

The puzzle of language evolution

Luc Steels

VUB AI Laboratory - Brussels

SONY Computer Science Laboratory - Paris

steels@arti.vub.ac.be

Abstract

Linguistics must again concentrate on the evolutionary nature of language, so that language models are more realistic with respect to human natural languages and have a greater explanatory force. Multi-agent systems are proposed as a possible route to develop such evolutionary models and an example is given of a concrete experiment in the origins and evolution of word-meaning based on a multi-agent approach.

Official reference: Steels, L (1999) The puzzle of evolution. In: Kognitionswissenschaft, Vol 8 No 4.

1 Introduction

No one can doubt that language evolves. This evolution takes place at all levels: the sound structure (phonetics and phonology), the lexicon, the grammar (morphology and syntax), and the pragmatics. Nevertheless, the historical evolution of language has not been at the forefront of research in linguistics in the 20th century - as it was in the 19th century. Instead, the emphasis has mostly been on the description of the syntactic and phonological structures of a language at the present time. This paper argues that we should put the historical evolution of language as well as the still mysterious question of the origins of language again on top of the linguistic research agenda.

Language appears to be a complex adaptive system that is constantly constructed and reconstructed by its users. Taking a synchronic view and focusing on

the competence of idealised speakers completely misses this important characteristic and therefore a linguistic theory only providing a synchronic static description is vastly incomplete. It would be like taking away the evolutionary dimension out of biology, and consequently removing its explanatory force. Any attempt to define the current state of a language (even idealised) is perhaps even bound to fail because it is not possible to define an object that is constantly evolving.

There is also a practical consequence. Ignoring the obvious fact that language is a complex adaptive evolving system has given the false impression that one can build natural language processing applications by defining "the" grammar of a language and programming it in a computer system. The weakness of this technology has become very apparent. Natural language systems are brittle and need heavy maintenance to keep up with the evolution of the language. A more robust technology can only come from understanding how language users construct and reconstruct their language as they adapt to the language spoken in their environment and try to keep up with the ever changing communicative challenges arising in their community.

Taking the evolutionary dimension of language seriously raises some very deep issues for what a theory of natural language should be like. The evolution of language must be an obvious feature, falling out from the theory in a natural way, rather than something puzzling which needs to be explained after the fact. My strategy in this paper is to draw on biology, because it has a long history of evolutionary models. I begin by invoking similarities between language and species, which leads to an analysis of the distinction between typological and population thinking in biology and its application in linguistics. Through a concrete example, I then illustrate a new kind of language modeling and show some results from experiments with robotic agents.

2 Language and Species

Can a language be defined?

Linguists have been trying to pin down what kind of object a language is, but this has turned out to be far from obvious. Clearly we cannot define a language as that what is spoken by a particular group of language users, partly because often people speak different languages and partly because that would yield a circular

definition. To identify the group we first need to identify what it means to speak the same language, which brings us back to the initial question.

We cannot define a language by listing its utterances either. It is generally agreed that a language has an infinite set of possible utterances, many not yet realised. Linguists have therefore tried to define a language using an algorithm (a generative grammar) that generates all the utterances of a language. But this assumes that the rules of the grammar are stable. In the case of natural languages, they are not. Grammars evolve in unpredictable ways. Even if a set of rules can be written down that is valid at some point in time, these rules will become invalid as soon as the language starts to change. Moreover, observation of real verbal behavior shows that language users all the time deviate from what could be considered ideal performance, expressible as a set of rules. Linguists have therefore tended to focus on the grammar of a single (idealised) person engaging in perfect verbal behavior, but this ignores that different people have different grammars due to different histories of interaction with other members of the language community or to different social or geographical affiliations and that performance deviations may play a causal role in shaping a language.

Similar issues have arisen for defining species and it is not surprising that there is a vast literature in biology, similar to that in linguistics, discussing what the ontological status of a species is. In theory, a species consists of all organisms which have similar characteristics, specifically which can interbreed. However, it has turned out to be very hard to identify such defining characteristics. Even the seemingly absolute criterion of interbreeding is not always valid. For example, the *Platanus occidentalis* and *Platanus orientalis* are two species of trees found on different continents and with very different structures of leaves and inflorescences. Nevertheless they interbreed yielding a mixture of both and have fertile offspring (Genermont, 1998).

