Nineteen

LANGUAGE ORIGINATED IN SOCIAL BRAINS

Luc Steels

1. Introduction

The question of the origins of language is today one of the central and hotly debated areas in human-related sciences and is directly relevant to the question how and why are brains became so profoundly social. All disciplines have something to say, from anthropology and psychology to linguistics, artificial intelligence, and complex systems science. Tremendous progress has been made recently, but there is not yet a widely accepted consensus yet.[1] In my work, I have been engaged in different experiments with artificial systems (robots). They are endowed with several general cognitive mechanisms, for example categorization, associative memory, structure recognition, etc., and are then programmed to engage in language games, routinized local interactions which have a side effect that some sort of communication system with language-like features emerges. I have described these experiments in quite some detail in other papers and here I want to focus on the hypotheses that we tested and validated in these experiments.[2]

To develop an adequate theory for the origins of language, we must answer four critical questions:

(1) To participate in a language community you obviously need a highly sophisticated brain that can pull together a wide range of mental resources in the service of language. This configuration of mechanisms and knowledge structures is usually termed the language faculty. The first question is: What is the architecture and functioning of this language faculty, so that it can not only produce and interpret language but also participate in its creation? What does a “language-ready” brain have to be able to do?

(2) Language is a collective phenomenon, like paths and nests collectively built by an ant society or movements and trends in a stock market. It is a complex adaptive system shaped by many individuals, in turn influenced by the emergent language. This leads to the second question: What are the processes that govern the emergence and evolution of language at this collective level? For example, how does a lexicon become and remain shared, even if every language user may

invent new words or forget and misuse others? How can there be an increase in complexity, for example from a purely lexical language stringing together words without syntax, to a grammatical language that exploits syntax to express how the meanings of words are to be combined?

(3) Even supposing that we have a theory that can explain how “language-ready” brains create and coordinate languages of greater and greater complexity, a third question remains: Where does this language-ready brain come from? By what processes has it evolved? Can we find explanations that fit within the standard neo-Darwinian framework of genetic evolution by natural selection? Alternatively, are epigenetic processes at work, which rely on the recruitment and dynamic configuration of neural subsystems?

(4) The three questions above all relate to the issue how language may originate. They do not yet ask why this might have happened in our species and only in our species. This suggests a fourth question: What are the unique social and functional prerequisites that have supported the growth and evolution of language? Why did human beings start to speak?

All these research questions interact. Without an unambiguous idea of the language faculty, the study of its neurobiological embodiment and evolution is like groping in the dark. Without understanding how the laws of complex systems affects collective properties of language, we cannot be sure what the contribution of an individual language user can or needs to be, or what has to be available *a priori*. Without understanding what conditions might have prevailed at the dawn of language, we are at the risk of making the wrong assumptions of what mental mechanisms and strategies we can put into the models.

In the remainder of this paper, I discuss hypotheses for each of these four critical questions.

2. The Architecture of the Language Faculty

Everybody agrees that the language faculty includes three major subsystems:

(1) A *perceptual-motor system* that senses the world; segments incoming sensory data into different objects and events; extracts different features like size, speed, or loudness; and performs actions in the world, including the articulations required for speech;

(2) A *conceptual-intentional system* that sets and monitors communicative goals and conceptualizes reality for language or

interprets conceptualizations by mapping them back to perceived reality, and

(3) A *language system* proper that maps conceptualizations into utterances in language production and reconstructs conceptualizations from utterances in language parsing.[3]

Each of these subsystems is in itself extraordinarily complex, consisting of many different subsystems, which all interact intimately with each other. To build a theory of the language faculty means to elucidate what all these subsystems are doing, what kind of information they require and produce, and how they do it.

Even supposing that we manage to make all the processes in the language faculty operational, including the learning processes that fill in the language-specific inventories for each of the subsystems, we would still have no explanation for the *origins* of language. We need something more. We need to find out what mental processes cause *new* elements to enter into a language. I argue, and this is my first major hypothesis, *that we should see the introduction of new conceptual or linguistic material as similar to invention.*

I consider the elements of language as tools, and these tools have been invented —and are still being invented and adapted— by intelligent human beings, using similar skills as those required for the invention of other tools like hammers or medicine. Most linguistic inventions are not entirely new just as tool inventions in other domains. The first hammer was quite crude, a stone to push a pole in the ground. The human mind has this great ability to see how to improve a tool through experience. Some people must have noticed that you can get a better hammer when you introduce a handle, strengthen the head, cover the handle with some rubber coating against sweating, or change the size or shape of the head to fit better with a particular type of nail.

Sometimes an invention is lost or nobody understands it anymore, but then something else is invented to achieve the same functionality, perhaps in other ways. I argue that the same is true for concepts, words, and grammatical constructions. Speakers have invented them to deal with some particular need in communication and then continuously honed and fine-tuned them to achieve a higher rate of communicative success or to make the effort required for using them more manageable.

You might be surprised, but not all languages have grammatical markers for future tense. For example, in Chinese you have to hint at future tense by additional phrases or adverbs, like “yesterday” or “as I planned to.” No specific future tense marker is used with the verb. So how can a language ever develop a grammatical way to express future?

A good candidate for hinting at future tense is a verb for desire. If you say, “I want to write a letter,” you have not written the letter yet, you will do so in the future. Other candidates are verbs of movement (“I am going to London to see the queen”), verbs of possession (like Latin *habeo* (I have), which was the basis for the French *-ai* morpheme in *Je parlerai* (I will speak)), or verbs of obligation (like old German *skal* (owe), which was the basis for "shall"). Until about the sixteenth century, the English verb "will" was a main verb to express desire, as it still is in Dutch or German. If you say in Dutch, *Zij willen hun broer zien* (literally, They will their brother see), you mean that they want to see their brother and whether this will happen is uncertain. "Will" is still a main verb. In English, "will" has been grammaticalized into a future auxiliary verb. If you say, "She will write a letter," that person is going to do it. If you say "It will rain," nobody wants anything, "will" has lots its meaning of desire and shifted syntactic category from main verb to auxiliary. So that you do not say, "I don't will write him a letter," but "I will write him a letter."

