A First Report on Emergent Phonology

Bart de Boer
AI-lab
Vrije Universiteit Brussel

Abstract
This paper contains a description of two experiments that have been done to test self organisation and cultural evolution as possible mechanisms for explaining the origins of language in general and the origins of phonological systems in particular. Using computer simulations of agents that try to communicate using speech sounds, together with a selectionist algorithm where successful phonemes are kept and unsuccessful ones discarded, an adaptive and open-ended phonology is created.
The first experiment involves a rather abstract model of human speech, while the second experiment involves a (simplified) physical model. In both experiments the same selectionist algorithm was used. The first experiment was meant to test the basic idea of self-organisation and cultural evolution, while the second one was meant to test these ideas in a more complex and realistic setting. Both have shown that communicative success can increase while communicative complexity also increases.
The implications of this for linguistics are discussed, as well as possible and necessary future research.

1. Introduction

Phonology is the study of sound systems of languages [9]. It studies the set of sounds in a language, as well as the rules that govern the possible combinations of sounds. Phonology is based on phonetics, the study of the production and perception of human speech sounds in general [13]. Where phonology is language dependent, phonetics is language independent. Phonology itself is essential for the description of languages. It also provides a basis for, for example, morphology and of the study of language change.
Current research into phonological theory investigates mostly the processes that transform internal representations of what a speaker wants to say into the actual sounds that are produced. Usually a number of rules and a number of levels of operation are postulated for this. Highly influential is the theory of Chomsky and Halle[5], but other theories exist as well [21].
Phonologists also study the patterns that underlie the sets of sounds in human languages. The sounds that occur in any given language are not a random subset of the possible speech sounds. On the contrary: highly regular patterns determine the occurrence and co-occurrence of speech sounds throughout the world's languages. These are called phonological universals [7][16].
Phonologists try to explain this using distinctive feature theory [11]. Features are (usually) binary oppositions between groups of phonemes. For example, /p/ and /t/ are [-voiced], whereas /b/ and /d/ are [+voiced]. Languages "choose" a number of features to make distinctions between speech sounds. Phonological distinctions are learnt by learning to distinguish the features. It is found that some features are always learnt before other features, and that some features imply the presence of other features. Opinions differ as to what exactly the different features should be, but agreement exists that there should be one universal set for all languages[1].
Where phonological features come from, and how we learn and use phonological rules is still a subject of study and debate in phonology, although most phonologists assume that a lot of structure is innate and therefore the same for all humans. Unfortunately, the question where the structures came from in the first place is not often addressed but see [17][19].
In this paper a new point of view on phonology is proposed. This point of view is based on the theories developed by Luc Steels [23]. The basic point in which his theory differs from the more current theories in linguistics is that language is a collective phenomenon. Nobody is in control of the language, and no speaker has a complete and accurate image of the language. According to Steels, there is not such a thing as de Saussure's *langue* [6] nor does an idealised language competence as assumed by Chomsky [4] exist.

[1] Although, obviously, not all languages use all features.

In Steels' theory language is something that emerges through the interaction of language users trying to communicate with each other and learning the language from each other. Coherence is created by self-organisation, not by innate language competence[2]. Language complexity comes to be through processes such as evolution (cultural, not genetic), co-evolution and level formation, each investigated in complex systems' theory and observed in other complex systems, such as biological ones [17].

The consequence is that language is seen as an open system, in which none of the actual language users has complete control, nor complete knowledge of the language, as it is slowly, but constantly changing. Language is also a complex adaptive system, as it adapts to new circumstances. New words as well as new ways to express things come to be, and the language adapts itself to the changing needs of the speakers and to the new ways in which the speakers apply it. Examples can be found in the emergence of pidgin and creole languages, in the splitting of a language in a written form and a spoken form once writing systems are introduced (for example in Tibetan[8]) or in the emergence of group languages and slang [24].

