

# The origin of the human language faculty: the language amoeba hypothesis

*Eörs Szathmáry*

## 1. Introduction

The origin of the eye, once considered one of the hardest problems for evolutionary science, now looks almost trivial compared with the problem of the origin of human language. What makes the question of eye origin now seem to be so easy? There are at least four reasons. First, the anatomy and physiology of the eye are relatively well understood. Second, there are many eye forms that can serve as plausible evolutionarily intermediate stages (even the vertebrate eye has a close analogue in cephalopods). Third, although apparently related to an ancestral light-sensitive spot, eyes with various anatomical and physiological structures evolved about forty times. Fourth, the genetics of eye development is now better understood.

Compare these with the corresponding list for language origin. First, our knowledge of the “language organ” is fairly limited. Second, there seems to be no extant intermediate between “protolanguage” (broadly understood as limited vocabulary without syntax) and human language. Third, even the much simpler protolanguage evolved only a few times (candidate species where protolanguage could be used in the wild include bottle-nosed dolphins, grey parrots, and chimps), and natural language evolved only once. Fourth, the genetics of human language is largely unknown.

It may seem, then, almost impossible to say anything useful about the origin of natural language. Nevertheless, there are at least five sources of scientific input that can, I think, contribute to enhancing our understanding of this conundrum:

- The spread of elements of vocabulary and language rules can be analysed using the tools of population biology (Nowak and Krakauer 1999, Nowak et al. 2000).
- A better understanding of primate communication and social life greatly contributes to narrowing the evolutionary and intellectual gap between our ancestors and us.
- The anatomy, physiology, and genetics of specific language impairment (SLI) expand our knowledge of language’s biological foundations.

- New, noninvasive brain analysis methods (such as PET) reveal a surprisingly dynamic neural foundation for human language.
- We are in the process of formulating a selectionist picture, connected to brain epigenesis, of what happens inside the brain during at least some forms of learning.

In this essay I will elaborate on the last four points. I will argue that our increased understanding in these areas points toward a “language amoeba” in the human brain. The language amoeba is the neural activity pattern that essentially contributes to processing of linguistic information, especially syntax. It is a sort of dynamic manifestation of Chomsky’s language organ. It finds its “habitat” in the developing human brain, whereas the brains of other primates apparently cannot sustain it. Why? I attempt to outline a tentative answer. Since the gap between other primates and us is fairly narrow, the genetic and correlated functional changes could not have been too numerous. What were the critical changes? I argue that an appropriate and rather widespread connectivity pattern in the immature human brain renders it a habitat for the emerging language amoeba. This change does not require too many altered (probably regulatory) genes, but there are great risks involved, which make this “major transition” truly difficult.

---

## **2. Primate communication and cognition: the narrowing gap**

Primates are especially relevant, since they are our closest evolutionary relatives. Bickerton (1992) correctly emphasises that the right starting point for understanding the origin of natural language is something like proto-language. Our closest relatives seem to possess it in some form. They transmit signals that may well turn out to be Saussurean signs: there is an orderly and reversible relation in the brain that links object to concept to sign. Sensitivity to conspecific vocalisations, but not to other types of auditory input, seems to be lateralised (Ghazanfar and Hauser 1999). Apes can be taught some elements of human language, but their syntactic abilities are clearly limited. This is even true for “action grammar”: the hierarchical structure of embedded actions. Chimps seem to be severely limited with recursive embedding (a crucial element of natural syntax): they usually cannot go beyond three levels. Remarkably, “linguistically” trained chimps do somewhat better (Greenfield 1991). The case of Kanzi demonstrates that childhood learning is better than adult learning and that understanding

---

is easier than production. But this should not surprise a human too much. In the social context, apes seem to carry all the necessary Machiavellian weapons that we do.

In an insightful paper, Worden (1995) deduced a “speed limit” for evolution that seems surprisingly uniform despite the fact that the relative contributions of component processes (mutation, recombination, selection, and drift) are allowed to vary. He argues that the number of genes specifically associated with the “language organ” must be rather limited, which prompts one to think that it was not the number of genes, but their identity (i.e., mainly regulatory effects) that was peculiar during the evolution of natural language.