Invention is only one side of the coin. Inventing a new tool for communication, such as a new word or a new grammatical marker, is fine enough. But if others cannot figure out its meaning and learn to use it themselves, the word will not propagate and it will die out. So learning (by understanding the intention) is crucial. It is even entirely possible that language learners unwillingly invent new features by believing that some aspect of an utterance, for example a particular word order or a particular intonation pattern, carries meaning, whereas it did not yet for the speaker. If the hearer then reuses the construction, another listener may pick it up and propagate it further in the population.

Invention requires that you can see the need for a tool and either create an entirely new one or adapt an existing tool for a new purpose. Intentional learning requires enough "theory of mind" to be able to guess what communicative goal somebody else wants to achieve and how she may have conceptualized the world, so that you can reconstruct the meaning and function of a tool's components, or what meanings and effects the words and grammatical constructions are supposed to have. A common reaction to the hypothesis that language originates through invention and intentional learning is that it requires too much intelligence and a conscious mind and so it assumes a much greater mental capability than we are normally willing to ascribe to early human beings or young children or maybe to most people period. I take this objection seriously.

Yet the kind of invention and intentional learning I have in mind is achieved by processes which do not involve consciousness nor impossible leaps of imagination. I am not at all thinking of intelligent conscious invention the way you would design a house or a new computer chip. Instead, the processes I have in mind take place at a subconscious level, even though you may occasionally consciously choose what your subconscious processes suggest.

Another common objection coming from some linguists is that language is not like a tool. They argue that the structure of language is not related to its function.[7] This is a surprising objection and counter-intuitive to non-linguists, because common sense tells us that language is a tool for communication and, derived from that, for representation and thought. I think this objection comes from not taking function and processing considerations into account. It is like looking at a car from a purely structural and descriptive point of view. Why it has wheels or a gear stick becomes impossible to understand. The function of the wheels becomes only obvious when you take into account that the car is designed to move. A gear stick makes sense only when you understand that a car has different "gears" to which the motor transmits power and that an engine's performance needs to be optimized by choosing the size of the gear depending on the speed of the car or the inclination of the road. The more you want to understand a car's components and the reason why they are installed, the more you have to delve into the operation and mechanics of the car. Why would this be different for understanding language and the language faculty that enables it?

3. Language as a Complex Adaptive System

My second hypothesis concerns collective dynamics. A few decades ago, the natural sciences discovered that quite often systems consisting of many elements interacting locally with a few others at the microscopic level could undergo a phase transition suddenly and spontaneously, in which a macroscopic structure emerges. For example, in the 1950s, Boris P. Belousov and Anatol M. Zhabotinsky discovered a chemical reaction, which spontaneously shows surprising spatial and temporal patterns. Order emerges out of chaos. The mixture of molecules they studied changes regularly among three colors, among three global states, like a clock. Poured into a dish, it produces the kind of beautiful spiral patterns, which are now on the cover of many books on chaos theory. Matter, which until then, people assumed to be dead, in the sense of moving to an equilibrium state and staying there, turned out to be full of activity, switching between unordered states, spontaneous order, and chaos, depending on contextual parameters. Molecules, which have only local interactions and no "awareness" of the global state, become coordinated.

The trick lies in a hidden positive feedback loop. Once things start moving in a designated way, the rest comes down like an avalanche. After investigators understood these phenomena and modeled them mathematically, scientists recognized similar self-organizing phenomena and discovered them in other chemical reactions and in physical systems like lasers or magnetic spin systems.

Self-organization is not the only morphogenetic process that creates structure in Nature. Another is selectionism, probably more familiar to many readers, as Charles Darwin appealed to selectionism to explain the origins of species. Selectionism, like self-organization, is completely general and applicable at many levels and to different types of systems, including economic or cultural systems. Another mechanism that has turned out to be quite useful is structural coupling, originally introduced by Humberto

Maturana and Francisco Varela. It occurs when two autonomous systems become intimately attuned to each other without a central coordinator or prior determination, but only because input for one is output for the other and vice-versa. We see structural coupling in the tight integration between different cells in the body or in the coordination between independently evolving organisms, which still strongly depend on each other.

Yet another mechanism is level formation. It explains how higher order structure may arise because individual units start to co-occur together and develop additional "glue" so that the whole becomes more than the sum of the parts. Given the ubiquity of all these generic morphogenetic mechanisms in Nature, scientists and mathematicians began to wonder whether any universal laws underlay them. This turns out to be the case and a body of knowledge has gradually developed into a true science of complex systems, which is still booming with excitement today, especially with the discovery of intriguing properties of evolving networks.

Is of all this relevant to our subject? I strongly believe so. This is my second hypothesis: *The same universal laws that explain the origins of order in natural systems are at work in socio-cultural systems, and especially in the origins and evolution of language*.[8] The elements at the microscopic level are the communicating agents, engaged in local verbal interactions with others and changing their internal states as a side effect of each interaction. Out of these local interactions, large-scale structures and behaviors emerge at the macroscopic level, the level of language itself.