These theories have been applied to a number of linguistic fields, particularly in the formation of semantics and a lexicon [22]. In the research described in this paper a first attempt is made to investigate the possibilities to apply the theory to phonology. In the next section the framework of the experiments will be described. In section 3 a first, simple experiment is described. In section 4 a more open-ended and realistic experiment is described. In section 5 the implications of these experiments are discussed, and in the last section the future possibilities of the experiments are discussed.

2. The Experiments

The experiments described in this paper are based on an artificial life like approach to linguistics. Artificial life[14**,**15] is the study of complex (biological) processes using computer simulations. More in general it is a synthetic approach to science, as opposed to the more common analytic approach. In the analytic approach, complex systems are observed, described and analysed in order to understand them. In the synthetic approach, possible mechanisms are implemented using computer simulations. Mechanisms that produce behaviour that is similar to the observed behaviour of the real system are regarded as plausible models of the real system. In practice both methods will be used at the same time: analysis provides possible mechanisms that can be tested by computer simulations, which then provide feedback for further analysis.

Our computer simulations are based on two mechanisms: selectionism and simulation games. Selectionism operates analogously to natural evolution: possible solutions to a given problem are generated and tested, and the most successful solutions are kept. The least successful ones will be thrown away. In our system the problem is communication under certain constraints. The possible solutions are phonemes and combinations of them. The generation of phonemes in our system is random, but it is quite possible, and probably even necessary to generate phonemes in a more intelligent way. This will be discussed in section 6.

Language games[3] in Steels' terminology [23] are interactions between agents in which the agents use language to communicate certain information, together with a number of rules of how the interaction should be structured and a definition of when such a game is successful. Agents can take part in multiple simultaneous language games. The goal of language games is to communicate information, either information about the agents' environment or about the language itself. If the communication is about the language itself, usually a non-verbal means of communication, such as gestures or observation of the other agent's actions will be necessary.

A language game in Steels' and Wittgenstein's terminology are also always about an outside reality, about which the agents try to tell each other something. In our experiments this was not (or only marginally) the case. Our agents played games only with sounds that were not related directly to an outside reality (except the physical reality in which the sounds are produced and perceived). We will therefore talk about *imitation games*.

In our experiments two agents are involved in the imitation game. The purpose of playing the imitation game is getting a phonology that is more or less the same for both agents, starting from scratch. In order to

[2] There remains the possibility that things that had to be learnt at first, later become innate through genetic evolution and the so-called Baldwin effect [1].

[3] The term language games is used somewhat differently by us than by the inventor of the term, Wittgenstein[25].

play the imitation game both agents build up a list of words and a list of phonemes (see Figure 1). Both are initially empty.

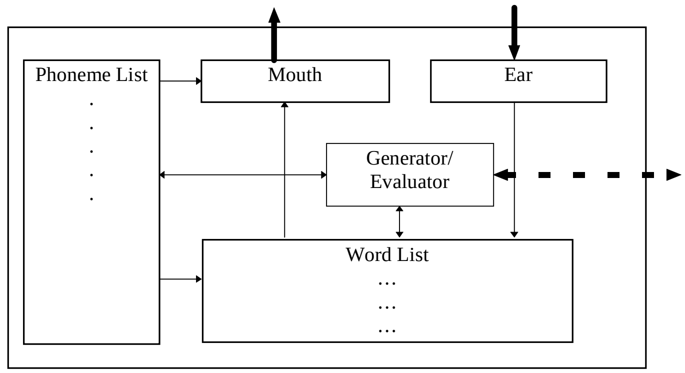

Figure 1: The agent

The short-term goal of the imitation game is imitation. If one agent succeeds in imitating the sound produced by the other agent, the game is considered successful. The rules of the game are as follows: an agent, the initiator, generates a word, which is built up of a number of phonemes. This word is then heard by the other agent (the replicator) and analysed into the phonemes this agent knows. The replicator then produces the word that is built up of these phonemes. The initiator listens to this word and analyses it into its own phonemes, which are then compared to the original word. Depending on the similarity, non-verbal communication is used to signal the success of the imitation game to the replicator. The phonemes and words that were used in the exchange are rewarded according to the overall communicative success. Agents take turns in playing initiator and replicator. The imitation game is illustrated in Figure 2.