All this leads me to two conclusions. First, apes can do everything that we can do – expect for things that require language or neural processing abilities that are likely to have evolved through natural selection for language. Second, there are likely to be surprisingly few genes associated with language’s neural substrate that are critically different from ape genes/alleles.

### **3. Specific language impairment (SLI): phenomenon and genetics**

There is no disagreement that SLI is real. What is contested is how closely it is limited to, or rooted in, a specific grammatical impairment. The famous Gopnik (1999) case has been very stimulating because of its characterisation as “feature-blind” dysphasia and its obvious genetic background (a single dominant allele). Whether other cognitive skills are also, or even primarily, affected has been debated ever since (Vargha-Kadem et al. 1998). Recently, more evidence involving other linguistic groups has accumulated (Dalalakis 1999, Rose and Royle 1999, Tomblin and Pandich 1999). A recent paper (Van der Lely et al. 1998), sadly without genetics, claims to demonstrate that grammatically limited SLI exists in “children”, notwithstanding the fact that it only studies one child.

Clearly, more genetics is needed. But the task is formidable owing to the obvious restrictions. I expect some crucial input to come from the genome projects. This input includes a list of genes and alleles that act as liability genes (Flint 1999) for language and that can cause SLI when mutated. Moreover, extension to a primate genome project could reveal genes and alleles that are closely and specifically associated with human language.

It is worth remembering that the genetics of human cognitive tasks is a notoriously difficult problem. The clinical characterisations are usually not

sufficient as descriptions of phenotypes (Flint 1999). Nevertheless, a consensus seems to be emerging that the genes involved are so-called “liability genes” that, when present in the right allelic form, significantly enhance the probability of developing the respective cognitive skills. We can safely conclude the following:

- Specific language impairment does exist, and some of its forms seem to be familial. There is disagreement as to whether it is restricted to grammar. The likely explanation behind these contrasting views is that the genetic factors in question distort smaller or larger parts of the brain. When language becomes localised, the morpho-physiological distortions involved partially overlap with the habitat of the language amoeba. It is easy to see that the relative size of the overlap with the outlying regions and the nature of the affected regions of the habitat will influence the degree to which the impairment is specific to language or, more specifically, to grammar (see Figure 1).
- Since many of the genes are likely to influence the epigenesis of a certain brain area, incomplete penetration can cause more or less severe forms of impairment.

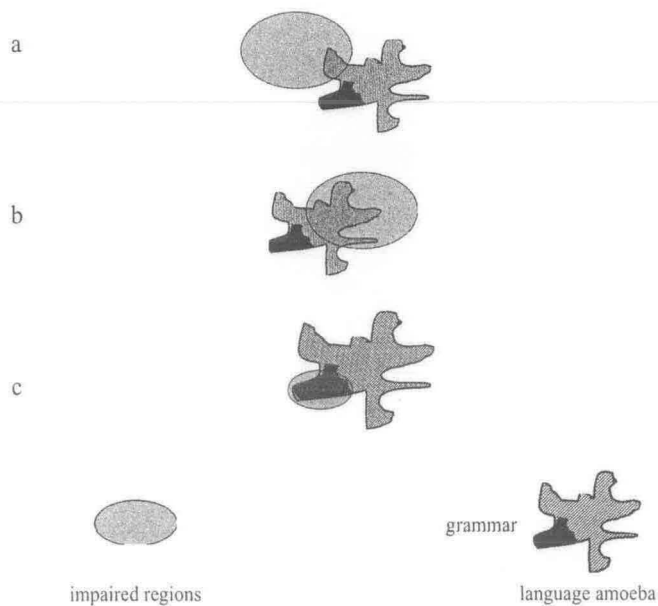


Figure 1.

#### 4. Neural dynamics and localisation of language: noninvasive studies

The analysis of neural activity during the performance of cognitive tasks has become a growth industry. The equipment's sensitivity has increased over the years. These methods are increasingly used to record brain activity during linguistic performance.