Species evolve, like languages, so it is not possible to define once and for all what the characteristics of a species are either. No-one can predict what elephants will look like in a million years, or even if there will be elephants. The members of a species exhibit important natural variation, even though certain characteristics are much more typical than others. So a species is rather a moving cloud of possible members with some more typical than others, as opposed to a cleanly definable enumerable set.

Circular causality

Because it is so difficult to pin down a language, it has been suggested that it is a mere epiphenomenon (Chomsky,1981). But this is clearly not the case. True, language is an emergent phenomenon, the result of the behavior of thousands or millions of language users, but this emergent phenomena has a strong causal impact on the behavior and learning of each individual. So there is a causal circularity: The individual language behaviors determine "the" language and the language co-determines the behavior of individuals. This top-down influence is established in two ways: (1) The language already in existence in the group is a strong constraint on the behavior of an individual. He or she has to abide by the system at the risk of not being understood, even if there is always the flexibility to expand the existing system if the need arises. (2) The language learner will be exposed mostly to examples from the existing language so that he or she will preferentially acquire its structures, even if they are in constant flux.

Circular causalities are not unusual in living systems. For example, the path formed by an ant society is an emergent phenomenon of the actions of the individual ants. There is no global coordination nor supervision and the individual ants cannot oversee the total path. Nevertheless the path is more than an epiphenomenon. It plays a causal role in the behavior of the individual ants. The path is formed by pheromone deposited by the ants as they follow the trail already existing. The more ants deposit pheromone the stronger the path becomes and the more the path causally impacts the behavior of the individual ants. Without the path the ants would move in all directions at random.

A similar circular causality holds for species. Because members of a species must be able to interbreed, their genes must be sufficiently compatible. Because they live in the same competitive ecosystem, they must be able to cooperate with each other and share cultural conventions. Even if we have a hard time to pin down unequivocally the boundaries of a species, it nevertheless acts as a strong and real constraint on what the future members of a species can be and how the present members can behave.

3 Typological versus population thinking

In the early 19th century, principally due to the work of Lamarck, it was well accepted that species evolution occurred, but explanations why were not satisfactory. As is well known, Darwin put forward natural selection as a possible explanation and this has since been largely confirmed by observation and laboratory exper-

iment. In order to formulate such an explanation, Darwin proposed a profound paradigm shift from a typological view of species to a population view (Mayr, 1975). I will argue that such a shift of viewpoint is also required in linguistics.

The typological or essentialist viewpoint, exemplified by the work of Linnaeus, considers all the members of a species as belonging to a certain (ideal) type. This viewpoint is directly in line with a Platonic or Cartesian philosophy which assumes the existence of ideas (perfect forms) underlying the perceived variety in real objects. The types are distinct and there is no gradation between them. Individuals are imperfect reflections of their types. The trouble with such a scheme is that evolution becomes unexplainable. In nature, we observe a gradation between types and new species evolve stepwise from existing ones. If natural selection operates on the type, all instances of a type would be wiped out at the same time, which is clearly not happening.

Darwin introduced instead a population viewpoint in biology, which considers every organism as unique and undergoing change during its lifetime. Rather than viewing a species as a type, it is seen as a relatively loose collection with fuzzy boundaries that can only be circumscribed statistically. Natural selection does not work on types but on individuals. Because every individual is unique, variation is normal. Collective properties of a species shift because individuals which carry certain variants of a trait survive in larger numbers in the total population and so their genes or cultural behaviors proliferate.

The differences between the two views are very profound and have motivated many debates in biology. As Mayr (1975:27) points out:

The ultimate conclusions of the population thinker and of the typologist are precisely the opposite. For the typologist, the type (*eidos*) is real and the variation an illusion, while for the populationist the type (average) is an abstraction and only the variation is real. No two ways of looking at nature could be more different.

Given the strong tendency towards a Cartesian philosophical stance in 20th century linguistics, it is not surprising that typological thinking has been pervasive in linguistic theory. This can be seen from the emphasis on competence (idealised behavior) as opposed to performance and the focus on an ideal speaker-hearer abstracting away from the natural variation among speakers clearly observed in a language community. But most important of all, it has lead to the conviction that there is a universal grammar from which real grammars diverge only in superficial

details. This universal grammar is assumed to be innate, the same way Plato thought that idealised forms are innate.

