

**VANDIVER:** Good afternoon and welcome. I'm Kim Vandiver, the chairman of the MIT faculty. We gather here today to honor two people-- James R Killian Jr and professor Noam Chomsky.

The Killian Award is the highest award that the MIT faculty bestows upon one of its own. It's named after the late Jim Killian, 10th president of MIT. He came to MIT as a student in the mid 1920s. He joined the staff of *Technology Review* in 1926 and he retired as chairman of the MIT Corporation 45 years later in 1971.

He brought great honor and distinction to MIT through his many years of national service. I'll mention two examples, which are just the tip of the iceberg. In the late 1950s he was the first full-time science adviser to a President of United States and provided important leadership in the period that set the nation's course to putting a man on the moon. He also served as chairman of the Carnegie Commission on Educational Television, which recommended the formation of the Corporation for Public Broadcasting.

He was also a very genuine and likable person, as I discovered quite by accident one day about 20 years ago when I was a new graduate student at MIT. I sat down at a table in the Lobdell Dining Room with an elderly gentleman that I didn't know and we struck up a conversation, one that was to begin a long series of enlightening conversations I had with him until the time of his death in 1988. I think his easy way with students and with other members of the MIT community are one of the things that define Jim Killian to us.

This award is named in his honor. We bestow it upon a faculty member each year who has also brought honor and distinction to MIT. Professor Chomsky has been doing so for 37 years as a member of our faculty. Professor Chomsky, would you come to the stage, please?

In preparing for today, I did one small piece of research on Professor Chomsky, which I'd like to share with you. With the help of the MIT library staff, I was able to determine how often he is cited in the Arts & Humanities Citation Index. As of the last time the statistics were compiled, we discovered that he is in the top 10. It happens that none of the others are still living.

[LAUGHTER]

He is the most cited living author. I thought you might like to know the names of the company he keeps. They are in order Marx, Lenin--

[LAUGHTER]

[APPLAUSE]

--Shakespeare--

[LAUGHTER]

--Aristotle, the Bible, Plato, Freud, Chomsky, followed by Hegel and Cicero.

[LAUGHTER]

[APPLAUSE]

I would like to read from the award. It says the president and faculty of the Massachusetts Institute of Technology have the honor to present the James R Killian Jr Faculty Achievement Award to Noam Chomsky and to announce his appointment as Killian Award Lecturer for the academic year 91-92 in recognition of his extraordinary contributions to linguistics, philosophy, and cognitive psychology. One of the foremost thinkers of this century, he revolutionized the study of language and of the human mind. His books, monograph, and articles constitute major milestones in the history of linguistics. An inspiring teacher, he has trained scholars who are among the most important contributors to the discipline. His presence here has greatly enriched MIT. Professor Chomsky, the microphone is yours.

**CHOMSKY:** Thanks.

[APPLAUSE]

It was lucky he didn't read what those citations say.

[LAUGHTER]

The title of this talk, the announced title, has three parts. It refers to language, of the cognitive revolutions, and then there's the quote around cognitive revolutions. The relation between language and cognitive science is natural-- that's a reasonable combination. Modern linguistics developed as part of what's sometimes called the cognitive revolution of the 1950s, actually to a large extent here in Building 20, thanks initially to Jerry Wiesner's initiatives. It was also a major factor in the development of the cognitive sciences ever since their modern origins. So those two words are sort of obvious.

The quotes are intended to express a lingering skepticism. And when I say cognitive revolution from here on in, you're supposed to hear so-called cognitive revolution. I'll explain why as I proceed. The skepticism that I feel does not have anything to do with the actual work that's done in the various areas of cognitive science-- so vision, language, cognitive development, a few others. In these areas there have been quite significant advances over the past 40 years. The questions in my mind arise at a second order level that is thinking about the nature of the disciplines that are concerned with what were traditionally called mental acts or mental faculties. And I want to come back to a few of those questions-- some of them are substantive, some of them are historical. First a few words about the contemporary cognitive sciences and the cognitive revolution as they look to me. And let me stress that others don't see it this way at all. This is not intended to be a survey, but rather a rather personal point of view.

George Miller, who was one of the leading figures and founders of the modern field, in a recent article traced the modern cognitive sciences back to a meeting in 1956 held in Cambridge. I can't remember whether it was here or at Harvard, but one or the other. It was a meeting of the Institute of Radio Engineers, which had a series of papers on experimental human psychology which made use of information theory, and signal detection, and other-- then-- pretty new ideas. I'd had a paper of mine on language which outlined some of the basic ideas of generative grammar. It had another paper on problem solving and reasoning. It was a paper by Herb Simon and Al Newell, a simulation of theorem proving in elementary logic. And he argues that that confluence of research papers essentially initiated the modern field.

Well there were some shared ideas among the participants. One shared idea was a shift in perception or perspective toward the disciplines, a shift which is not terribly surprising now but was pretty striking at that moment. The shift was away from the study of behavior and the products of behavior, such as words, or sounds, or their arrangements in text, a shift from that focus toward the study of the states and the properties of the mind and the brain that enter into human thought and action.

From the first point of view, which overwhelmingly prevailed at the time, behavior and its products were regarded as the object of study. From the second point of view, what later came to be called the cognitive sciences, behavior and its products were of no particular interest in themselves, they simply provided data, data which would be interesting insofar as it served as evidence for what really concerns us now, the inner workings of the mind, where the phrase mind here simply is a way of talking about relevant properties, and states, and processes of the brain. The data of behavior, and the texts, and so on from this point of view are to be used alongside of other data, whatever might turn out to be useful-- say, electrical activity of the brain-- but it has no inherently privileged status and it's not the focus of interest.

That's a shift. It's a shift away from what was called behavioral science, or structuralism, or statistical analysis of texts and a shift toward something that might properly be called cognitive science. In my personal view, it was a shift from a kind of natural history to at least potential natural science. Now that shift was highly controversial at the time, and it still is. And in fact, if you were to use such measures as amount of government funding, it would still be very much in the minority. But I think it was the right move. Well that was one shared interest from independent points of view. There was no communication among the people prior to the meeting.

A second shared interest was an interest in what we might call computational representational systems. That means a way of looking at mental activities and mental faculties of what the brain is doing, to looking at them as a kind of a software problem, that is the study of the mechanisms but viewed from a particular abstract perspective. Analogies can always be misleading. But one might compare this approach to, say, chemistry about a century ago, where there were notions like valence, or periodic table, or organic molecule and so on, but they were not grounded in the basic physics of the time. The results, however, and the ideas did provide guidelines for what turned out to be quite radical changes in physics that were necessary to accommodate these ideas. Chemistry was never actually reduced to physics. It would be more accurate to say that physics was expanded in a certain sense to chemistry.

Similarly, computational representational series of the brain attempt to establish states and properties that have to be accommodated somehow in the brain sciences and might well provide, as far as they're correct, they might well provide guidelines for inquiry into mechanisms. It's possible, again, that they may provide guidelines for what will have to be radical changes in the way that the brain and its processes and states are conceived. How two approaches to the natural world will be unified, if ever, can never be guessed in advance, really. Could be reduction, as in, say, the biology/chemistry case, more or less. It could be what you might call expansion, as in the physics/chemistry case. Or it could be something else.

Crucially, there's nothing mysterious about any of this. It seems, to me at least, to be just normal science, normal science at a typical stage before unification of several approaches to some topic in the natural world before they are unified. Well, again, that was and remains highly controversial.

Assuming it anyway, how can we proceed? You can look at the brain or the person from a number of different points of view. Let's begin by taking what you might call a receptive point of view. So we think of the brain, say, as being in a certain state. A signal reaches the brain. It produces some kind of symbolic representation of the signal internally, presumably. And it moves to some other state. If we focus on the interstate transitions, we're looking at problems of learning or, maybe more accurately, growth. If we focus on the input signal and the symbolic representation of it, we're looking at perception. If we focus on the initial state of the system prior to any relevant signals, we're looking at the biological endowment, the innate structure. We could trace the processes back further. At that point we'd be moving into embryology. And one should stress that these distinctions are pretty artificial, actually.