It is now widely accepted that neural localisation of language can be plastic (Nobre and Plunkett 1997, Neville and Bavalier 1998, Musso et al. 1999). Studies of brain injury have revealed that damage to the left hemisphere before a critical period is not necessarily debilitating because the right hemisphere can take over the necessary functions (Müller et al. 1999). This does not contradict the finding that in most people Broca's area does seem specialised for syntax (Embick et al. 2000). It thus seems that the common left-hemisphere localisation of language is just the most likely outcome when there is no genetic or epigenetic disturbance. PET studies have revealed a truly shocking feature of language development: the localisation of linguistic processing shifts during normal ontogenesis. The outcome in "normal" people is also highly variable.

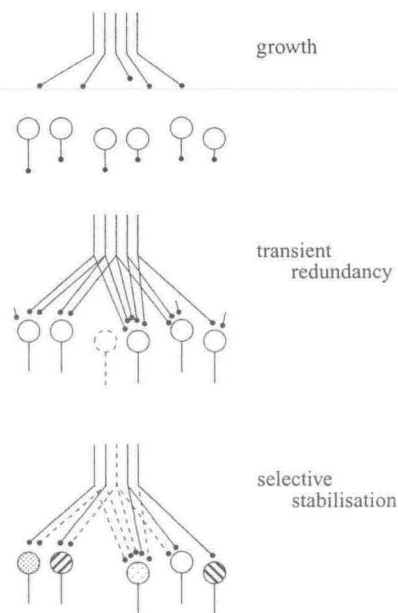


Figure 2.

Analysis of a particularly interesting genetic syndrome called Williams syndrome also reveals surprisingly dynamic manifestations during ontogenesis. Whereas affected children seem to be bad at language and good with numbers, adults perform the other way round (Paterson et al. 1999). Needless to say, the classic characterisation of the disease has been based on adult performance.

I draw the following conclusions from brain studies. First, language localisation is not fully genetically determined: even major injuries can be tolerated before a critical period. Second, language localisation to certain brain areas is a highly plastic process, both in its development and its end result. Third, a surprisingly large part of the brain can sustain language. True, there are recognised areas that seem to be most commonly associated with language, but these are by no means exclusive, either at the individual or the population level, or during either normal or impaired ontogenesis. Fourth, whereas a large part of the human brain can sustain language, no such region exists in apes.

## **5. Brain epigenesis: plasticity and selection**

It's time for a confession: on the whole, we don't understand how the brain works. Nevertheless, some crucial elements seem to be apparent. One is that the normal brain's development is enormously plastic, even though the power of genetic factors is obvious. One classic example is that in the same brain area of identical twins the two hemispheres of the same individual resemble each other more closely than the twin's corresponding hemispheres (Changeux 1983).

Another insight is that a tremendous amount of variation and selection is going on during brain ontogenesis. This is doubtless a Darwinian process. As the psychologist William James recognised a long time ago, natural selection of heritable variation is the only known force that can lead to adaptations, which has led to its application to both brain ontogenesis and problem-solving. There are several studies that all regard the brain, in one way or another, as a "Darwin machine" (see Calvin and Bickerton 2000). I think Changeux's formulation (1983) is the most relevant to the language problem. According to this view, the functional micro-anatomy of the adult cortex results from the vast initial surplus stock of synapses and their selective elimination according to functional criteria (see Figure 2).

The previous section indicated that a large part of the human brain can process linguistic information and perform syntactical operations. This means that there is no fixed macro-anatomical structure that is exclusively dedicated to language, but that the micro-anatomical structure must be appropriate, otherwise it could not sustain language. There is consequently some statistical connectivity feature of a large part of the human brain that suits it to linguistic processing. From the selectionist perspective (see Figure 2) there are three options: the initial variation in synaptic connectivity is novel; the means of selecting functional criteria are novel; or both. I think it is likely that both component processes are different in the relevant human brain areas. I would not dare to speculate on their relative importance.