The theory of complex systems suggests mechanisms for which we might look and add as ingredients to the artificial agent experiments. The universal laws make it possible to predict whether macroscopic structure will emerge. To take a mechanism discovered in some subfield of science, abstract from it the domain-specific details and then re-instantiate the same principles and apply them to the full glorious complicated detail of language is not so simple to do, but this is precisely the kind of exercise that I am advocating.

Here is, for example, how we can apply self-organization to the case of speech sounds. Suppose each individual agent produces sounds. If they ignored the others, no shared speech system would ever emerge. Suppose that agents align their speech sounds to those of others by slightly adapting their motor control programs and shifting their perceptual prototypes. This creates a hidden positive feedback loop and self-organization: The more some sounds are used, the more agents encounter them and because they align themselves to these sounds, a shared repertoire inevitably emerges and persists even if the population changes. So a global order, a shared sound system, appears, without one agent dictating to all others what sounds to use, and without any global overview by any of them, nor any prior genetic predisposition towards some sounds.[9]

Here is concisely how selectionism applies. Selectionism explains how "solutions" to a task can arise without explicit conscious intelligent design. It needs four ingredients: storage of the blueprints of solutions, instantiation of a solution from a blueprint, replication and recombination of blueprints with variation, and preferential choice to keep blueprints that lead to better solutions based on testing them for the task. Several ways exist in which we can apply this framework to language. One way, originally suggested by Richard Dawkins and elaborated by the linguist William Croft, is to view language elements as memes that get replicated.[10]

We can view the basic elements of language (concepts, words, constructions) as "solutions" to the problems that speakers and hearers encounter if they try to communicative successfully. This is similar to viewing the structures and behaviors of an organism as solutions to the problems of survival in a particular ecosystem. Blueprints for solutions (like lexical entries) are stored as part of the language faculty of each agent and instantiated every time a sentence needs to be formulated or interpreted. This is similar to the genome, which is a kind of blueprint that guides the development and behavior of an organism.

In both cases, we should not take the term blueprint too literally, definitely not like the architectural blueprint of a house, because contextual processes and self-organization also play a big role. As a side effect of learning, the blueprint for a particular communicative solution is stored in the language faculty of the individual, and invention may recombine existing solutions to build up more complex hierarchical structures.

Invention and learning play the same role as replication and recombination of genomes. Invention, intentional learning, and instantiation of blueprints in actual language use with unavoidable inaccuracies, errors, and slips of the tongue, all conspire to introduce variation in language use. Seeing that there is also selection is not difficult. Language constructs that lead robustly and steadily to successful communication and that minimizes the cognitive effort of the individuals survive and the rest wither away.

So the selectionist framework fits perfectly for studying language as a dynamically evolving complex adaptive system. Its application is not just metaphorical because the framework predicts some of the properties the language faculty of communicating agents needs to have. The mathematical tools theoretical biologists have developed for analyzing the dynamics of selectionist processes become applicable.

Given that many conflicting constraints operating on language exist (for example, less effort for the hearer may imply more effort for the speaker), and that historical contingencies obviously exist, we cannot expect that language users will ever arrive at an optimal communication system. On the contrary, overwhelming evidence from the historical development of all human languages supports the conclusion that language users move around in the space of possible solutions. Sometimes languages optimize one aspect (for example dropping a complex case system), which then forces another solution with its inconvenience (for example using a large number of verbal patterns with idiomatic prepositions —as in English). Language, as a complex adaptive system, exploits the available physiological and cognitive resources of its community of users to handle their communicative challenges, but without ever reaching a stable state.

Here is a third example how universal notions of complex systems are relevant to understanding language emergence: the application of the notion of structural coupling to the coordination of different subsystems of the language faculty and especially to the coordination of the conceptual-intentional system and the language system. Suppose that the conceptual system has the capability to create new color distinctions. For example, you get two types of bananas, a green one and a yellow one, and discover that the yellow banana tastes better than the green one. So your conceptual system cuts up the continuous color space to have a boundary between these two color regions so that you can pick out the tasty bananas in the future. The same process occurs when you encounter delicious red apples and rotten brown ones, orange and black carrots, or purple and green (unripe) plums. But if individuals keep inventing new color distinctions to serve their individual needs, based on their personal history of interactions with the world, we would have no guarantee that the color distinctions become shared. Miscommunication is inevitable.

This has lead some people to suppose that color distinctions must already be genetically fixed in advance, which is unlikely given the profound cultural and individual differences in color categorization I discussed earlier. So we have a paradox for which structural coupling can come to the rescue.[11] The conceptual system suggests categories offered to the lexical system to be put into words, so an influence from conceptualization to language during production exists. But the lexical system also influences the conceptual system. Speakers may get feedback whether their hearers properly understood the categories they used, and if not, the speaker's conceptual system might adjust them.

The hearer has to figure out how the speaker uses words. So the hearers' lexical system tells their conceptual-intention system what categories should apply to the present case. If they do not, the conceptual system must adjust categories or create new ones. Our models have shown that this mutual influence is enough to efficiently coordinate both subsystems inside the language faculty. Perhaps more remarkably, the same mechanism is enough to coordinate the conceptual inventories across agents, even though each agent independently develops his own. Again, we see a macroscopic structure, a shared repertoire of perceptually grounded categories, emerges through local autonomous interactions between agents. No central control or prior disposition is required.

4. Recruitment Theory

My third hypothesis concerns the origins of the language faculty. Showing the information processing that a language faculty must be able to do to see the emergence of symbolic communication systems through universal morphogenetic processes, such as self-organization or selectionism, is insufficient. I also need to explain how this language faculty could have evolved. We know little about this extraordinarily difficult question despite intense speculation and debate. Only a few coherent theories exist, but one framework, defended by many linguists and evolutionary psychologists, does emerge. It starts from the hypothesis that the "human language faculty is a complex adaptation which evolved by natural selection for communication"; the interconnected areas of the brain involved in language form a highly specialized neural subsystem, a kind of language organ, which is genetically determined and came into existence through neo-Darwinian genetic evolution.[12] So I apply selectionism here at the level of the biological embodiment of the language faculty and not at the level of the language itself as I advocated in the previous section.