This typological, essentialist thinking in linguistics generates the same problems as in biology, specifically it makes it difficult to understand and explain evolution. If the grammars of all languages satisfy the schemata of a (non-trivial) universal grammar and if this is genetically encoded in the genes of an individual, then there is no reason why this grammar would evolve, nor why there is so much variation between languages or individuals. Of course, innate grammars could still evolve through genetic evolution, but genetic evolution is glacially slow, whereas linguistic evolution is sometimes very rapid. In one or two generations, a language can lose its case system for example, or shift from Subject-Object-Verb to Subject-Verb-Object word order, or develop a new category such as prepositions.

It has been argued that universal grammar only covers the core and that there is a periphery which can change (Lightfoot, 1991). Changes to the periphery may then become incorporated within the core by a process similar to genetic assimilation. However genetic assimilation is still too slow to explain how new grammatical traits spread so rapidly and no clear delineation criteria have so far been given to decide which traits are peripheral or core.

If we adopt instead a population viewpoint, language evolution becomes obvious and can be understood using the framework of natural selection. Every individual speaks differently and has a different state of knowledge of the language which is constantly evolving. "The language" is a statistical abstraction, even if it causally influences the behavior of the group as discussed earlier (through circular causality). Just as in biological species, there is a natural variation in verbal behavior and the internal knowledge of individuals. This variation may get amplified in a selectionist process. For example, increased regularities in word forms makes it easier to remember them and makes it easier to recognise word boundaries. This will cause these wordforms to propagate in the population.

Although Darwin's explanation of evolution in terms of natural selection of spontaneous variation applies, there is a big difference. Language evolution takes undoubtedly place at a cultural level rather than a genetic level. The variation comes from performance deviations and different learning histories. The selectionist pressures are related to achieving robust communicative success despite the many sources of stochasticity making real world communication so difficult. Additional pressures come from keeping the language learnability. Conventions which cannot be learned will not survive.

The emergent, evolutionary view of language is of course not new. Already

in the 19th century there was a strong current - around the same time as Darwin worked on the *Origins of Species* - viewing language as a living system in constant evolution. Contemporary work on grammaticalisation has generated similar points of view (Hopper, 1987), (Heine, et al., 1991). The main contribution of our own research is to attempt to construct formal models with the same exactness as typologically inspired, generative grammars, and to study language evolution with the same rigor as biology has studied species evolution.

4 Multi-agent systems and language games

Adopting a population view in linguistics requires profound changes to the way we build models of language. It no longer suffices to circumscribe the competence of an idealised speaker. We need at least a population and we need to implicate performance as part of the model. In the last decade a potential candidate framework for constructing this kind of models has come forward, namely multi-agent modelling.

Components of a multi-agent system

The most basic component of a multi-agent model is the agent. An agent is defined to be an entity that has a particular internal state which determines its behavior in interaction with the environment. For our purposes, the internal state contains the grammar and any other knowledge relevant for engaging in verbal behavior. There are other scientific domains where multi-agent modelling has currently gained prominence, particularly economics and biology. The internal states of the agents then contain other things, such as survival strategies.

When the state of the agent can change based on learning mechanisms inside the agent, the agent is called autonomous (it makes its own (*auto*) laws (*nomos*)). When the agents do not have access nor any influence on each other's internal states, they are called distributed. When the agents have a body and sensors and actuators to interact with the environment they are called grounded. Clearly language users must be modeled as grounded distributed autonomous agents. Although distributed autonomous agent models can be tested using software simulations only, testing models for grounded agents requires that we introduce robot technology, such as a camera for visual sensing, arms and hands for tactile sensing, etc. and that we confront the agents with the real physical world.

The second component of the model is a population of agents. This population can have a particular structure, for example agents can be spatially distributed influencing the amount of contact they have with each other. The population can also have an in- and outflux of members so that we can study transmission of knowledge from one generation to the next, either in a genetic way or in a cultural fashion.