It is necessary to see this idea in close connection with Rapoport's observations (1999) on the bottom-up coevolution of brain and cognition. According to the bottom-up view, a genetic change of some neural structure is subjected to selection, and based on its performance it either spreads or does not. There is, however, a so-called top-down mechanism that could have contributed more significantly to the evolution of human cognitive skills, including – especially, I argue – language (see Figure 3). The top-down view has three crucial components. First, due to the plasticity of brain development, enhanced demands on a certain brain region lead to less synaptic pruning (a known mechanism). Second, less synaptic pruning leads to more elaborate (and more adaptive) performance. Third, natural selection will favour any genetic change contributing to the growth of that brain area.

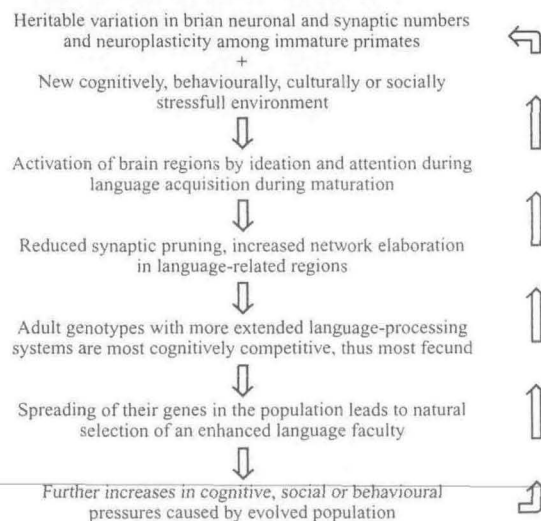


Figure 3.

Two additional observations are necessary. First (as Rapoport himself conceded), the top-down mechanism is a more detailed exposition of the late Alan Wilson's theory. According to Wilson, a larger brain, due to its more complex performance, alters the selective environment (in social animals it is composed largely of conspecifics), which selects for an even larger brain, and so on. Second, and perhaps more importantly, this mechanism is also a neat example of a Baldwin effect (or genetic assimilation), in which learning guides evolution. As Terry Deacon (1997) points out, applying the notion of genetic assimilation to language is trickier than is usually thought. The reason for this is that the performed behaviour must be sufficiently long lasting and uniform in the population. It is thus hard to imagine how, say, specific grammatical rules could have been genetically assimilated. Here, however, we are dealing with a different phenomenon: the genetic assimilation of a general processing mechanism that is performed via the connectivity of the underlying neural structures. My claim is that the network's ability to process syntactical information was the most important – and largely novel – faculty selected for. The specific hypothesis is that linguistically competent areas of the human brain have a statistical connectivity pattern that renders them especially suitable for syntactical operations. This leads me to three conclusions. First, the origin of human language required genetic changes in the mechanism of epigenesis in large parts of the brain. Second, this change affected statistical connectivity patterns of the neural networks involved. Third, due to the selectionist plasticity of brain epigenesis, coevolution of language and the brain resulted in the genetic assimilation of syntactical processing ability as such. But if this is true, then why isn't language more common?

## **6. The origin of language: a difficult transition?**

Some major transitions in evolution (such as the origin of multicellular organisms or that of social animals) occurred a number of times, whereas others (the origin of the genetic code or of language) seem to have been unique events. One must be cautious with the word "unique", however. Because we lack the "true" phylogeny of all extinct and extant organisms, we can give it only an operational definition. If all the extant and fossil species that possess traits due to a particular transition share a last common ancestor after that transition, then the transition is said to be unique. It is obviously possible that there have been independent "trials", but we do not have comparative or fossil evidence for them. Two factors can lead to "truly" unique transition.



- 1) The transition is variation-limited. This means that the set of requisite genetic alterations has a very low probability. “Constraints” operate here in a broad sense.
- 2) The transition is selection-limited. This means that there is something special in the selective environment that favours the fixation of variants that are otherwise not particularly rare. Both abiotic and biotic factors can contribute to this limitation.

There are interesting sub-cases for both types of limitation. For 1), one can always enquire about the time-scale. “Not enough time” means that given a short evolutionary time horizon the requisite variations have a very low probability, though this could change with a widened horizon. An interesting sub-case of 2) is “pre-emption”, meaning that the traits resulting from the transitions carry out selective overkill and sweep through the biota so quickly that further evolutionary trials are competitively suppressed. The genetic code could be a case in point.