According to this vision, language is an instinct, not learned as such. The language organ grows like the liver or the fingers of a hand. The environment sets some growth parameters, for example, the way the liver of a heavy drinker becomes "fatty" to deal with excess alcohol. But none of the kinds of intelligence, invention, and intentional learning is ascribed to the language faculty the way I suggested earlier. Nor do we see any key role seen for the rapid cultural evolution of new language traits through grammaticalization as we view language change as surface variations on a static *Bauplan* (building plan; common properties of a systematic group) that captures the essence of language.

Steven Pinker, with like-minded psychologists and linguists, advances two types of arguments in favor of this language-as-adaptation hypothesis. The first type is based on examining the structure of language. Human languages exhibit some non-trivial universal trends, which could be the logical consequence of an innate language acquisition device that imposes quite specific structures on human languages. Otherwise, seeing how extremely young children based on apparently poor data so easily and routinely acquire the intricate complexity of human language is difficult. In addition, these same universal trends show up when new languages form, as in the case of Creoles. So far, so good —but two problems arise with this argument.

The first problem is that the universal trends are statistical trends, not universal laws. That means we observe some features in many languages, but we find many other languages (often the majority), which do not have them. A speech repertoire with the vowels *a*, *e*, *i*, *o*, and *u* is found in 28 percent of the world's languages, but the other 72 percent have other systems. Some have only three vowels; others have nine. R-colored vowels (like in American pronunciations of "sir") occur only in 1 percent of the world's languages. The language of the Khoi-San in the Kalahari Desert uses clicks and produces some vowels while inhaling. Mandarin Chinese or Somali use tones. French exploits the nasal cavity (as in the vowel *un* (one)). Ian Maddieson and Peter Ladefoged, who have been the principle experts in this area and built the

UPSID database on which we usually base discussions on speech trends, point out that although there are 177 vowels and 645 consonants, this classification groups phonemes which are not always exactly identical. They argue that any sound the vocal tract can produce and the auditory system systematically recognize (after training) is used in one language or another. So if you have to design a language faculty with an innate set of speech sounds, how would you do that? Which ones would you include or exclude?

Similar issues arise with any other aspect of language including conceptualization. If you want to create a bias for color categories, what would that bias be? Although some statistical trends appear to exist, as first shown by Brent Berlin and Paul Kay in the 1960s, tremendous differences not only among languages but also among individuals speaking the same language are also evident. For example, British psychologists Debi Roberson and Jules Davidoff studied the Berinmo of Papua, New Guinea, who conceptualize the yellow-blue-green region of English into two regions named *nol* and *wor* along lines that do not coincide with those of English at all. Roberson and Davidoff replicated their study with the Himba, a semi-nomaid culture in equatorial Africa, and established yet another profoundly different color map. So if you want to bias color concepts, whose categorization are you going to chose? Why would the color categories of English have any bigger claim on being innate than those of the Himba or the Berinmo?

The situation is not much better for grammar. Adaptationists assume that minimally the syntactic and semantic categories of languages, such as the parts of speech (noun, speech) or possible cases (nominative, dative, accusative) are innately determined and provide a strong bias on language acquisition. They argue that these have to be part of the language instinct that has evolved through natural selection. But what about the observation that many languages do not make a distinction between adjectives and nouns, do not have a notion of subject or direct object, categorize events and event types into quite different semantic roles, maintain quite different subtle usages for the dative, and so on. Decades of research in generative linguistics to capture the universal grammar that underlies all human languages have not at all yielded a consensus view, despite that most studies focus only on English or a few other European languages. Some linguists have made the unavoidable conclusion that syntactic and semantic categories are language and maybe even speaker dependent. Like pieces of chess, grammatical categories only make sense within a designated system. If we pull a chess piece from a chess game and insert it into a set of checkers, what does it mean? If I wanted to design a language system, which has the “innate” syntactic and semantic categories programmed in but is not ridiculously biased towards English, which ones should I program into it?

A second problem arises. Even assuming that universal trends exist, does this necessarily mean that they are the result of innate predispositions? My earlier discussion on complex adaptive systems should alert the reader that this is not necessarily the case. That universal trends in language are a consequence of universal morphogenetic processes (like self-organization, selectionism, structural coupling) is entirely possible. These processes would operate within the constraints that human beings have: the physical constraints of the real world, human embodiment, the resources and limitations of the human brain (for example, memory, and processing speed), and the nature of the communication task. Without going to the bottom of this line of explanation, we cannot know for sure what has to be innately specified and what is emergent. Many highly complex patterns in nature are not the result of innate predispositions but a side effect of epigenetic, behavioral, and environmental factors. Is the same style of explanation relevant to language? It would be foolish and narrow-minded to exclude it.

For example, are the trends we find in the vowel systems of the world's languages a consequence of innate predispositions? Or, are they the consequence of morphogenetic processes constrained by the human vocal tract, the auditory system, the physics of the sound medium with its unavoidable noise, the limits to the precision of motor control, the nature of human categorization processes, and memory? A lot of basic biological structure develops under genetic influence, such as the structure of the larynx, the vocal chords, the glottis, and the circuitry to perform fine-grained high-speed control or auditory categorization. The point I am arguing is that the vowels do not have to be innate. The generic machinery exists, but not the contents on which it operates. Once a group starts to self-organize a speech system with this kind of machinery and subject to the same constraints as human beings, they arrive automatically at the sort of speech systems we find in human languages.