Typically, the systems that we're looking at are like physical organs, or we might say more accurately like other physical organs, because we are studying a physical organ from a particular point of view. They're like physical organs, other physical organs in that-- let's say like the digestive system, or the circulatory system, or the visual system, or whatever-- which, again, we often look at abstractly, as one would expect in the study of a system of any degree of complexity. We can call these particular systems mental organs, not implying, again, anything more than that they are involved in what were traditionally called mental states, mental faculties, or mental processes.

Well, like other physical organs, these mental organs quite typically reach a kind of a stable state at a certain point. That is, the input signal that's causing an interstate transition moves the system around within some kind of equivalence class of states that are alike for purposes at hand. That's obviously an idealization, but a reasonable one. When you look at the stable state, after a point at which there isn't significant change, then you're very clearly studying perception when you look at the signal-result relationship.

So for example, if you take the visual system, if you look at the initial state, the state prior to relevant signals, you're looking at the basic structure of the visual system. If you look at the interstate transition, you're looking at the growth of the visual system, the changes in the distribution of vertical and horizontal receptors in the striate cortex or the onset of binocular vision after a couple months of age in humans, and so on. If you look at the stable state, you're studying adult perception. Adult perception, of course, has some kind of access to systems of belief about the world. Not much is known about this, but it seems reasonable to assume that that involves some kind of computational procedure that specifies an infinite range of beliefs which are called upon when needed, although here we move into a domain of almost total ignorance.

Turning to the language system, you can look at it the same way. The initial state of the system, the initial state of the language faculty-- the study of it is sometimes called universal grammar these days-- the initial state of the language faculty, the study of that is a study of biological endowment. It's an empirical fact which seems pretty well established that this element of the brain is pretty uniform across the species. The variations appear to be quite slight apart from extreme pathology. And it also appears to be unique in essentials-- there's no other known organism that is anything like it. Also, it's rather different from other biological systems in many respects. Systems of discrete infinity, which this plainly is, are pretty rare in the biological world, so it appears.

If you're looking at the interstate transition, as signals, some relevant information comes in, you're studying language acquisition. It apparently reaches a stable state at about puberty. This stable steady state incorporates a computational procedure which characterizes an infinite set of expressions, each of them having a certain array of properties-- properties of sound, properties of meaning, properties of structural organization. That computational procedure constitutes a person's language.

To have such a procedure in the appropriate part of the brain is to have a language or, as we sometimes say, to know a language. When I speak of a language, I simply mean the computational or generative procedure that specifies the properties of a infinite range of linguistic expressions, such as, say, the one I just produced. Since your language is similar enough to mine, when the signal that I just produced reaches your brain, your generative procedure constructs an expression, an expression with these properties, which is similar enough to the one that I form so we understand one another more or less.

Unlike the system of belief that's accessed in visual perception, in this case we happen to know quite a lot about the generative procedures that are involved in the use of language, and even about their innate invariant basis. And these, in fact, are the areas of major progress in the past generation.

Well, what about perception of language? There are several common assumptions about this. One major assumption is that there is something called a parser. A parser is something that maps a signal into a symbolic representation, paying no attention to any other aspects of the relevant environment. Of course it is assumed, if people are rational, that this parser accesses the generative procedure-- that is, it's assumed that you use your knowledge of English when you interpret signals.

Note however several assumptions. First, one assumption is that such a thing exists, that a parser exists, there is some faculty in the mind that interprets signals independently of any other features of the environment. That's not at all obvious. The existence of an internal generative procedure is far more secure and far more obvious than the question of the existence of the parser. And it's also a far better grounded assumption, far better grounded empirically, and with much richer theoretical structure. This is contrary to common belief, which usually considers the other way around. Nevertheless, let's assume that a parsers exists, though it isn't obvious.

It's further assumed that the parser doesn't grow the way a language grows from some initial state to a stable state. Now that assumption, again, is not very well founded. It's based mainly on ignorance. It's the simplest assumption, and we have no reason to believe anything about it, so we might as well believe that. There is some kind of evidence for it, indirect perhaps but interesting. There is a forthcoming book coming out from MIT Press by Robert Berwick and Sandiway Fong-- a former student of his-- on what you might call a universal parser, a fixed invariant parser, which has the property that it if you flip its switches one way it interprets English and if you flip its switches another way it interprets Japanese. And if that's the right idea, then there is a universal parser which is unaffected by environmental factors apart from the properties that serve to identify one or another language, that is the array of flipping of switches, and we'll come back to that.

There's a further-- so that assumption, the assumption that the parser doesn't grow, while strange, could be true. And there's only this indirect evidence for it. A further common assumption is that parsing is easy, something that you do very quickly and easily. In fact, that's often held to be a criterion for the theory of language. The criterion is supposed to be that the theory of language must provide generative procedures that satisfy this requirement, that parsing is easy and quick.

Well that one's flatly false, uncontroversially. Parsing often fails completely and quite systematically, and it's often extremely difficult. Many categories of what are called parsing failures are known, even with short, simple expressions. Parsing failure just means a case where the parser works fine but it gives some kind of an interpretation which is not the one that's assigned to the signal by the language, so it's wrong in that sense. It's perfectly true-- and very often you can easily find categories of expressions, even pretty simple ones, where the parsing system just breaks down. People don't know what they're hearing, sounds like gibberish, even though it's perfectly well formed and has a fixed, strict meaning, so on and so forth.

Now it's perfectly true that one uses expressions that are readily parsed, but that emerges on tautology. What it means is we use what is usable. Other resources of the language we just don't use. There's a related assumption that's often alleged, that languages are well adapted to the function of communication. It's not clear that that statement is even meaningful, or any similar statements about biological organs. But to the extent that one can give some meaning to that statement, this one again looks just false. The design of language appears to make it largely unusable, which is fine. We just use those fragments that are usable for us. It doesn't mean that the design made it usable somehow.

There are similar assumptions in learnability theory, formal learnability theory. It's often claimed as a starting point that natural languages must be learnable. In fact natural language is often defined, natural languages are defined often as the set of languages that are learnable under standard conditions. Now noticed that that could be true. We don't know that that's false as we know that the comparable statement about parsing is false. It could be true that languages are learnable.

But it's not a conceptual necessity. Language design might be such, our language faculty could be such, that it permits all kinds of inaccessible languages. These would be realizable in the brain, just as, say, English is, but not accessible to humans. Again, that wouldn't stop us from doing what we do. We would just pick up the accessible ones. In fact there is some recent work, quite recent work, which suggests that natural languages are learnable. But if true, that's an empirical discovery, and in fact a rather surprising.

Well, I've said nothing so far about production of language. There's a reason for that. The reason is that apart from very peripheral aspects it's almost a complete mystery. We can learn a good deal about the mechanisms that are put to use when we speak. And we can study these certain aspects of perception, the mapping of signals to percept using the internalized mechanisms, the generative procedure of the language. But no one has anything sensible to say about what I'm doing now or what two people are doing in a normal conversation.

It was observed centuries ago that the normal use of language has quite curious properties. It is unbounded, it's not random, it is not determined by external stimuli-- and there's no reason to believe that it's determined by internal states, it's uncaused but it's somehow appropriate to situations, it's coherent, it evokes in the hearer thoughts that he or she might have expressed the same way. That collection of properties, we can call them the creative aspect of language use, that set of properties was regarded by the Cartesians as the best evidence that another organism has a mind like ours. not just a body like a dog, or a monkey, or a robot.

Actually if you look back at their argument it was quite reasonable. The science happened to be wrong, but it was very reasonable. And within the framework of scientific reasoning, we might put matters in different terms, but we haven't made any progress in understanding these critically important phenomena. The phenomena that concerned them provide much of the basis for the traditional mind-body distinction. In my opinion, in fact, contemporary thought about these topics, usually phrased in terms of notions like the Turing test, involves a considerable regression from the much more scientific approach of Cartesian psychology. I'm not going to have time to elaborate properly, I merely mention this to provoke some possible interest or maybe outrage. We'll see.