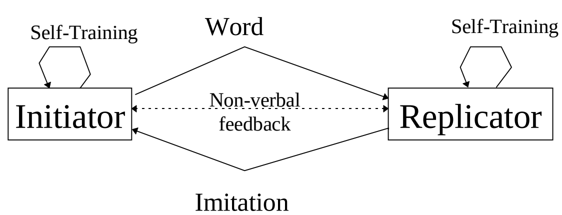

Figure 2: The imitation game

The selectionist mechanism operates on the words of an agent. Words with a high score (compared to the other words of the agent) are kept, and words with a low score are thrown away. If, after the words have been thrown away, there are phonemes that are not used by any word that is left over, these will be thrown away as well. From time to time, new words are generated, and if it is too hard to generate a new word with the existing phonemes, a new phoneme is generated.

Some remarks should be made on the assumptions that have been made in our experiments. The first assumption is that speech is produced and analysed in terms of phonemes. This is not a very controversial assumption, but there exist theories that claim that speech is processed on either the lower level of features [5][11] or on the higher level of syllables [3][18]. Both theories have their advantages and disadvantages, and we will not discuss them here, but we should remark that the way humans analyse speech in terms of phonemes is probably also an outcome of a selectionist (evolutionary) process. The fact that we start with an analysis into phonemes already present should therefore be considered as a theoretical bias.

Two other assumptions, that are not linguistic assumptions, but just simplifications of our computer simulation are that there is no noise and agents never speak at the same time. Speech sounds are perceived as well as they are produced. Also we assume that there is no cost in producing speech. This is unrealistic, but the utterances of our system are so short that a cost would probably not make much difference.

3. A First System

The first system that will be described was designed to test the ideas of self-organisation and co-evolution. In this system no modelling of the processes of articulation or perception was undertaken. The agents communicate with each other by transferring strings of features that represent phonemes. In order for the system to remain interesting, neighbouring phonemes influence each other. From this system one cannot yet expect results that are interesting for the study of human phonological systems: for this the possible variation is too limited.

		[-fric]		[+fric]	
		[-v]	[+v]	[-v]	[+v]
[-back]	[-sh]	p/m̥	b/m	f/m̥	v/m
	[+sh]	t/n̥	d/n	s/n̥	z/n
[+back]	[+sh]	c/ɲ̊	ɟ/ɲ	ʃ/ɲ̊	ʒ/ɲ
	[-sh]	k/ŋ̊	g/ŋ	x/ŋ̊	ɣ/ŋ

Table 1: Example of [+cons] phonemes. Per entry [-nasal]/[+nasal] are shown

Agents try to communicate with each other by transferring a sequence of phonemes which are represented by a number of features. Agents do not have to worry about when a phoneme starts and when it ends. In a realistic system this would be a big problem (see section 4). However, in our system, phonemes are transferred one by one.

		[-round]		[+round]	
		[-low]	[+low]	[-low]	[+low]
[-back]	[+sh]	i/ĩ	ɛ/ɛ̃	y/ỹ	œ/œ̃
	[-sh]	e/ẽ	a/ã	ø/ø̃	ɶ/ɶ̃
[+back]	[-sh]	ɤ/ɤ̃	ʌ/ʌ̃	o/õ	ɔ/ɔ̃
	[+sh]	ɯ/ɯ̃	ɑ/ɑ̃	u/ũ	ɒ/ɒ̃

Table 2: Example of [-cons] phonemes. Per entry [-nasal]/[+nasal] are shown

Phonemes are represented by a list of eight features. For consonants: [+consonantal], [+/-nasal], [+/-back], [+/-sharp], [+/-voice], [+/-fricative], [+/-long] and [+/-high tone] and for vowels [-consonantal], [+/-nasal], [+/-back], [+/-sharp],[+/-low], [+/-rounded], [+/-long] and [+/-high tone]. It must be said that the features presented here have been chosen rather to make it possible to present the phonemes to a human observer of the experiment than to provide a firm theoretical basis. The possible consonants are given in Table 1, the possible vowels are given in Table 1.