The interaction between two agents is usually called a game and game theory has become an area of intense research in all fields where a multi-agent approach is fruitful (Maynard Smith, 1982). A game has an outcome which in the case of autonomous agents may change the future behavior of the agent. Because language behavior takes place at many levels, we can envision many different kinds of games. The most complete one would be a fully grounded language game, involving perception, conceptualisation, lexicon lookup, grammatical parsing and production, speech production and perception, interpretation, and action in the world. But much simpler games can be defined to focus on subparts of language. For example, phonetics and phonology can be studied using imitation games, in which one agent tries to imitate the sounds produced by another agent. Lexicon formation can be studied using naming games, in which agents try to transmit meaning through individual words. Syntax can be studied by agents sending each others strings to be parsed and generated until all agents share a common grammar.

A specific investigation consists of defining the architecture of the agents, creating a population, which can possibly change its constellation if one wants to study language transmission, and then letting the agents play a successive series of language games. In such investigations, it becomes quite natural to study language evolution. For example, one can test whether agents with a particular architecture enabling them to construct and acquire a lexicon, indeed arrive at a shared lexicon, whether this lexicon is resistant to changes in the population, whether it scales up to large numbers of meanings and agents, under what conditions shifts in meaning might occur, etc.

There is at present a growing group of researchers constructing these kinds of multi-agent language models and this is giving rise to a totally new type of language theory, one that takes a population view and considers adaptation and evolution as an integral part of linguistics. Overviews and collections of papers can be found in Steels (1997), Briscoe (1999) or Hurford, et al. (1999). Results are so far presented mostly in conferences on artificial life and adapted behavior, but it would be highly beneficial if this methodology also becomes adopted by the

linguistic community as an alternative to generative grammar style modeling.

The Talking Heads experiment

One example of this multi-agent approach is the Talking Heads experiment, which we have been conducting during the summer of 1999. The Talking Heads experiment focused on the origins and evolution of grounded word-meaning and therefore involved robotic agents in the form of steerable cameras in which different agents can be loaded (figure 1). Agents can travel through the Internet and install themselves in robots in different locations so that they are exposed to many different environments. An agent can only interact with another one when it is physically instantiated in a body and thus perceive the shared environment. For this experiment, the shared environments consist of magnetic white boards on which various shapes are pasted: colored triangles, circles, rectangles, etc. The interaction between agents takes the form of a language game, called the guessing game.

The guessing game is played between two visually grounded agents. One agent plays the role of *speaker* and the other one then plays the role of *hearer*. Agents take turns playing games so all of them develop the capacity to be speaker or hearer. The objects located on the white board at the beginning of the game constitute the *context* (figure 2). Agents are capable of segmenting the perceived image and of collecting various characteristics about each object, specifically the color (decomposed in RGB channels), grayscale, and position in pan/tilt coordinates. The speaker chooses one object from the context, further called the *topic*, and gives a linguistic hint to the hearer.

The linguistic hint is an expression that identifies the topic with respect to the other objects in the context. For example, if the context contains [1] a red square, [2] a blue triangle, and [3] a green circle, then the speaker may say something like "the red one" to identify [1] as the topic. If the context contains also a red triangle, he has to be more precise and say something like "the red square". Of course, the Talking Heads do not say "the red square" but use their own language and concepts which are never going to be the same as those used in English. For example, they may say "malewina" to mean [UPPER EXTREME-LEFT LOW-REDNESS]. Such words are atomic units, although the agents may produce and recognise multi-word phrases, but without any syntactic properties.

Based on the linguistic hint, the hearer tries to guess what topic the speaker



Figure 1: Picture of one 'Talking Head' being a steerable camera capturing images about scenes with geometrical figures in front of them. The captured image is shown on a separate monitor.




Figure 2: Image captured by a camera containing various geometric figures. The speaker describes one of the pictures and the hearer has to guess which one was chosen.

has chosen, and he communicates his choice to the speaker by pointing to the object. A robot points by transmitting in which direction he is looking. The game succeeds if the topic guessed by the hearer is equal to the topic chosen by the speaker. The game fails if the guess was wrong or if the speaker or the hearer failed at some earlier point in the game. In case of a failure, the speaker gives an extra-linguistic hint by pointing to the topic he had in mind, and both agents try to repair their internal structures to be more successful in future games.