Well, I've mentioned several kinds of shared interests in the early modern cognitive sciences, say back to the 1956 meeting that George Miller pointed to. One common interest, again, is in computational representational systems, looking at the brain in terms of software problems. For linguistics at least, this turned into the study of the generative procedures that are put to use in perception, and presumably in expression of thought-- speaking-- although the latter, again, remains a complete mystery.

There's another shared interest-- it's what we might call the unification problem. That's a problem that has two different aspects. One aspect is the relation between the hardware and the software, the relationship between computational representational theories of the brain and the entities that they postulate, and what you might call implementation. And as I suggested, I think you can regard that as similar to the problem of the relation between, say, chemistry and physics 100 years ago, or genetics and chemistry 50 years ago, and so on. That's one problem, one aspect of the unification problem, the deepest aspect.

A less deep aspect has to do with the relation among the various domains of cognitive science, say language and vision, whatever may be involved in problem solving. This is the question what do they have in common. Notice that that's a much less profound problem. If it turns out the answer is nothing, just the way it is. On the other hand, we'd be kind of upset if the answer to the first one was no answer.

Well, at this point some of my skepticism begins to surface. So let's take the first one, the deep problem, the serious scientific problem, the relation between roughly the hardware and software, chemistry and physics. Here there are two different approaches that you can identify.

One approach, let's call it naturalistic, goes kind of like this. It says the mental organs are viewed as computational representational systems, and we want to show that they're real, that their statements in these theories are true, that is that there are certain states in the properties of the brain which these theories exactly capture, try to show. The mental organs can also be viewed in other ways. For example, you can look at them in terms of cells or atoms. Or more interestingly for the moment, you can look at them in terms of electrical activity of the brain, so-called event related potentials, ERPs.

Now rather surprisingly, there's some recent work which shows quite dramatic and surprising correlations between certain properties of the computational representational systems of language and the ERPs, the potentials. For the moment these ERP measures have no status apart from their correlation with categories of expressions that come out of computational representational theories. That is, they're just numbers, picked at random. Because there's no relevant theory about them that says you look at these numbers and not at some other numbers. In themselves, in other words, they're curiosities. Still, it's interesting, because there are correlations to rather subtle properties that have emerged in the attempt to develop computational representational theories which explain the form and meaning of language.

So that suggests an interesting direction for research, namely to try to unify these quite different approaches to the brain and to place each of them in an appropriate theoretical context. For the moment, that's primarily a problem for the ERPs. But if you can pursue it, could be quite an interesting direction. One can think of lots of consequences.

Well the same is true of the relation between CR systems, computational representational systems, and physiology, at least in principle. Here it's just in principle, because in practice at present we don't really have much to look at there except from the theories of the computational representational systems. But here, whatever the difficulties may be, we seem to face ordinary problems of science before the typical situation, when there's no understood way to unify them or connect them.

Well, there's a different approach to this unification problem, which actually dominates contemporary thinking about cognitive science-- not necessarily the actual work, but the thinking about what it amounts to. This approach is to divorce the study of mental organs from the biological setting altogether and to ask whether some simulation can essentially fool an outside observer. So we ask whether the simulation passes some rather arbitrary test, called the Turing test, but Turing test is just a name for a huge battery of possible tests. And so we pick out one of those at random.

And we now ask whether some, say, extraterrestrial invented organism or some actual programmed computer passes the test, that is fools somebody. And people write articles in which they ask how can we determine empirically whether an extraterrestrial or a programmed computer, a machine, can, say, play chess or understand Chinese. And the answer is supposed to be if it can fool an observer under some arbitrary conditions called the Turing test then we say we've empirically establish that.

Actually you may have noticed last November there was a big splash about a running of the Turing test over at the Boston Computer Museum together with the Cambridge Center for Behavioral Studies and a quite distinguished panel that a bunch of programs which were supposed to try to fool people. There was I think a \$100,000 prize offered, a big story in the *Boston Globe*, front page story in the *New York Times*. *Science* had a big splash on it. I wasn't there, but I'm told that the program that did best in fooling people into thinking that it was a human was one that produced cliches at random.

[LAUGHTER]

Apparently it did pretty well. I don't know. It may tell you something about humans, but I'm not sure what.

As far as I can see, all of this is entirely pointless. It's like asking how we could determine empirically whether an airplane can fly. The answer being, if it can fool someone into thinking that it's an eagle, say, under some conditions.

[LAUGHTER]

But the question whether an airplane can fly, and whether it's really flying-- or for that matter, the question whether high jumpers were really flying at the last Olympics-- that's not an empirical question. That's a matter of decision, and different languages make decisions differently. So if you're speaking Japanese, I'm told airplanes, eagles, and high jumpers all are really flying. If you're talking English, airplanes and eagles are really flying, but humans aren't, except metaphorically. If you're talking Hebrew, eagles are flying and neither airplanes nor humans are. But there's no substantive issue involved, there's nothing to determine, no question to answer. Simulation of an eagle could have some purposes, of course-- say, learning about eagles, or maybe solving some problem in aerodynamics. But fooling an observer under some variety of the Turing test is not a sensible purpose, as far as I can see. As far as I'm aware, there's nothing like this in the sciences.

Incidentally, if you go back to the 17th and 18th centuries, scientists were very much-- who we call philosophers, but they didn't distinguish-- were greatly intrigued by automata, by machines, for example automata that simulated, say, the digestion of a duck. But their goal was to learn about ducks, not to answer meaningless questions about whether machines can really digest, as proven by the fact that they can fool an observer, who's, just say, looking behind a screen or something, however they may be doing it.

Well this is one of the respects in which, in my opinion at least, there's been regression since the 17th and 18th century. Chess-playing programs are a perfect example of this in my opinion. Well, anyway, that's two approaches to the first question of unification, the deep one.

What about the more shallow question of unification, that is the question of unification among the various branches of the cognitive sciences? Is there any real commonality among them? Well, that doesn't seem entirely obvious.

So take what you might call physical organs below the neck, metaphorically speaking, meaning not the ones we call mental. Say, take the circulatory system, and the digestive system, visual system, and so on. These are organs of the body in some reasonable sense viewed abstractly. And that's a normal way to study any complex system. The study of these abstract organs does fall together at some level, we assume, presumably at the level of cellular biology. But it's not at all clear that there's any other level, that is that there's some kind of organ theory at which these systems are unified. Maybe, but it's not obvious.

Well, the brain is scarcely understood. But to the extent that anything's understood about it, it appears to be an extremely intricate system, heterogeneous with all kind of sub-parts that have special design and so on. And there's no special reason to suppose that the mental organs are unified in some kind of mental organ theory above the level of cellular biology, that is above the level at which they may fall together with the circulatory system and so on. If there is no intermediate level, no such unification, then there is no cognitive science, hence, there was no cognitive revolution, except for similarities of perspective, which could be quite interesting in fact.

Well, there are various speculations about this question. In fact, they go right back to that 1956 meeting. One general line of speculation is that the mind is what's called modular, that is that mental organs have special design, like everything else we know in the physical world. The other approach assumes uniformity, that there are what are called general mechanisms of learning or thinking which are just applied to different domains. They're applied to language, or to chess, and so on. And as I say, these differences, they didn't exactly surface, but they were more or less implicit at the 1956 meeting.

The paper on generative grammar took for granted modularity, that is a rich specific genetic endowment, the language faculty, special system, and no general mental organ theory. The Newell-Simon approach, the study of the logic machine, that took for granted the opposite position, that is that there are general mechanisms of problem solving which are quite indifferent as to subject matter-- in their case, theorem proving, but equivalently chess, or language, or whatever. And you can trace these different intuitions back quite a way, back hundreds of years in fact. They're very lively today, in the discussions about connectionist work for example.