I argue that the same explanatory framework is true for all other trends observed in languages: color categories, syntactic and semantic categories, and typical grammatical patterns. I realize that we have to show precisely how that is possible and many of our experiments do that. Claiming that something is innate (for example, an inventory of speech sounds) does not resolve the difficulty. You cannot merely shift responsibility to the evolutionary biologist. You must show scenarios that demonstrate how the distinction between dative and accusative, active and passive, voiced and voiceless, or past perfect and future continuous may become innate. I have never seen a convincing evolutionary argument or a simulation for any of these.

Arguing that a quasi-optimal system is in place (for example that the vowels are optimally distributed over the space of possible sounds the vocal tract can make or that the ear can perceive) is not sufficient either. We have to show how linguistic agents can autonomously discover these solutions without central control, a global view, or prior bias. Or, if you are an adaptationist, you have to show how genetic evolution and natural selection have found them and how the genes can exert sufficiently control over the fine-grained structure of the brain so that the brain circuits become biased towards them.

The second type of arguments in favor of the language-as-adaptation hypothesis comes from molecular and population genetics. They rest on the identification of genes that have undergone selection in the human lineage, and have an unmistakably identified effect on language. The most promising (and only) example in this respect is the FOXP2 gene which is linked to a unique constellation of language impairments identified in the multi-generational KE-family.[13] At the time of its discovery, it was hailed as "the" or "the first" language gene by the popular press. Since then, a growing number of questions have been raised which show that things are not that simple, as is so often the case in biology.

First, pinning down exactly the language organ is supposed to be in the brain is difficult. Brain imaging studies overwhelmingly show that speaking or listening involves the activation of many brain areas, not just those traditionally associated with language: Broca's and Wernicke's area. Even for just speaking, the cerebellum and thalamus gets involved along with other motor areas. Conversely, the traditional "language areas" are also active in many other non-linguistic tasks; they are definitely not only specialized for language. For example, Broca's area is heavily involved in selecting and monitoring all sorts of motor control patterns, not only those related to speech production. Interestingly, different brain areas can take over the functionalities of these areas if brain damage occurs, as long as the brain still has enough vitality and plasticity to recuperate. The most amazing example in this respect is that the neural activities associated with Broca's and Wernicke's area can shift from the left hemisphere of the brain where they are typically located in right-handed people to the right hemisphere. So the picture is blurred. No obvious language organ exists akin to, say, a patently delineated liver.

Second, we have to be careful about a direct mapping between genes and brain areas.[14] Those involved in this research know this, but that message usually gets lost in popularization. The FOXP2 gene is not uniquely relevant for language but plays a crucial role in the development of many areas of the brain and leads to many impairments if affected, especially in the motor domain. The gene plays also a role in the development of the lungs, the gut, and the heart. To say that FOXP2 is a language gene is misleading, you might as well call it a gene for lungs. FOXP2 is a transcription factor, which has the potential to affect a potentially large number of genes. In the brain, its function is probably more generic, perhaps regulation of post migratory neuronal differentiation.

Does all this mean that the language-as-adaptation hypothesis is invalid? Not necessarily. It just means that we cannot elevate this hypothesis to dogma yet. Science can only benefit from exploring and comparing many hypotheses. An alternative hypothesis exists: *the recruitment theory of language origins*. It argues that the origins of the human language faculty might be explainable as a dynamic configuration of brain mechanisms, which grows and adapts, like an organism. Instead of being genetically pre-wired, the brain may be recruiting available cognitive/neural resources for optimally achieving the task of communication. Achieving this task implies maximizing expressive power and communicative success while minimizing cognitive effort in terms of processing and memory. The mechanisms that get recruited are not specific for language. They are exaptations in Stephen Jay Gould's terminology and instantiated and configured dynamically by each individual. Consequently, genetic evolution by natural selection is not the causal force that explains the origins of language. Genetic evolution still plays a role, partly to evolve the basic building computational blocks such as categorization or detection of hierarchical structure, and partly to evolve brains that have the kind of intense plasticity that recruitment requires. I personally do not believe that genetic evolution has fixed the vowels of languages or whether the subject of a sentence is likely to come before or after the verb, or any of the other constraints on language that are usually ascribed to an innate language faculty.

One example of recruitment concerns egocentric perspective transformation (computing what the world looks like from another viewpoint). This activity is normally carried out in the parietal-temporal-occipital junction of the brain and used for a wide variety of non-linguistic tasks, such as prediction of the behavior of others or navigation. All human languages have ways to change and mark perspective (as in "your left" versus "my left"), which is only possible if speaker and hearer can conceptualize the scene from the listener's perspective. That, in turn, implies that they have recruited egocentric perspective transformation as part of their language system. By recruitment, I mean that information can flow from one subsystem to another and so that the signals make sense on both ends.

Another example of a universal feature of human languages is that the speaker can express emotional states by modulating the speech signal. For example, in case of anger, the speaker may increase rhythm and volume, use a higher pitch, or a more agitated intonation pattern. This requires that the neural subsystems involved in emotion (such as the amygdala) are linked into the language system so that information on emotional states can influence speech production and so that information from speech recognition can flow towards the brain areas involved with emotion. I would argue that the specific links are not predetermined, but are the result of a recruitment process that relies on the plasticity of the human brain to establish links in all directions and then choose the ones needed to participate in language.