The architecture of the agents has two components: a conceptualisation module responsible for categorising reality or for applying categories to find back the referent in the perceptual image, and a verbalisation module responsible for verbalising a conceptualisation or for interpreting a form to reconstruct its meaning. Agents start with no prior ontology nor lexicon. A shared ontology and lexicon must emerge from scratch in a self-organised process. The agents therefore not only play the game but also expand or adapt their ontology or lexicon to be more successful in future games.

The Conceptualisation Module

Meanings are categories that distinguish the topic from the other objects in the context. The categories are organised in discrimination trees where each node contains a discriminator able to filter the set of objects into a subset that satisfies a category and another one that satisfies its opposition. For example, there might be a discriminator based on the horizontal position (HPOS) of the center of an object (scaled between 0.0 and 1.0) sorting the objects in the context in a bin for the category 'left' when $HPOS < 0.5$, (further indicated as [HPOS-0.0,0.5]) and one for 'right' when $HPOS > 0.5$ (further written as [HPOS-0.5,1.0]). Further sub-categories are created by restricting the region of each category. For example, the category 'very left' (or [HPOS-0.0,0.25]) applies when an object's HPOS value is in the region [0.0,0.25].

A distinctive category set is found by filtering the objects from the top in each discrimination tree until there is a bin which only contains the topic. This means that only the topic falls within the category associated with that bin, and so this category uniquely filters out the topic from all the other objects in the scene.

Discrimination trees grow randomly by the addition of new categorisers splitting the region of existing categories. Categorisers compete in each guessing game. The use and success of a categoriser is monitored and categorisers that are irrelevant for the environments encountered by the agent are pruned, that is, they are deleted from the discrimination trees.

Verbalisation module

The lexicon of each agent consists of a two-way association between forms (which are individual words) and meanings (which are single categories). Each association has a score. Words are random combinations of syllables. When a speaker needs to verbalise a category, he looks up all possible words associated with that category, orders them and picks the one with the best score for transmission to the hearer. When a hearer needs to interpret a word, he looks up all possible meanings, tests which meanings are applicable in the present context, i.e. which ones yield a possible single referent, and uses the remaining meaning with the highest score as the winner. The topic guessed by the hearer is the referent of this meaning.

Based on feedback on the outcome of the guessing game, the speaker and the hearer update the scores. When the game has succeeded, they increase the score of the winning association by a fixed amount (equal to 0.1 for the experiments in this paper) and decrease the competitors with the same amount, thus implementing lateral inhibition. Scores are bounded between 0.0 and 1.0. When the game has failed, they each decrease the score of the association they used. Occasionally new associations are stored. A speaker creates a new word when he does not have a word yet for a meaning he wants to express. A hearer may encounter a new word he has never heard before and then store a new association between this word and the best guess of the possible meaning. This guess is based on first guessing the topic using the extra-linguistic hint provided by the speaker, and on performing categorisation using his own discrimination trees as developed thus far.

The conceptualisation module proposes several solutions to the verbalisation module which prefers those that have already been lexicalised. Agents monitor success of categories in the total game and use this to target growth and pruning. The language therefore strongly influences the ontologies agents retains. The two modules are structurally coupled and thus get coordinated without a central coordinator.

In the Talking Heads experiment, we gave human users the ability to log in through the Internet, create their agents, follow his progression and that of the group, and teach their agents words. This way we were able to study the interaction between human and artificial language dynamics. More details can be found in a series of technical papers on the various components of the experiment and its results (see for example Steels (1996a), Steels (1996b), Steels (1998), Steels and Kaplan (1999)).

The experiment has been a success in the sense that a shared communication system, in other words a set of words with associated meanings that were suf-

ficiently shared to enable successful communication, emerged after a few days. This system was preserved for the remaining duration of the experiment (four months), despite a steady influx of new agents. This performance was achieved even though many technical problems such as machine crashes and vandalism occurred during the course of the experiment and despite the fact that people could change the environment or create new agents and teach them new words. In total more than 1000 agents were created by human users from many different places in the world using the Talking Heads website (<http://talking-heads.csl.sony.fr/>) which received 10,000 hits per day. The agents played more than 300,000 grounded language games. A total of 8000 words was created covering 500 perceptually grounded categories. But the core lexicon was much smaller, consisting of about 200 words for basic categories like small, large, up, down, red, green, etc.