My own view is that the evidence is overwhelmingly in favor of the highly modular approaches. There are only a few areas where we have any real insight. And these invariably seem to involve quite special design. Possibly there are some unifying principles. But about these, if they exist, nothing much is understood. Again, that's a minority view, but I can only tell you the way it looks to me. Even internal to the language faculty there seems to be quite highly modular design.

So let's distinguish roughly between what we might call grammatical ability, meaning the ability to produce properly formed expressions, and conceptual ability, that is to produce expressions that somehow make some sense, and pragmatic ability, that is the ability to use expressions in some way appropriate to the occasion. Well these capacities have been discovered to be dissociated developmentally-- like one can develop and the others not-- and selectively impaired under injury. And their properties are quite different. And they seem to be radically different from what we find in other cognitive systems. So even internally, just at the most gross look, we seem to find modularity. And when you look more closely into subsystems of language, they just seem to behave differently.

I don't see why this should surprise anyone. It's exactly what we find about every other biological organism, so why shouldn't it be true of the brain, which is maybe the most complex one around. But it's considered a very controversial position. Again, I stress, as far as I know, all the evidence supports it and it shouldn't be surprising.

Well, just how modular is the language faculty? That's an interesting question. It doesn't seem so special that it's linked to particular processing systems. Here the study of sign language in the last decade or so, maybe 20 years, has been quite interesting. Rather surprisingly it seems that sign language, the normal language of the deaf, sign language is localized in speech areas of the brain in the left hemisphere, not in the visual processing areas, which you might expect since you see it, they don't hear it. That's shown pretty well by aphasia studies, and again ERP studies. Also turns out that children who are acquiring sign language in a natural setting go through steps, at least in the early stages which have been studied, that are remarkably similar to the acquisition of spoken language-- timing and everything else. They even override obvious iconic properties of visual symbols. So pointing is overridden and given a symbolic interpretation.

Furthermore, acquisition of sign language turns out to be quite different from acquisition of communicative gestures. There's even one case on record of spontaneous invention of sign language by three deaf children who had no environmental model and no stimulus at all as far as anyone can be determined. That's the perfect experiment. The system that these three kids constructed has very standard properties of natural language, natural spoken language, and the developmental stages also appear to be similar. That looks like about a pure expression of the language faculty unaffected by experience. And it suggests that however modular the system may be, it's not so modular that it's keyed to a particular mode of articulation or input.

What's now known suggests that the language faculty simply grows a language, much as other organs of the body grow. And it does so in a way largely determined by its inner nature, by properties of the initial state. External events doubtless have some kind of an effect, like I'm talking English not Japanese. But a rational Martian scientist who is looking at us would probably not find these variations very impressive. And the language that grows in the mind appears to be at least partially independent of sensory modality. Well at this point we're moving towards serious inquiry into modular architecture of mind and brain in some of the few areas where it's possible to undertake it and with some surprising results.

Well one last element of skepticism about the cognitive revolution has to do with just how much of a revolution it was. In fact, to a far greater extent than was realized at that time-- or is even generally recognized today-- the shift in perspective in the 1950s recapitulated and rediscovered some long-forgotten ideas and insights from what you might call the first cognitive revolution, which pretty much coincided with the scientific revolution of the 17th and 18th centuries.

Now the Galilean-Newtonian revolution in physics is well known. The Cartesian revolution and its aftermath in psychology and physiology is not very well known. But it was quite dramatic and it had important effects, had important consequences. It also developed some of the major themes that were taken up again since the 1950s, primarily in the areas of language and vision, which are perhaps the most successful branches of contemporary cognitive science as well.

Now this first cognitive revolution, though it had a great deal of interest and there's much you can learn from it, it did face barriers, and in fact they were insuperable barriers. So for example, in the study of language it came to be realized soon enough that human language is somehow a process, not just a dead product. It's not a bunch of text, it's some process that goes on. Furthermore, it's a process that makes infinite use of finite means as Wilhelm von Humboldt put it in the early 19th century. But no sense could be made of these ideas, and the inquiry aborted, aborted for well over a century.

By this century the formal sciences had provided enough understanding of all of this so that many traditional themes could be productively engaged, particularly against the background of other advances in science and engineering and anthropological linguistics. And one could understand quite well what it meant to make infinite use of finite means.

Well, let's turn to a quick look at some of the kinds of questions that arise and some of the kinds of answers that can be offered with respect to language specifically. So take something that everybody knows, start with something really trivial, very simple. So take a simple phrase like, say, "brown house". You and I have generative procedures which are more or less similar. And these generative procedures determine what we know about this expression. For example, we know that it consists of two words, we know that those two words have the same value for most speakers. We know also that if I point to something, I say that's a brown house, what I mean is its exterior is brown. If somebody paints a house brown, we know that they painted the exterior brown. If they painted the inside brown, they didn't paint the house brown. If you see a house, you see its exterior. So we can't see this building as a matter of conceptual necessity-- if we were standing outside, we might.

And the same is true of a whole mass of what are sometimes called container words, like box, or igloo, or airplane. You can see an airplane, for example, if you're inside it, but only if you can look out the window and see the wing, or if there's a mirror outside which reflects the exterior of the airplane, or something like that. Then you can see the airplane, otherwise not. The same is true with invented concepts, even impossible concepts. So, say, take a spherical cube. If somebody paints a spherical cube brown, that means they painted the exterior brown, not the interior. So a house and all these things are exterior surfaces, which is kind of curious to start with.

However, they're not just exteriors. So suppose two people, say, John and Mary, are equidistant from the exterior regarded as a mathematical object. Suppose they're equidistant from the exterior, but John is outside it and Mary's inside it. We can ask whether John, the guy outside, is near the house, or whatever it is. And there's an answer depending on current criteria for nearness. But we can't ask it about Mary. We're not near this building, no matter what the criteria are. So it's not just an exterior. It's an exterior plus something about a distinguished interior.

However, the nature of that interior doesn't seem to matter very much. So it's the same house if you, say, take it and fill it with cheese or something like that. You change the interior, it's the same house, hasn't changed at all. And similarly you can change the walls around and it stays the same house. On the other hand, you can clean a house and not touch the exterior at all. You can only do things to the interior. So somehow it's a very special combination of an abstract though somehow concrete interior with an abstract exterior. And of course the house itself is perfectly concrete. Same is true of my home. That's also perfectly concrete, but in a quite different way. If a house is a brown wooden house, it has a brown exterior surface, but it doesn't just have a wooden exterior. It's both a concrete object of some kind and an abstract surface, as well as having a distinguished interior with weird properties.

Well, proceeding further we discover that container words like, say, house have extremely weird properties. Certainly there can't be any object in the world that has this combination of properties, nor do we believe that there is. Rather, a word like "house" provides a certain, quite complex, perspective for looking at what we take to be things in the world. Furthermore, these properties are completely unlearned, hence they must be universal. And as far as we know, they are. They're just part of our nature. We didn't learn them, we couldn't have learned them.

They're also largely unknown and unsuspected. Take the most extensive dictionary you like, say the big *Oxford English Dictionary* that you read with a magnifying glass, and you find that it doesn't dream of such semantic properties. Rather, what the dictionary does is offer some hints that allow a person who already knows almost everything to pick out the intended concept. And the same is true of words generally, and in fact the same is true of the sound system of language. They are largely known in advance in all their quite remarkable intricacy. Hence, they're a universal, a species property. The external environment may fix some details the way the *Oxford English Dictionary* can fix a few details. But language acquisition, and in fact probably a good deal of what's misleadingly called learning, is really a kind of biological growth where you fill in details in an extremely rich predetermined structure that just goes back to your nature, very much like physical growth or, as I think we should say, other aspects of physical growth.