Another example relates to the core of grammar. Language is unique compared to animal signaling systems because it uses recursive structure: words group in phrases, which group in bigger phrases and a bigger phrase can reoccur as component in a phrase of the same type. For example, "the big box" is a noun phrase, which can be part of the noun phrase "the ball next to the big box," which can in turn be part of a still bigger noun phrase: "the speed of the ball next to the big box."

Hierarchical structure implies that language production must involve a hierarchical planning process and that parsing requires the ability to recover hierarchical structure. But is this capability unique for language? Definitely not. Whenever we plan a series of actions, we have to be able to conceive and monitor hierarchical structure. Whenever we are recognizing events in the world, and especially complex actions performed by other human beings, we need hierarchical event recognition. So we find it difficult to maintain that hierarchical processing is unique. It may not even be unique to the human species. The language faculty recruits the ability to do this kind of hierarchical processing and then instantiates it for the domain of syntax.

I argue that the brain functionalities necessary to create a symbolic communication system could be entirely epigenetic, driven by the need to build a better, more expressive, and more efficient communication system as opposed to being genetically pre-determined and evolved through natural selection. Because such profound differences between languages exist, we cannot exclude that different languages call upon different functionalities. If we compare the needed language faculty, we may end up with quite different neural systems. We may even see differences between individuals speaking the same language, as alternative ways to achieve the same functionality may exist. The historical choices a language community may have made force newcomers to recruit designated neural resources and subsystems, and so in some sense, a human community encourages the brains of newcomers to configure themselves in designated ways. If a language exploits the pitch differences between vowels, such as the tone system in Chinese, then everyone speaking that language must develop feature detectors able to sense and reproduce this, something not at all simple to do if you are not accustomed to it. If a language does not exploit word order for syntax, as early forms of Latin or present-day Australian aboriginal languages, then the language system has to set up the parsing process quite differently, keeping in memory disconnected individual units, which might only link into hierarchical structure when all words have been heard.

The recruitment theory resonates with other proposals for the origins of the language faculty. For example, the biologist Eörs Szathmary has proposed the metaphor of a growing "language amoeba," a pattern of neural activity essential for processing linguistic information and grows in the "habitat" of a developing human brain with its characteristic connectivity pattern.[15] Many researchers, including those who believe that some parts of the language faculty are innate, agree that a multitude of non-linguistic brain functions gets recruited for language —if we take the language faculty in a broad sense, including conceptualization of what to say.

5. The Human Revolution

The reason something like human language developed might appear obvious. Almost any task requiring cooperation can profit from having a powerful communication system: to plan joint action, coordinate effort and activity is ongoing, or discuss afterwards what went wrong and how we could improve things. Then why species closely related to us do not have language as well?

Part of it has to do with lacking the mental capability for language. But if the same selectionist pressures acted on chimpanzees, whose brains appear tantalizingly similar to those of human beings, with whom they share 98.77 percent identical genomes, did not both species undergo similar evolution?

Consider two additional puzzles: Communication gives information away which is potentially exploitable by the listener. Fine for the listener. But what does the speaker gain? Speakers may give information away that is useful for them or perhaps disclose their intentions so that others know what they are going to do, so the listeners can use that information for personal gain. In a Darwinian world driven by selfish genes, only those organisms survive that maximize their individual fitness. So cooperation and communication presents a problem. We cannot take cooperation for granted.

People chatter away on the bus with neighbors they have never seen before and may never see again, apparently unaware of such a potentially dangerous open attitude towards information sharing. Then how can hearers avoid manipulation by speakers? In a symbolic communication system, cheating is easy. For example, it does not cost you anything to say with great conviction that you are powerful or angry and in that way entice others to give you food or sexual favors. I am constantly amazed by how some people apparently get sucked into sending large amounts of money or disclose their personal information solely based on some Email message promising large gains. In a Darwinian world, such an attitude would be a disaster. A cheater could clobber up all resources to the detriment of the rest, leaving the others to be eaten alive.

So how do animals cope with this danger of a fully cooperative attitude? They cope in two ways. First, they only unselfishly cooperate and communicate freely under designated conditions, such as kinship selection —I help you because you share enough of my genes that I consider unselfish cooperation and communication to be beneficial in the propagation of my genes.

More directly relevant here is that if animals communicate, they invariably use signal-based communication instead of symbol-based communication. Signals are bodily changes or auditory and bodily gestures that have an influence on the behavior of others. The peacock tail or the alarm calls of vervet monkeys are two typical examples. Signals are analog in the sense that the intensity of the signal correlates with the intensity of what is expressed; the bigger the peacock, the stronger the bird, the louder the alarm call, the greater the threat.

Signals are not under conscious control, neither for the sender nor for the receiver; their evolution is nearly always genetically driven. The peacock cannot just grow a longer tail. The dog cannot refrain from barking when it sees another dog threateningly near. Emotional signaling in human beings (getting white from anger, red from shyness, crying from sadness) falls within this class. Unless you are an excellent actor, you cannot just turn white when you want to without feeling authentic anger; unless you have a heart of stone, you cannot remain untouched when a person cries in your arms. Signals are honest. If a peacock could grow a longer tail by will, it could signal, “Look how powerful I am,” without actually being sufficiently powerful. If you could just turn white when you want to, you could cheat other people into believing that you are angry when you are not.

In contrast, symbolic communication, like human language, does not link in a reflex-like way to internal states, but is produced and interpreted on demand, through a rich system of conceptualizations. This makes possible lying, doubt, or ignoring what someone else says, whereas these actions are impossible to do with signals. We establish the meaning of symbols through a cultural consensus. Whether we call fish “fish,” *poison* (French), or *oryu* (Chinese) is entirely arbitrary.