5 Language evolution

The Talking Heads experiment, and similar multi-agent experiments, exhibit several evolutionary phenomena found in natural language and they hence provide an explanation why these phenomena occur.

Reaching and maintaining coherence

First of all, despite the fact that the individual agents all have a different lexicon and must acquire their lexicon as well as its underlying ontology autonomously, we see that coherence arises in the system, just like a natural language tends towards lexical coherence. This coherence is necessary for achieving successful communication, even though it does not have to be absolutely complete. Synonyms (many words for the same meaning) arise naturally in a group of distributed agents because agents do not have a global view and therefore sometimes create new words not knowing that there are already other words in the population. These synonyms get damped due to positive feedback loop between use and success, causing a progressive self-organisation of the language. This illustrates the effect of causal circularity discussed earlier. Figure 3 shows the damping of synonymy for one word. The graph shows the frequency of the different words used with the same meaning. The meaning identifies a region in the horizontal dimension, and thus corresponds to "to the left". After a struggle among different words, one word ("wogglesplat") stands out and gets preserved subsequently. The other




Figure 3: Word-competition diagram showing how synonymy gets damped. It graphs the word frequency of words competing for the same meaning for a series of 100,000 language games.

words are partly artificial (e.g. bevuwu, bozopite, danuve, fibofure, etc.) partly coming from human languages (e.g. gauche, links, but also red, rouge, yellow). Humans teach by correcting the use of a word in a specific situation, but they only see the situation not how the agent has conceptualised the situation. Because humans may guess another (but compatible) meaning than the one used by their agent, they may accidentally teach their agents ‘wrong’ words (from the viewpoint of their own language). This is how words like ”red” or ”yellow” become associated with ”to the left”.

Polysemy (one word having many meanings) also occurs naturally because agents have to guess the meaning of unknown words and often more than one way to conceptualise reality is compatible with a scene. Polysemy gets damped as well because situations will arise where the different interpretations of the same meaning are no longer compatible. The damping is not complete because in some cases disambiguating situations are hard to come by. Thus polysemy involving a more specific and a more general meaning (for example large and very large) might be maintained. The evolution of polysemy is illustrated in figure 4. It shows the dif-



Figure 4: Meaning-competition diagram showing how polysemy gets damped. It displays the word frequency of meanings competing for the same word for a series of 100,000 language games.

ferent meanings competing for the same word, "droite", another word introduced by humans. We see periods that one meaning ("to the right") dominates, then a shift to another meaning ("to the bottom"), then again a shift to a more specific meaning ("very much to the right"). This kind of evolution from the more general to the more specific is typical for words that are taught to the artificial agents by human users participating in the experiment.

Evolution due to performance deviation

Another well known source of lexical evolution is due to errors in the transmission of word forms. The speaker may erroneously produce a slightly different word form or the hearer may hear something else than transmitted by the speaker. We have performed other experiments varying this source of stochasticity showing an evolution in wordform (see figure 5 from Steels and Kaplan (1998)). The same experiment also shows the impact of a flux in the population. Initially the




Figure 5: Word-diagram showing the impact of population fluxes and errors in wordform transmission. In a closed population (phase 1) a single word dominates. If there are errors in transmission, evolution in the wordform may start to occur.

population stays constant and there is no error on wordform transmission. There is a rapid winner-take-all situation with the word "bagi". Then new agents are entering the population - but still without errors (phase 1). There is no evolution as the word gets learned correctly by the new agents. Next (in phase 3) we have increased stochasticity in wordform transmission. Competitors to the word "bagi" now form: dagi, bogi, pagi, kagi, etc. A period of instability follows after which a new dominant winner "kugo" emerges. The higher the errors in transmission the quicker an existing convention can destabilise.