Well, if you look at the structure of more complex expressions, then these conclusions are just reinforced. Again, let me take pretty simple cases to illustrate. So take the sentence, say, "John ate an apple." and drop out the word "apple" and you get "John ate." And what that means, everybody knows, is John ate something or other. So if you eliminate the phrase from the sentence, you interpret it as something or other. "John ate" means John ate something or other, more or less.

Take the sentence, little bit longer but not much, "John is too stubborn to talk to Bill." Well that means John is so stubborn that he won't talk to Bill. And just as in the case of "John ate an apple.", drop the last phrase, "Bill", and we get "John is too stubborn to talk to." And by analogy, or by induction, if such a thing existed, people ought to interpret it as meaning John is so stubborn that he won't talk to someone or other. However, it doesn't mean that. John is too stubborn to talk to means John is so stubborn that nobody's going to talk to him, John. Somehow it all inverts if I drop the last word.

Suppose you make it a little more complex-- "John is too stubborn to expect anyone to talk to Bill." Well, that means John-- you have to think here maybe a little bit already, because parsing is not easy and quick-- "John is too stubborn." it means John is too stubborn for him, John, to expect that anyone will talk to Bill. Now drop "Bill" again. Now we get "John is too stubborn to expect anyone to talk to." And reflect for a moment, and you'll see that the meaning shifts radically. It means John is so stubborn that somebody or other doesn't expect anyone to talk to him, John. Everything shifts.

Now all of this is inside you. Nobody could never learn it, nobody could possibly learn it. None of it is mentioned, even hinted at in fact, even in the most comprehensive grammar books. It's known without experience. It's universal as far as we know. In fact, it better be universal or there are genetic differences crucially among people. It just grows out from our nature.

Now a huge variety of material of that sort has come to light in the last 30 years or so as a direct consequence of something new, namely the attempt to actually discover and to formulate the generative procedures that grow in the language faculty, the brain, virtually without experience. That had never been done before. It was always assumed that it's all obvious. How could there be anything complicated here? If you go back, say, to the early '50s, the problem about language acquisition was supposed to be why does it take so long. Why does it take so long for such a trivial thing to happen? Why do children need so much exposure material to do this simple process of habit formation?

As soon as you try to describe the facts, you discover that it's quite the opposite. There's extremely intricate things that are known even in the simplest cases. There's no possible way of picking them up from experience. And consequently, they're universal, they're part of our nature. These are just part of the generative procedures that the brain determines.

Well, what do these generative procedures look like? Apparently, they look something like this. A language can be looked at as involving a computational procedure and what you can call a lexicon. A lexicon is a selection of concepts that are associated with particular sounds. The concepts appear to be pretty much invariant. They're things like house, let's say. And they have very weird properties, which means that they must be invariant. Because if they have very weird properties, they've got to be unlearned, which means that they're going to be invariant, just part of our nature. The sound variety also seems quite limited.

On the other hand, the association between a concept and a sound looks free. So you can say tree in English, and baum in German, and so on. So that tells us what one aspect of language acquisition is-- it's determining which of the given set of concepts, with all their richness and complexity, is associated with which of the possible sounds with their limited variety. Notice that that's a pretty trivial task. So you can imagine how a little bit of experience would help fix those associations if the concepts we're given and the structure of the sound was given.

Well, the lexicon also contains what are sometimes called formal elements, things like inflections, like tense, or plural, or case markers of the kind that you find, say, in German, or Latin, and so on. And languages do appear to differ quite substantially in these respects. However, that seems to be mostly illusion-- at least, the more we learn, the more it looks like illusion. So for example, take, say, cases. English doesn't have them. You've got to memorize them when you study German and Latin.

But that looks illusory. It appears to be the case that English has them fine, in fact exactly the way Latin has them. It's just that they don't happen to come out the mouth, they're not articulated. They're in the internal mental computation. And their effects, which are quite ramified, are detectable. But that part of the computation just doesn't happen to link up to the vocal tract. So the computation is running along, but where it gets spelled out of sounds, it just doesn't have to look at this stuff.

Well, that's another aspect of language acquisition-- that is, to determine what aspects of the computational system receive an external articulation. There are also some differences among these formal elements that also have to be fixed. But they appear pretty narrow.

Notice, incidentally, that if you have a very intricate system with just a few possible differences in it, which could be right at the core somewhere, making those changes may produce something that looks phenomenally quite different. Just look at it from the outside and they look radically different, even though it's basically the same thing with really only trivial differences. And that's pretty much what languages seem to look like.

Well, what about the computational system-- that's the lexicon. What about the computational system? It's possible that it's unique, that is it's invariant, there's just one of them, there's nothing at all to be learned. It's not certain, but it's at least possible, in fact plausible. If this picture is on the right track, notice we're very close to concluding that there's only one human language, at least in its deeper aspects.

And incidentally, going back to this rational Martian scientist, that's what he would have guessed in the first place. Otherwise, acquisition of these highly intricate systems on the basis of fragmentary experience and the uniformity of the process would just be a miracle. So the natural assumption is, well, there's basically only one of them, and it's fixed in advance. If, in fact, acquisition of language is something like a chicken embryo becoming a chicken or humans undergoing puberty at a certain age, then it's at least within the realm of potential scientific understanding, not a miracle. And that's the way it seems to be-- the more we learn, the more it looks like that.

More precisely, there does seem to be a certain variety of languages, apparently a finite variety. That it, it's a good guess now that each possible language that the language faculty makes available, each possible language is fixed by answering a finite number of fairly simple questions. You can look at the system as being sort of like a complex wired-up thing which is just given with a switch box associated with it and a finite number of switches-- you can flip them up, you can flip them down. You do that on the basis of the answer to one or another of these questions, and they've got to be pretty simple questions.

If we could understand the way that works-- bits and pieces of it are understood-- suppose we could understand it. Then we ought to be able to, say, deduce Hungarian by setting the switches one way and deduce Swahili by setting the switches a different way. The questions must be easily answered for empirical reasons. Children do it highly efficiently and very fast with very little evidence, pretty much the way they grow. It follows that languages are learnable.

That is, if this is correct, it follows that, as I mentioned before languages in fact are learnable, which is an empirical discovery and a pretty surprising one. No biological reason why it would have had to be true, but it looks as though it may be true, largely because the variety is so restricted. Well, assuming that to be true, we now have the following situation-- languages appear to be learnable, which is surprising, but languages appear to be unusable, which I don't think is surprising at all. What that means is that only scattered parts of them are usable, and of course those are the parts we use, so we don't notice it.

Now actually this unusability property may be somewhat deeper than what I suggested. Since the origins of modern generative linguistics, there have been attempts to show that the computational system is constrained by certain very general principles of economy which have a global character to them. And there's some recent work that carries these ideas a long step forward. It's still pretty tentative-- in fact, it's unpublished. But let me sketch some leading ideas that indicate where things might go.

Suppose we think of a linguistic expression in something like the following way. The language faculty selects a certain array of items from the lexicon-- that's what it's going to talk about. And it begins computing away using its computational procedure, which I'm now seeming to be invariant across languages. It computes in parallel, just picks out these items and goes computing along with them. It occasionally merges pieces of the computation-- so you bring them together and maybe it picks something out and starts going in parallel again. At some point in the computation, they've all been merged-- it just keeps computing. It proceeds merrily on its way, in other words. At some point after they've merged, the language faculty decides to spell it out, meaning to provide instructions to the articulatory and perceptual system. Then the computation just keeps going on. And it ultimately provides something like a representation of meaning, which is probably to be understood as instructions for other performance systems.

Well, proceeding in this way, the computation ultimately proceeds to paired symbolic expressions, to paired output, instructions for the articulatory perceptual apparatus on the one hand and instructions for language use, what are called semantic sometimes. So these are the two, let's call them, interface levels of the language faculty. The language faculty seems to interface with other systems of the mind, performance systems, at two points, one having to do with articulation and perception, the other having to do with the things you do with language, referring to things, asking questions, and so on, roughly semantic. So let's say that's what happens, looks like it does.