Symbols are not costly; the cost of producing a signal does not keep it honest. Whether you say “I did not cheat at the exam” or you say, “I cheated at the exam” requires equal amounts of efforts even if statement is true and the other is false. It can be adapted easily to new communicative needs; it can handle great semantic complexity; we can learn it quickly. These features make symbolic communication powerful. But the power comes with enormous risks that animals generally cannot take.

All this helps us to understand why animals do not have symbol-based communication, but not yet why human beings do have it. Here I propose my fourth hypothesis: *Human beings developed symbol-based communication because they underwent a kind of revolution that allowed them to surpass the Darwinian world of selfish genes. They became social, in the sense of fundamentally cooperative and capable to establish and follow a rule of law, even if it goes against their interest*. A game of chess is an example of the rule of law. Players agree to follow the rules; everything is allowed, as long as you follow the rules. You may not hit somebody else on the head with a chess piece or use telepathic forces.

All this helps us to understand why animals do not have symbol-based communication, but not yet why human beings do have it. Here I propose my fourth hypothesis: *Human beings developed symbol-based communication because they underwent a kind of revolution that allowed them to surpass the Darwinian world of selfish genes. They became social. They became fundamentally cooperative and capable to establish and follow a rule of law, even if that law goes against their interest*. A game of chess is an example of the rule of law. Players agree to follow the rules; everything is allowed, as long as you follow the rules. You may not hit somebody else on the head with a chess piece or use telepathic forces.

Language is like the rule of law. You agree to abide by rules of the language game that you tacitly accept. You accept the conventions and contribute to them yourself. In jointly creating a language, you have to be entirely honest (in your use of the language) and cooperative, otherwise symbol-based communication system collapses.

The sociality assumption is crucial for understanding why we profusely speak not only with our kin, but with anyone. We hold a tacit assumption that others will not exploit us, that we can give away information, and that eventually everybody will benefit from the free flow of information. Those that display a greater capability to acquire new information gain prestige in the group when they make the information available to others. This does not mean that no one will exploit other people. Unfortunately, this happens all too often and language plays a key role in manipulation and cheating. But if human language would be an honest, costly signaling system, it would not be possible to use it for this purpose. Granted, human beings do not always follow the rule of law. We are constantly tempted, especially those of us who have amassed more power or have fewer moral objections compared with others, to try to set up or change the rules to our advantage or to circumvent the rules. But if we are caught, the communal reaction is severe.

All this raises the question how sociality might have developed in our species. How have our brains become so profoundly social? Where does this unusual cooperative attitude and capability to set up the rule of law come from? How did it develop and why? I personally have no particular theory, although I believe the theories of some anthropologists such as Leslie Aiello, Chris Knight, and Camilla Power go a long way.[16] They base their answers on biological arguments. As brain size expanded, energy cost for bringing up helpless children grew. Mothers had to break through the male dominance structures that are still the norm in chimpanzee groups in order to entice males to help feed their children. This necessarily required a cooperative attitude among females to raise the children and form coalitions against males but also among males to hunt the large game that became necessary. Once language or other forms of symbolic representation became established even in a very simple form, it was an enormously powerful instrument to help organize human society, formulate and enforce the rule of law, and help everybody to see the world in similar ways.

6. A Revolution in Linguistics

I have introduced hypotheses for each of the four issues introduced in the beginning of this paper. All of them are obviously controversial and some readers may have become sufficiently upset that they have torn pages out of this book in total disagreement! The experiments that we have been doing with artificial agents attempt to operationalize each hypothesis and prove that if we endow artificial agents with the mechanisms they imply, we see the emergence of communication systems with features similar to human languages. So what is the nature of language and its origins that is beginning to emerge? Language did not jump from nothing by a freak genetic mutation that suddenly caused a human child to have a language organ, enabling it to speak a grammatical language which its parents presumably could not understand (as apparently Derek Bickerton suggested).[17] Neither did human beings stumble on language by accident and then transmitted culturally by imitation, the way human beings presumably stumbled on the idea of roasting coffee beans to make a stimulating drink or on chewing ginger to alleviate stomach pain.

Language is an invention created by intelligent human beings and propagated based on intentional learning. To engage in language, the brain self-configures an enormously complex language faculty, recruiting the many functionalities required from perceptual-motor processing and conceptualization to parsing and production. Human embodiment, the challenges of the communication task, and the ecologies and environments in which human beings find themselves all conspire to drive the emergent communication system in particular directions and explain why we see some universal trends.

We cannot underestimate the depth of change I am advocating here compared with mainstream linguistics. Noam Chomsky famously set the goals of the field in the late fifties and sixties as follows:

> Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who know its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of this language in actual performance.[18]

These goals were entirely appropriate at a time when the cognitive revolution was beginning to make everybody aware that we could best study mental processes as information processes. A good first step was to isolate a single individual and study the knowledge this individual might appear to have at any given moment under ideal circumstances. But if we keep pursuing only these goals, as many mainstream linguists continue to do, we discard the possibility of finding explanations for the origins of language.

We can no longer apply or observe the universal morphogenetic processes like selectionism or self-organization that underlie the origins of so many other complex structures in nature. First, these collective processes always rely on a population, for example a population of organisms in natural selection or a population of molecules in molecular self-organization. So we also need to take a population view in our models of language. Doing so is entirely natural because language is a phenomenon found in a community of people, not a single individual. Second, morphogenetic process always requires variation among the elements and their behavior. So we should not consider the kind of population we need in our models to be homogeneous. Again, doing this is entirely natural because studies have proven that in a human language community, neither the competence nor the performance of all individuals is homogeneous. Third, we cannot restrict our attention to a synchronic view, only looking at a snapshot of a language at a designated point in time. If we did so, we would not see that language is forever changing.

