leap forward. All of a sudden there's an explosion of complex artefacts which reveal a lot of creativity. If you compare humans with Neanderthals, that's the closest living, almost living relative, up till thirty thousand years ago, they were all over the world, they were very successful, biologically successful organism, all over the place,

had very sophisticated abilities. Like Neanderthal could make tools so refined that modern humans can't make them except with instruments, but they were all the same. You go to Asia and Europe and Africa, wherever they are, it's always the same tool. There's no indication of any creativity. It's a little exaggerated, but overwhelmingly there's essentially no creativity, and that's the record of early human origins too: complicated tools, but they're made in a particular way a million years ago, and then they're made in a different way two hundred thousand years ago. There aren't signs of major change. The big change seems to be roughly in this window and it's indirect evidence, it's coming from artefacts. There is some evidence about the brain size. It appears to be the case – you can tell from the skull – that the brain did undergo some big increase around a hundred thousand years ago. And it's been speculated that whatever that brain increase was, nobody knows why it happened, but whatever it was, it allowed for changes in the internal structure of the brain – which we know almost nothing about, incidentally – that permitted this development to take place. I mean, it's indirect evidence. Yeah, it's kind of like cosmology, you're not there, you have to reconstruct from indirect evidence, but it doesn't mean you can do nothing.

Q6 [Question that's unclearly formulated]

...

NC. There's a biological question. Is there a genetic component to what you and I are doing.

...

There is a straight scientific question, you and I are doing something, okay, does that trace back to some genetic factor that's common to you and me and every other member of the species and is non-existent in any other species? That's a scientific question. If the answer to that is 'yes', language exists. If the answer is 'no', then it's just like today's weather, it's just some odd collection of things going on. It's not an easy scientific question, but it's a scientific question.

...

Q7. Hello Professor, thank you for coming, and you still look fantastic, I must say. So, with language being so closely associated with power, do you agree or disagree with language being gender constructed. For example, the rule of his/her in a sentence rather than her/his, and is this based on a patriarchal culture circumstance or is there some fundamental principle that can explain it.

NC. It's an interesting fact, but I don't think it's really a fact about language, it's a fact about society. I mean, if we had a matrilineal society, it'd probably be the other way around. So, it's a fact about the way language is used in particular social and cultural circumstances. That's one of those many variable problems – at that point you're studying a lot of different factors, not just whatever the language is, which is probably uniform for all humans, but what happens when this kind of biological instrument, if you like, is used under various conditions? And, of course, it'll be used in different ways. And those are interesting questions, but not in a narrow sense linguistic questions. Those are socio-cultural questions involving language, which is fine, nothing wrong with those questions, they're just different ones.

Q8. Dr. Chomsky, given the ubiquitousness of information communication technology and how that seems to be so much a part of human development now, if you look at children and what they're exposed to, can you offer some of your speculations on how, moving ahead generations, that may impact language and thought?

NC. It's a good question. Actually there's been some study of it. So, you know, children now are just immersed in complex video games, Twitter, text messaging. I mean, you know, like my daughter, when she tries to communicate with her children, she can't use email anymore, that's too old-fashioned. For a while, she used to

use – I don't know what, I think now, maybe there's something beyond text messaging. You have to use something that's extremely fast and very short and it doesn't take more than thirty seconds to look at, cause there's too much other stuff going on. In fact, there was just a study that came out, some hospital, I've forgot where, I think it was in Boston, was studying a new disease that's developing among teenage kids – almost all teenage girls – sleep deprivation. There's a very serious epidemic of sleep deprivation and it turns out that it traces to the fact that these kids are holding their cell phones, or whatever they use, because they don't want to miss the possibility that at three a.m. some friend of theirs will say 'I walked across the street', you know, or something like that. So, you have to be awake to catch up with that. These are pretty interesting facts.

...

There are studies of what effects this will have, and nobody really knows.

Q9. Hi, I wanted to ask you about the proper relationship between language documentation and careful linguistic analysis. Because, if we're doing this kind of normal science mode of needing crucial experiments, it seems some of the crucial data might be disappearing. As endangered languages disappear, we might be losing access to some of the data that could inform our linguistic analyses. So, what do you think the proper role of theoretical linguistics should be with respect to language documentation? Like the kind of old fashioned linguistics that preceded generative grammar.

NC. I mean, it's reciprocal. First of all, language documentation requires linguistic theory. You have to know what you're documenting. So, like when I was a student, the work that we did as students was language documentation. You know, you take some Native American Indian language and try to get data from it. But the data you were looking for was dependent on what linguistic theory tells you to look for. There's nothing else you can do. All documentation is theory-guided. Well, as theory changes, people do different documentation. So, field linguists today just look for totally different things than they looked for, say, fifty years ago. I mean, then you were looking for phonetic contrasts and so on, now you're looking for whether there are transformational rules meeting certain island conditions. I mean, documentation is just data collection. And yeah, fields of course require data, but the kind of data you look for depends on what you think exists, what you think is important. And a lot of the data are just contrived. So, these examples that I mentioned, say, these ECP examples, you'll never get them in a corpus of data. I mean, almost inconceivable. But those are examples you construct on the basis of theoretical ideas, which themselves are based in data, and then you see what happens in those cases. That's exactly like the normal sciences. An experiment is very remote from anything that actually happens. You know, you're interested in very exotic circumstances when you do an experiment. That's even true for Galileo, obviously for modern science. So, I don't really think it's an issue. There's a constant interaction between data collection and theoretical development and they influence each other.

42
Shares

CHOMSKY.INFO