Now of these various computations, only certain of them converge in a certain sense. The sense is that the two kinds of outputs, the phonetic instructions-- the articulatory perceptual system and the semantic-- actually yield interpretable instructions, like they might compute along and end up with an output that doesn't yield interpretable instructions. In that case, if either one of them fails, we say that it doesn't converge.

Now looking just at the convergent derivations, the conversion computations, those that yield interpretable paired outputs, some of them are blocked by the fact that others of them are more economical-- that is, they involve less computation. Now here you've got to define amount of computation in a very special and precise sense, but it's not an unnatural sense. Well, if you can do this, it's going to turn out that many instructions that would be perfectly interpretable just can't be constructed by the mind, at least by the language faculty, because the outputs are blocked by more economical derivations.

What is a linguistic expression then? Well, a linguistic expression is nothing other than the optimal realization of certain external interface conditions. A language is the setting of certain switches, and an expression of the language is the optimal realization of universal interface conditions, period. All the work would now be done by special properties of the nature of computation and very special notions of economy, which do have a global character. They have a least effort character, but of a global sort.

Well, it turns out that quite a variety of strange things can be explained in these ways in terms of a picture of language that is really rather elegant. It's guided by conditions of economy of derivation with very simple and straightforward operations. They're not only elegant, but they're pretty surprising for a biological system. In fact, these properties are more like the kind what one expects to find, for quite unexplained reasons, in the inorganic world.

Well, these economy conditions, as I mentioned, they have a global character to them, though not entirely. They have a tendency to yield high degrees of computational complexity that would render large parts of language unusable. Well, that's not necessarily an empirical objection. We then turn to an empirical question-- are the parts that are rendered unusable the parts that can't be used. We know there are plenty of parts that can't be used, so if it turns out that parts of language involve irresolvable computational complexity, irresolvable by a reasonable device, and those are the parts you can't use, fine. That's a positive result, not a negative result.

Well, what we might discover then, if this continues to look as promising as I think it does now, what we might discover is that languages are learnable-- because there isn't much to learn-- that they're unusable in large measure, but that they're surprisingly beautiful, which is just another mystery if it turns out to be true. Thanks.

[APPLAUSE]

I understand there's some time for discussion. So if anybody feels like, there's a couple mics up here.

**AUDIENCE:**

Hi.

**CHOMSKY:**

Can't see.

**AUDIENCE:**

Should I go ahead?

**CHOMSKY:**

Oh, OK. Do you have a mic? Oh, you can go ahead. Go ahead.

**AUDIENCE:**

OK. You touched on my question with your last couple of statements. How well do you think what we want to express maps with what is expressible through language? Do you think there are maybe easily computable things that can't get through? And do you think this has some implications for other forms of expression-- music, et cetera?

**CHOMSKY:** Well, that's an interesting question. The question is how much of what we think can we articulate. In order to have an answer to that question, we'd have to have a way of getting at what we think that's independent of what we articulate. So we'd have to have a way of getting at what we think which is independent of language. And that's hard. There's really no way to get at what we think other than what we can articulate somehow.

Now nobody knows anything about this, so your guess is as good as anyone else's. But if you introspect about it, just to the extent that introspection means anything, it looks as though the kind of thing you're talking about happens all the time. So it's a common experience at least, whatever it means, in a conversation for you to say something and realize that's not what you intended. And you try it again, you listen to that one, and maybe it was a little closer but not quite. And then you try it again, or maybe somebody suggests something, finally you get it out. And, yeah, that was what you intended.

Well, it's hard to make sense out of that experience unless there's something going on in there that's not expressed linguistically, some kind of the thought or something, and you're trying to figure out a way to articulate it. However, what that means, nobody knows, because we don't have any way of getting at the thoughts other than by the articulations. If you can figure out a way of doing it, that would be interesting. Phil?

**AUDIENCE:** [INAUDIBLE]

**CHOMSKY:** Yeah, Yes I will. I just talked about speech and hearing. I also talked about sign and seeing. But the question is, well what about writing. That's obviously an input, doubtless. Whether it's a natural system is not so obvious. There isn't, as far as I'm aware, there's no evidence-- maybe someone can correct me if I'm wrong, but I don't know of any evidence that children can learn language from writing. They can learn it from sign, clearly, and that's just like speech. They can learn it from hearing. But as far as I know, learning from writing is always secondary. That is, it's mapping some visual sign system or some auditory sound system into another mode of representation. And learning it is not so simple-- it takes some work. It does feed the same systems, as far as we know, so it would be left hemisphere represented, and so on and so forth. You get left hemisphere aphasia, that will cause writing problems sometimes.

**AUDIENCE:** [INAUDIBLE]

**CHOMSKY:** Right, right, that's absolutely true. Writing gives you all sorts of options that speech and sign don't give you. Because you can look back, and you can think about it, and you can make much more complex expressions. And if you look at what's in writing, it has much more complexity than what would ever happen in a left to right system. I mean both speech and sign are restricted by the fact that time is linear. It's left to right processing, so all kinds of memory constraints start to enter and all sorts of other computational constraints. A parser for speech and sign has got to be left to right. A parser for writing, you can take your time. You can pore over a sentence for a long time and try to figure it out. Sometimes you even can, sometimes you can't. Depends who you are. If I read a text in postmodern literary criticism, I can look at the sentence for about a year and I still don't understand it.

[LAUGHTER]

But presumably you can at least figure out that it's got some grammatical structure.

[LAUGHTER]

And there are some interesting questions about-- there's some work actually suggesting that the creation of writing systems may have had an impact even maybe on language change, in allowing certain more complex options to be developed and so on. I don't think that's impossible. You should really ask other people here about that, because they know a lot more about it than I do, including some of them who have worked on it.

But another point that ought to be mentioned is that there are invented systems-- all of you know them. When you study at MIT, you're learning invented, created systems all the time. Now those things are called languages, like, say, the language of mathematics or something, but that's just a metaphor. They're very different from languages. Even the language of ordinary scientific informal discourse is radically different from language when you begin to look at its properties. All of these weird, complex properties of language, like what I mentioned about house-- when you develop a system that you want to use for understanding the world-- you know, science-- you try to divest the concepts of all of these strange properties of natural language.

So for example, the other day, I can't say I read, let's say I looked at a paper in *Science* on the granular structure of matter. It was asking whether a heap of sand is a liquid or a gas, or is it some other form of matter, something or other. Well, I couldn't understand the article, but I could see what he's talking about. But the point is, the author of the article didn't care how the term liquid is used in English or in any other language-- that didn't come up. What he cared about is the special technical concept created in the course of this effort to understand the world, which may have the same sound as the English word "liquid" but is going to have very different properties.

So these things that are constantly created, you might say they are accretions to language. And the interactions with them are typically in written form. But they're very different systems. I expect that they're created by totally different faculties of the mind. If people would like to come up to the mics, it would be a little easier.

**AUDIENCE:** One of the Japanese philosophers said the cerebral physiology, the physical study of brain, showed to some extent the physical computational processing of languages is very different--

**CHOMSKY:** Of what? Physical computational processing of?

**AUDIENCE:** Processing of language of a computer is very different from that of a physical brain. And he also said it's safer for us to call a machine a computer rather than call a machine artificial intelligence. On the other hand--

**CHOMSKY:** I'm sorry, I didn't entirely understand. If you look at the computer processing of language--

**AUDIENCE:** Yeah.

**CHOMSKY:** --the way the computer works is simpler than the brain?

**AUDIENCE:** The computer processing system of language is different--

**CHOMSKY:** Different from the way the brain does it?

**AUDIENCE:** Yeah.

**CHOMSKY:** Yeah, that doesn't really mean anything. It's like saying that if you build a bulldozer it's different from a person digging a ditch. Yeah, if you write a computer program to do something, whether it happens to be like the way humans do it-- it depends what your purposes are. If you're writing a computer program to simulate some aspect of human thought, say translating English or something, then it better be like the brain. Otherwise, you're just doing the wrong thing. It's like for creating a model for some other physical system and getting it wrong. On the other hand, if you're building a computer program to do something because you just want the thing done, well then it's like constructing a bulldozer to dig a hole-- you don't care whether people do it that way or not.