Linguistic theory should, therefore, incorporate change at its foundation, not only change over long periods of time as studied in historical linguistics, but rapid change even within the confines of a single dialogue. Every verbal interaction changes the language and through these cumulative changes, language grows and evolves.

Apart from leading to a more naturalistic approach to language, this complex adaptive systems approach brings language more in line with contemporary science. Instead of a science of being, trying to capture the essential properties of a static idealized entity, it becomes a science of becoming. To use the terminology of the French philosopher Gilles Deleuze, language is a multiplicity. It is forever in the making.

Acknowledgment

The research underlying this paper was conducted at the Sony Computer Science Laboratory in Paris and the Artificial Intelligence laboratory of the University of Brussels (VUB). I thank Oscar Villaroya for his continuous encouragement, the organization of the tremendous Social Brain symposium, and his deep thinking on the issues I discussed in this paper.

Notes

1. Alison Wray, ed., *The Transition to Language* (Oxford: Oxford University Press, 2002); E. J. Briscoe, ed., *Linguistic Evolution through Language Acquisition: Formal and Computational Models* (Cambridge, UK: Cambridge University Press, 2002); Angelo Cangelosi and Domenico Parisi, eds., *Simulating the Evolution of Language* (Berlin: Springer-Verlag, Berlin, 2001); James W. Minett and William S.-Y. Wang, *Language Acquisition, Change and Emergence: Essays in Evolutionary Linguistics* (Hong Kong: City University of Hong Kong Press, 2005); and Caroline Lyon, Chrystopher L Nehaniv, and Angelo Cangelosi, eds., *Emergence of Communication and Language* (Berlin: Springer Verlag, 2006).

2. Luc Steels, “A Self-Organizing Spatial Vocabulary,” *Artificial Life*, 2:3 (1995), pp. 319–332; Luc Steels, F. Kaplan, Angus McIntyre, and Joris Van Looveren, “Crucial Factors in the Origins of Word-Meaning,” *The Transition to Language*, ed. Alison Wray (Oxford: Oxford University Press, 2002); and Luc Steels, “Evolving Grounded Communication for Robots,” *Trends in Cognitive Science*, 7:7 (July 2003), pp. 308–312.

3. Steven Pinker and Ray Jackendoff, “The Faculty of Language: What’s Special about it? *Cognition*, 95 (2005), pp. 201–236.

4. Rosalind Thornton and Kenneth Wexler, *Principle B, VP Ellipsis, and Interpretation in Child Grammar* (Cambridge, Mass.: The MIT Press, 1999).

5. Elizabeth Closs Traugott and Bernd Heine, *Approaches to Grammaticalization*, vol. 1 and 2 (Amsterdam: John Benjamins Publishing Company, 1991).

6. Michael Tomasello, *The Cultural Origins of Human Cognition* (Cambridge, Mass.: Harvard University Press, 1999).

7. Frederick J. Newmeyer, *Language Form and Language Function* (Cambridge, Mass.: The MIT Press, 1998).

8. Luc Steels, “Language as a Complex Adaptive System, Lecture Notes in

Computer Science," *Parallel Problem Solving from Nature–PPSN VI*, ed. Marc Schoenauer, et al. (Berlin: Springer-Verlag, Berlin, 2000).

9. Bart De Boer, "Self-Organization in Vowel Systems," *Journal of Phonetics*, 28:4 (2000), pp. 441–465.

10. William Croft, *Radical Construction Grammar: Syntactic Theory in Typological Perspective* (Oxford: Oxford University Press, 2001).

11. Luc Steels and Tony Belpaeme (2005) "Coordinating Perceptually Grounded Categories through Language: A Case Study for Color," *Behavioral and Brain Sciences*, 24:6 (2005), pp. 469–489.

12. Steven Pinker, "Language as an Adaptation to the Cognitive Niche," *Language Evolution: The States of the Art*, eds. Morten H. Christiansen, and Simon Kirby (Oxford: Oxford University Press, 2004), p. 16.

13. W. Enard, M. Przeworski, S. E. Fisher, C. S. Lai, V. Wiebe, T. Kitano, A. P. Monaco, and S. Paabo, "Molecular Evolution of FOXP2, a Gene Involved in Speech and Language," *Nature*, 418 (2002), pp. 869–872.

14. Karin Stromswold, "The Heritability of Language: A Review and Metaanalysis of Twin, Adoption, and Linkage Studies," *Language*, 77:4 (2001), pp. 647–723.

15. Eörs Szathmary, "Origin of the Human Language Faculty: The Language Amoeba Hypothesis," *New Essays on the Origins of Language*, eds. Jurgan Trabant and Sean Ward, *Trends in Linguistics, Studies, and Monographs*, 133 (2001), pp. 41–55.

16. Luigia Carlucci Aiello, "Terrestriality, Bipedalism and the Origin of Language," *Evolution of Social Behaviour Patterns in Primates and Man*, eds. Walter Garrison Runciman, John Maynard-Smith, and Robin Ian MacDonald Dunbar (Oxford: Oxford University Press, 1996), pp. 269–289; Chris Knight, *Blood Relations: Menstruation and the Origins of Culture* (New Haven, Conn.: Yale University Press, 1991); and Camilla Power and Leslie C. Aiello, "Female Proto-Symbolic Strategies," *Women in Human Evolution*, ed. Lori D. Hager (New York: Routledge, 1997), pp. 153–171.

17. Derek Bickerton, "The Language Bioprogram Hypothesis," *Behavioral and Brain Sciences*, 7 (1984), pp. 173–222.

18. Noam Chomsky, *Aspects of the Theory of Syntax* (Cambridge, Mass.: MIT Press, 1965), p. 3.