Factors in language evolution

These various experiments and simulations show how a population framework can be used to study language evolution. There is individual variation among language users as well as deviations in performance due to errors in transmission of word forms, errors in feedback, etc. Variations may spread in the population,

either in a completely arbitrary fashion (there is often no reason why a particular word is used to express a certain meaning) or due to selectionist pressures naturally arising in the system. Thus meanings which are more robustly recognised in the visual image will tend to have an easier time to maintain themselves as the dominant meaning of a word, or, if two words are competing for the same meaning, they weaken each other giving a chance to a newcomer to overtake both. There is an additional impact from the influx of new agents because they are most sensitive to the most common lexicon and hence they tend to simplify the total lexicon for the future.

In a typological view, universal tendencies in language (which undoubtedly exist) are ascribed to Universal Grammar, which has the status of a Platonic ideal form, innately known. In an evolutionary populationist view, universal tendencies are the consequence of various constraints within which the evolutionary process operates. The constraints come from three sources: the physiological and neural apparatus available to humans for language, the system already in existence, and pressures coming from using and maintaining an adequate communication system, which implies robustness against errors, sufficient power to handle complex meaning, and learnability. So we arrive at a whole new research program for linguistics, one that tries to explain the characteristics and universal tendencies in language rather than merely describe them and then declare them to be innate.

6 Conclusions

Language evolution should be at the core of linguistics. It is our only hope for developing an explanatory rather than a descriptive theory of language, and for capturing the full richness of human natural languages. But a sensible study of language evolution requires that we take a population view, as opposed to a Cartesian typological view, which in turn implies that we construct multi-agent models based on populations of agents whose internal states reflect evolving knowledge of the linguistic conventions emerging and maintained by the group.

Experiments like the Talking Heads experiment illustrate this approach and already show it to be extremely promising, both from a theoretical viewpoint because they enable us to test evolutionary theories of language and from a technological viewpoint because the methods, techniques and technologies required for doing evolutionary multi-agent simulations and experiments help us develop an alternative language technology, which is no longer based on implementing a

frozen state of a language but which adapts continuously to the language present in the environment of human or artificial agents.

7 Acknowledgement

Many important contributions to the Talking Heads experiments have been made by members of the VUB Artificial Intelligence Laboratory in Brussels and the Sony Computer Science Laboratory in Paris. I am particularly indebted to Angus McIntyre, Frederic Kaplan, and Joris van Looveren. The work in Brussels has been financed by a GOA grant of the Belgian government.

8 References

Briscoe, T. (1999) *Linguistic evolution through language acquisition: formal and computational models*. Cambridge Univ. Press, Cambridge.

Chomsky, N. (1981) *Lectures on government and binding*. Foris Publications, Dordrecht.

Genermont, Jean. (1998) Qu'est-ce qu'une espece? *Pour la Science*. p. 30-32.

Mayr, E. (1975) *Evolution and the diversity of Life*. Harvard University Press, Cambridge.

Heine, B., U. Claudi and F. Huennemeyer (1991) Grammaticalisation. A conceptual framework. The University of Chicago Press, Chicago.

Hopper, P. (1987) Emergent Grammar. *Berkeley Linguistic Society* 13, 139-157.

Hurford, J. C. Knight and M. Studdert-Kennedy (eds.) (1999) *Approaches to the Evolution of Human Language*. Cambridge Univ. Press. Cambridge, 1999.

Lightfoot, D. (1991) *How to set parameters: arguments from language change*. The MIT Press, Cambridge Ma.

Maynard Smith, J. (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge.

Steels, L. (1996a) Self-organising vocabularies. In: Langton, C. (ed.) *Proceedings of Artificial Life V*. The MIT Press, Cambridge Ma.

Steels, L. (1996b) Perceptually grounded meaning creation. In: Tokoro, M. (ed.) (1996) *Proceedings of the International Conference on Multi-Agent Systems*. AAAI Press, Menlo Park Ca. pp. 338-344.

Steels, L. (1997) The synthetic modeling of language origins. *Evolution of Communication*, 1(1):1-35.

Steels, L. (1998) The origins of syntax in visually grounded robotic agents. *Artificial Intelligence* 103 (1998), 1-24.

Steels, L. and Kaplan, F. (1998) Situated Grounded Word Semantics. In *Proceedings of IJCAI-99*, Stockholm. Morgan Kaufman Publishing, Los Angeles. p. 862-867.

Steels, L. and Kaplan, F. (1999) , In *Proceedings of COLING-ACL*, Montreal, August 1998, p. 1243-1249.