It is hard to assess at the moment why language is unique. Even the not-enough-time case could apply, which would be amusing. But pre-emption, due to the subsequent cultural evolution that language has triggered, may render further trials very difficult. There is, however, another indication that language could be variation-limited in a deeper sense.

The habitat of the language amoeba is a large, appropriately connected neural network. Most of the information processing within the network elaborates on information coming from other parts of the network. There is a special type of processing required, namely that of hierarchically embedded syntactic structures. I see the following difficulties:

- Neural networks contain a large number of cycles. Syntactic structures of language are tree-like. It seems difficult to process large trees without getting into loops.
- Overproduction of initial synapses or decreased pruning, both implied in the origin of language, may easily lead to “solipsist” network dynamics. This has two consequences. First, the network’s activity is detached more than optimally from external sources of information. Second, exaggerated internal processing leads to too much “internal talking”: linguistic processing for its own sake.

I think schizophrenia is a case in point. It may be, as some have contended, the “price for language” (Crow 2000, Maddox 1997). True, this syndrome affects a surprisingly large portion – up to 5 percent – of the human population. But this is exactly the type of pattern one would expect from an evolutionarily recent major transition. Insufficient adaptive fine-tuning causes the device to malfunction fairly frequently. Brain lateralisation is partly a means to decrease the spontaneous fission of the language amoeba, which could otherwise happen more easily. After all, a large habitat can sustain two amoebas. Moreover, there is evidence that schizophrenic people display reduced lateralisation (Gold and Weinberger 1995, Shapleske et al. 1999, Zaidel 1999, Niznikiewicz et al. 2000).

The idea that a schizophrenic language amoeba is split into two (or more) functional offspring can be tested by, for example, PET and related methods. In practice, however, this might be hampered by the fact that the two entities’ spatial separation may not be apparent enough to be observable.

## 7. Modelling

One of the ways to test my ideas would be to model neural network architectures that are particularly suited to syntactic processing. Obviously, it cannot be a clumsy application of some current connectionist model (see Elman et al. 1996). There are three reasons that discourage me from applying such models. First, they tend to use some variant of the back-propagation algorithm, which is known to be grossly unrealistic (Crick 1998). Second, they cannot figure out abstract rules in the way people seem to do readily (Marcus 1998, Marcus et al. 1999). Third, they do not incorporate the implementation of neuronal selection.

My suggestion for a more fruitful modelling of linguistic operations by neural network is straightforward: implement the selectionist picture (Figure 2) as an algorithm. The plasticity involved could more easily lead to the appropriate architecture.

My recommendation contains an additional twist. It comes from the recognition that even the rules of epigenetic plasticity are themselves evolutionarily plastic; that is, they can be moulded by variation and selection. This evolutionary plasticity is probably also necessary for a successful research project. Very recently, Rolls and Stringer (2000) presented such a novel method for the design of neuronal networks. They specify a number of gene types influencing the development of neuronal networks. A small number

of genes specify the generic properties (e.g., the number of neurons in the class and the firing threshold) of a certain neuron class. A larger number of genes influences synapse formation from neurons in other classes onto neurons in the given class. Genetically determined traits include whether the connection is excitatory or inhibitory, the nature of the learning rule at the synapse, and the initial synaptic strengths. A genetic algorithm was used to evolve networks performing a particular task, including competitive learning, pattern association, and auto-association. Such an approach would also be welcome in language studies.

### **8. Where to start?**

It is not obvious where to start evolving a linguistic neuronal network. This problem relates to the question of whether the connectivity pattern in the language amoeba's habitat is unprecedented among primates. Probably not. First, because it is a statistically determined property, a few such areas must appear by chance. Second, my guess is that apes use similar areas for social cognition, but that limited cortical size does not provide a sufficient habitat for the language amoeba. The idea that social cognition (including the mental simulation of actions in a partially hierarchical network) is a good prerequisite for language is not new (Calvin and Bickerton 2000). Deacon's (1997) observation that an increase in the size of the prefrontal cortex could have critically liberated parts of the brain from visceral tasks is another case in point.