**AUDIENCE:** And on the other hand, some current theory of syntax too much care for the--

**CHOMSKY:** Is too?

**AUDIENCE:** Some current theory of syntax care too much for the computational rationality rather than to think for the [INAUDIBLE] of our brain. And--

**CHOMSKY:** Sorry, I didn't understand. Somebody understand who can-- somebody say it again so I'll understand it? That just shows that the generative procedures aren't always identical.

[LAUGHTER]

**AUDIENCE:** OK, then.

[INAUDIBLE]

**CHOMSKY:** Oh, the theories of languages or the computer theories?

**AUDIENCE:** OK, in other words, they care they can apply the grammatical theory to the computer. They put priority the application of the computer rather--

**CHOMSKY:** OK, let me see if I understood. There's three different things here. One is what the brain is actually doing. A second is what some linguist is claiming the brain is actually doing, a linguistic theory. A third is a computer program which is doing something or other. Now we want the first two to be alike. If they're not, got the wrong theory and you go back to the drawing board. When you're working with a computer program, you just have to decide what you're trying to do. If you're trying to simulate, you're trying to give a model, you're trying to understand what the physical system you're interested in is actually doing, then you want your computer program to identify processes and states which are real. For example, you want the program to fail where the person fails. If a person has a parsing failure, you want your parser to fail, and so on.

There's no issue of rationality that I can see. A computer program and the brain, they're not rational any more than the legs are rational when you walk. I mean they do what they do. The brain is designed to do what it does. A computer program that's trying to capture aspects of that will, to the extent that it's correct, do what it does the right way. But the issue of rationality doesn't arise.

**AUDIENCE:** OK. Another question is in your current theory of syntax you assume that modules are principals. And I wonder you assume there exists such a kind of system in actual?

**CHOMSKY:** In reality, yeah.

**AUDIENCE:** In our brain.

**CHOMSKY:** Right. So when you construct a theory of syntax, you assume you're probably wrong, like when you're doing any kind of science. But you make a guess that says, I hope I'm right. And then you check it out and see where you went wrong, and if somebody can make it better, and so on. But if somebody writes a paper in some journal, they are, if they're serious, claiming, look, this is the way it works, this is the way the brain really works. Probably they're wrong, because we usually turn out to be wrong. But you try to get better the next time around.

On the other hand, you certainly are claiming that's the way it really works. Same if you construct a theory of organic molecules or anything else. You're climbing, yeah, that's the way it really works. I know I'm probably wrong, but I'll claim that's the way it works and then I'll see if I make some progress.

**AUDIENCE:** I really enjoyed your comment when you said there was one human language. And you also indicated that if we were Martian and we came to the Earth we would hear one sound in a sense. So I was wondering, with all of the diversity that we have in the world today and historically-- with Asia being one way linguistically, and Africa another, Europe another-- and with inflectional the languages that we've developed from Indo-European influences and the agglutinative languages that come from Africa and the Creoles, from the Caribbean area, in the future, say 100 years from now, could you comment on which tendency would you feel that we would more likely be to go in, that of one human language for the Earth or diversity with many of them?

**CHOMSKY:** Well, see when I say there's one language, I don't mean that we should try to all speak Esperanto or something. What I'm claiming is the actual diversity, even the possible diversity-- there's a lot of possible languages around that just haven't come along for historical accident-- but if you could consider the possible diversity, what I'm saying is a rational Martian scientist who looks at us the way we look at frogs would say, well, there's only one of them. If we look at frogs, they're all frogs.

[LAUGHTER]

If a frog looks at frogs, they're all wildly different from one another. In fact, a frog would just take for granted that every reasonable thing is a frog and wouldn't even ask any question about what it's like to be a frog. A fraud would presumably be precisely different, interested in whatever minuscule difference there may be among frogs that tells them act this way or that way to them. They don't worry about the fact that they're all frogs, that's taken for granted.

Or if you're doing a genetics experiment, let's say, you pretend that all the fruit flies are identical. I mean if somebody really asks you, you'll say, yeah, they're a little bit different here and there, but they're basically identical for all I care. And whatever's different about them, I forget about it. And again, from the point of view of the fruit flies the world may look quite different.

Now we're no different from any other organism I presume-- I think we're part of the natural world. So we take for granted that anything around is a human being. I mean, if other people use the word "house" with this weird collection of properties that we use, we don't even notice it-- that's just like breathing, because that's just the way any reasonable creature is. And if you go over and you hear somebody speak Swahili or something, that looks wildly different.

On the other hand, I'm suggesting that if a Martian were to look at us, or if we were capable of looking of looking at ourselves, as we can-- you abstract yourself away from being a human being and you try to become. You can't get out of your own skin completely, but you try. Look at us the way we would look at frogs or fruit flies, and I think what you discover is we're remarkably alike. All of this diversity is really pretty superficial. It goes back to very similar structures, in fact so similar that this diverse range of languages may very well have only one computational procedure and only slight variations among the set of concepts with all their richness and intricacy, most of which is undreamt of-- nobody even knows what it is. That probably is going turn out to be alike.

Now of course there will be differences. And in our lives as human beings, we'll be very much interested in the differences. But in our lives as scientists we'll say they're all the same. That's what I mean.

**AUDIENCE:** I see what you mean. In other words, were you saying that-- you made the example of the container, the container [INAUDIBLE].

**CHOMSKY:** Yeah.

**AUDIENCE:** I believe you said if we look upon it from the exterior then we see that--

**CHOMSKY:** Well we have a very weird way of looking at containers, a way which is extremely complex, so complex in fact that there couldn't be a physical object that has the collection of properties that we assume when we talk about them. Well, it doesn't mean we're confused or anything like that. It just means that the language resources we have compel us to look at the world from a very strange perspective. And we all do it that way because we're human. Frogs do it their way because they're frogs. Now there are also slight differences among us, like English and Swahili. But they look slight.

Now about the diversity of languages, chances are it's going to move towards uniformity. But that's for completely other reasons. That's because a lot of people get murdered, and power spreads, and that sort of thing. So languages are in fact disappearing, but that's like the way biological diversity is being reduced.

**AUDIENCE:** Thank you.

That's the question I want to ask you about, the future of language. I talk to people from other countries and I ask them what they listen to on the radio in terms of music. And it's all English-based pop music. The publishing industry seems to be going in the direction of English language universal throughout the world. There are countries that have set up organizations to protect their languages. I know the French government has such an organization, particularly when-- what was it-- pinball machines in France became [FRENCH]. Are we going toward a future where English will be the universal language?

**CHOMSKY:** Well that's a question at a different dimension altogether. That's a question of power. And a different set of issues arise. My guess is not. In some respects the world's becoming more diverse.

For example-- I don't have any statistics, maybe somebody knows-- but I wouldn't be at all surprised if the number of, say, scientific papers being written in German and Japanese is increasing, not decreasing. You go back to, say, early '50s they'd probably be writing in English, because the US is so overwhelmingly dominant, if you're not writing in English, you don't exist. But I'm sure some of you know-- maybe Phil knows or something-- but I suspect it's more diverse now. And if real development takes place in other places, I would expect it to continue to be diverse in that respect.

Also, take, say, the European Community, or take Europe altogether. While Europe is at one level unifying, it's also splitting all over the place. We read about it in Eastern Europe where there's a different country emerging every village by now. But it's happening in Western Europe too. I mean Scotland is moving towards, has a big pressure towards, independence.

Actually, it's even happening in the United States. I was giving some talks in Alaska a couple of weeks ago and I discovered, to my amazement, that the governor of Alaska was elected on an independence ticket. Apparently everybody thinks he's a total crook, but they voted for him because he was calling for Alaskan independence, and that's what they want. They want to get away from this business down here. So I think one also finds very disintegrating factors in the world system towards people trying to identify communities that they can belong to, feel part of, or something like that. And there's been a lot of revival of languages that were dying. I think it's awfully hard to predict. It depends on complicated economic, and social, and political processes that we don't have any-- at least I don't have much-- understanding of.