The features are used to determine the influence of phonemes on each other. Phonemes in sequence in all human languages are found to influence each other. This has been explained by rules that determine the influence of phonemes on each other or by reference to dynamic aspects of the articulatory process [2]. In our system the influence of phonemes on each other is determined by a subset of the features of the phonemes, their *crucial* features.

The crucial features of phonemes, which are enumerated for every phoneme, are those features that are crucial for the recognition of a phoneme. If these features are present in a sound, the phoneme is said to be recognised. For example, in nasals in English, in many circumstances only their nasality is crucial. Thus they assimilate in place of articulation (“impossible” [ˌɪmˈp^{h}ɔsəbl̩], but “incorrect” [ˌɪŋkəˈɹɛkt]) or they loose their consonantal character completely and nasalise the preceding vowel (“can’t” becomes [k^{h}æ̃t]). In our system crucial features are considered strong: they are able to transfer their values to other phonemes.

The influence works as follows: if a phoneme appears between two other phonemes, or between one phoneme and the word boundary, and if the neighbouring phoneme(s) have the same value for a feature that is marked crucial for both, the intermediate phoneme takes over the value of this feature.

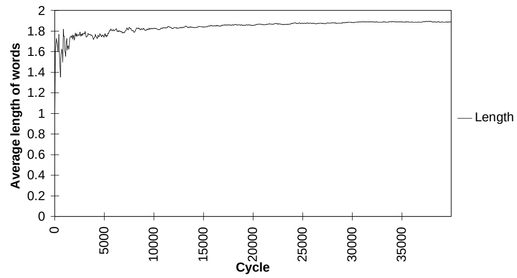

Figure 3: Word length in the short word experiment

From the simulation a number of observations can be made: the success of communication increases, even though the number of phonemes and the number of words increases. Also the length of the words rises to a certain maximum, which depends on certain parameters, such as the likeliness with which the system creates new phonemes, and the speed with which words are allowed to grow in length.

The results of two experiments will be presented and compared in some detail. Both experiments consisted of letting two agents evolve a phonology based on the system described above. In the first experiment (the experiment with the short words) words grew slowly. In the second experiment (the experiment with the long words) words grew more quickly. In both cases the agents were allowed to communicate 40 000 times.

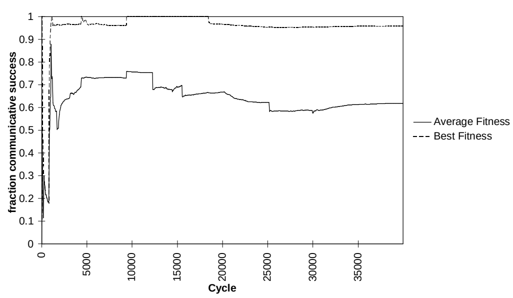

Figure 4: Fitness in the long word experiment

Let us first look at the short word experiment. The average length of the words stayed below 2 (see Figure 3), so in order to create the thousand or so words that are used in the system, some fifty phonemes were present in the end (see Figure 7). Communicative success of individual phonemes was slowly rising, and the best phoneme of the agent has a communicative success of nearly 100% (Figure 5). Of course an average communicative success of 100% for all phonemes can never be reached, as new phonemes are constantly created.