**AUDIENCE:** I heard a report of evidence on brain waves that if you're hearing in English or seeing in English a table, your brain wave pattern is different than if your natural language is German. Is that because we always give them an oral element, or is it a difference in visual management, or what causes this difference?

**CHOMSKY:** Well, I don't know those results, but if there's anything like that I suspect it would be picking up something about the input system. First of all, I don't think there's any technology or understanding around that would enable you to identify the conceptual structure of English versus the conceptual structure of German in anything like, say, electrical activity. Again, I'm sure some of you know much better than I do. But if there's anything involving, say, electrical activity of the brain that is distinguishing a German word from an English word, it's probably picking up the physical input, like the articulation, which of course are radically different.

**AUDIENCE:** Yes, these were experimenting in which they were asked mentally to think sentences--

**CHOMSKY:** Think about something-- but when you think about something--

**AUDIENCE:** And they're also given items to see.

**CHOMSKY:** Yeah, but when you think about something-- see, but you can't think about something without beginning to get your motor system working. Try to think about a word without saying it to yourself.

**AUDIENCE:** I'd like to ask two questions. The first one is Alonzo Church said once that *Principia Mathematica* was as much a part of the English language as *Paradise Lost*. Now speaking from my modest personal experience, I can report that in an hour of an average mathematics seminar one is likely to hear more metaphors than in a month of a Iowa poets workshop. Your distinction, on the other hand, between language of science and natural languages, seems to turn around the prescriptive ideal that a scientist supposedly sets for his language. I would like to ask to what extent actual empirical data about linguistic competence of speakers of languages of science has been incorporated into the conclusion that you make.

**CHOMSKY:** Well, it can't be studied, because the notion languages of science is just a metaphor. It's like asking how much evidence about airplanes flying has been incorporated into the study of eagles. It's not a meaningful question. When scientists construct the system that they're working in, they may not think about it, but what they're actually doing is divesting the system of many of the complicated properties of natural language. And they're also adding to the system other properties that natural language doesn't have.

So for example, take *Principia Mathematica*, or let's say go back to Frege, which starts the thing off. The systems that, say, Frege and the beginning of modern philosophy of language and philosophy of mathematics, the systems that he created-- which of course he called languages, but that doesn't mean anything-- the systems that he created had properties, stipulated properties that language probably doesn't have at all, in fact presumably some quite critical properties. Like a Fregean system, the kind that is then developed in *Principia Mathematica*, that has a property of reference-- that is, words refer to things. So the numeral three refers to the number three, period. It's just made up that way. Well that doesn't seem to be way natural languages work. They don't have what are called logically proper names.

Back 40 years ago or so, philosopher Peter Strawson talked about what he called the myth of the logically proper name, pointing out, and I think accurately, that languages just don't have them. They have other ways of carrying out the act of referring. But if you're trying to use the language ways in mathematics or in physics, it would just be a total screw-up. So what you do is create systems which have properties, like reference between a symbol and the thing, and which of course pick the things to make some sense. You pick the things so that they're going to be part of the natural world. Natural language doesn't care whether they're part of the natural world or not. Natural language is not a system which is given to us with the intent of coming to an understanding of nature any more than walking is. And you've got to keep trying to overcome your intuitions when you do science.

I mean you look at the sun set every evening. You can't help but seeing the sun set. You can know all the astronomy you like, and when you look at it you see the sun set. There's no other way to see if, that's just the way you're constructed. But insofar as you're a scientist, you try to take yourself out of your skin-- I don't want to see the sun set, I want to see relative motion, and so on and so forth. And it's the same in constructing these so-called languages, whether it's Frege, or *Principia Mathematica*, or anyone else.

So there's no meaningful question about are they like natural language. No, they're difference. I mean it's perfectly true, as you say, that if you go to a math class, math lecture, people are going to use metaphors all over the place. Similarly, open the "American Journal of Mathematics" and the things that are called proofs are not proofs in anybody's sense. They're hints that a smart mathematician could use to make up a proof. If you tried to construct a proof, every issue of the "American Journal of Mathematics" would be as big as the Library of Congress or something like that. You can't do that. So of course you use metaphors and hints, and you appeal to people's intuition and understanding, and so on. But that's just because you're a human being trying to do science. Insofar as you're doing science, you will-- very consciously often-- diverge from properties of common sense or natural language, which are just biological capacities, just as you'll try to diverge from your perception of the sun as setting.

**AUDIENCE:** Perhaps I was calling for a more descriptive rather than prescriptive notion of proof in analyzing the language of mathematics. But I would like to ask you something--

**CHOMSKY:** But, see, there's no issue of-- look, there is no such thing as a descriptive study of the practice of the "American Journal of Mathematics". It's just not an interesting topic.

[LAUGHTER]

I mean you do it the best you can, with the least effort, to make sure that people understand more or less what you understand will get the idea. That's the way you do it.

**AUDIENCE:** Are you saying then that the relation of reference is radically different in natural language or perhaps even nonexistent?

**CHOMSKY:** Well, it appears to-- I mean this is a very controversial issue again, and I don't want to be glib about it. But my own view at least is that the relation of reference which is stipulated for formal languages doesn't even exist in natural language.

**AUDIENCE:** Thank you.

I liked the metaphor about setting switches. My switches were obviously set to English. I tried to reset them, and reasonably well, to Japanese and found that quite different. I wonder if your studies have given any insight into second language acquisition and perhaps ways to facilitate it?

**CHOMSKY:** Mine haven't. There are people here who work on second language acquisition. Suzanne Flynn, for example, is a person who's done interesting work on second language acquisition. And her work suggests that there are certain parts of the system that are just fixed and universal. Then there's other peripheral parts that you've got to set the switches. And her work suggests that at the adult level the distinction is still there, so that when you're learning Japanese as an adult you're still using the. See, you might assume that it all just gets assimilated-- once you set the switches, it's just what it is, and you can't distinguish the switches from the principles anymore. But according to her work, you probably can distinguish. So that's interesting.

On the other hand, how people learn second language, that seems to vary a lot. Like, say, take our department, little small department over at MIT with half dozen, eight faculty members. We happen to have in the department two people who are at opposite extremes in the human race in their capacity to learn second languages. One is me, who is at the extreme of total incapacity. The other is Ken Hale, who picks up second languages the way a child does. It's like a miracle. And how this happens, nobody knows, nobody has the slightest idea. People just differ.

**MODERATOR:** I would suggest we take one last one--

**CHOMSKY:** OK,

**MODERATOR:** --and then give you a break.

**AUDIENCE:** I hate to be the last question. I hope I don't disappoint anyone.

**CHOMSKY:** It's a real burden.

**AUDIENCE:** I was wondering if there's some practical difficulty with constructing a theory about language that requires language to articulate the theory. Is it self-referential--

**CHOMSKY:** Is that some sort of contradiction, using language to talk about language?

**AUDIENCE:** Is there a problem.

**CHOMSKY:** I don't think so. I mean you can get confused. It allows for confusion. But if you keep your head straight, there isn't a real problem.

For one thing, when a linguist talks about language, a linguist is doing something like a chemist talking about the world. That is, you're making up other systems which are going to have special properties that scientific systems have. You're going to use English words. But of course, if you go to a physics class, they also use English words. A physicist will use the word work and energy, but you're supposed to understand that they don't mean what you mean out on the street when you say it's too much work to finish my homework or something. You don't mean that, you mean some other thing.

And that kind of problem does arise. And you're right, it does arise in a somewhat greater form when the thing you're talking about is your own language. But it's mostly just one of those extra burdens of possible confusion that you have to try to sort out. There's no intrinsic difficulty.

[APPLAUSE]

Thank you very much.

[APPLAUSE]