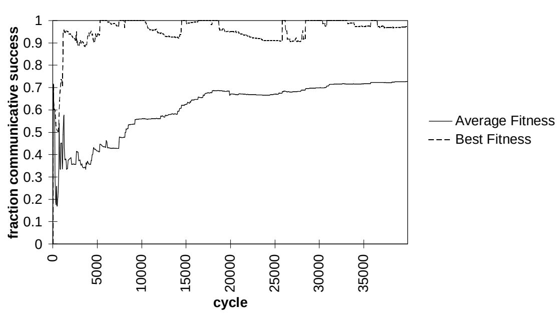

Figure 5: Fitness in the short word experiment

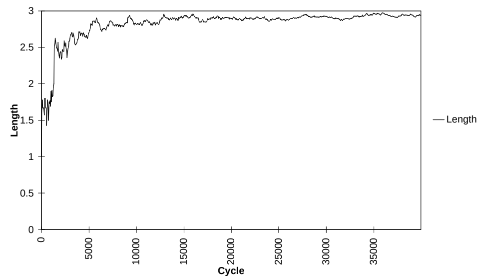

Figure 6: Word length in the long word experiment

These results can be compared with the ones of the experiment with longer words. The average length of words in this experiment approached 3 (Figure 6). In this experiment communicative success of the best phoneme was nearly 100% (Figure 4), but the average success was a bit lower and fluctuated more. This can be explained by taking into account that the influence of phonemes on each other is bigger in longer words. An interesting observation is that, because the words are longer, less phonemes are needed in order to make the words. For the 600 words that are present in the end, only 15 phonemes suffice(Figure 8). In the experiment with the short words, 38 phonemes were present for 600 words.

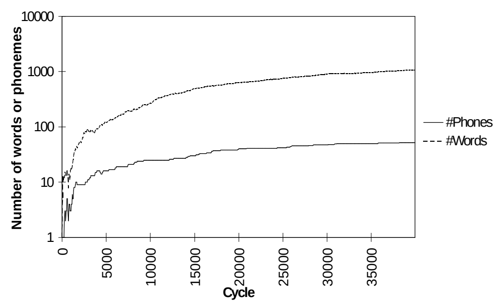

Figure 7: Number of words and phonemes in the short word experiment

The forms of words were also interesting: in some cases agents succeeded in imitating each others words perfectly. In other cases the imitations resulted in phonemes that were different in one or two features. For example, in one experiment a successful phoneme was pronounced [y] by one agent and [ø] by the other, the feature [sharp] not being crucially distinctive in this case. We can also compare the sets of the most used (more than 5 times) and most successful (>70%) phonemes that were found in the agents after a period (10 000 cycles) of evolution. In the first agent the phonemes were: [ø̀ː], [ɛ́], [ó], [ɲ̀], [ŋ́ː], [ɟ́ː] and [ɔ́ː]. In the

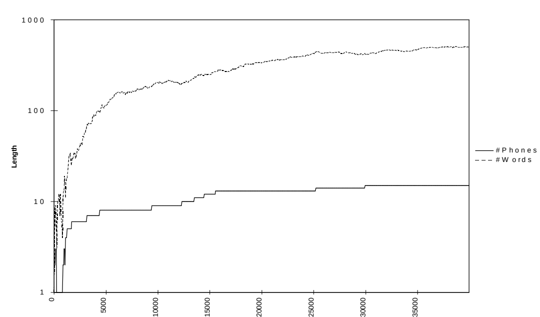

Figure 8: Number of words and phonemes in the long word experiment.

second agent the corresponding phonemes were: [ø̀ː], [œ́], [ó], [ɲ̀], [ŋ́ː], [ɠ́ː] and [ỳː] and [ɤ̀ː]. The diacritics [´] and [`] stand for [+high tone] and [-high tone] respectively, and [ː] stands for the feature [+long]. The correspondence between most of the phonemes is clear. The [ɛ́] and [œ́] differ only in the feature [round] and the [ʄ́ː] and [ɠ́ː] differ only in the feature [sharp]. The phonemes [ɔ́ː] in the first agent and [ỳː] and [ɤ̀ː] in the second agent are more problematic. It is not very likely that they correspond with each other, rather that they correspond with a set of phonemes of the other agent, depending on their context.

4. Towards a More Realistic System

The experiment that has been described in the previous section was based on a rather abstract model of human speech. In order to investigate more human-like patterns of speech , a more realistic system needs to be built. A first attempt to do this was based on a very simple articulatory speech generator, a dynamic model of articulation and a neural-network model of speech perception. The difficulties of signal processing in this experiment caused the results to be less than those of the previous experiment, but nevertheless some real signs of the emergence of a primitive coherence of speech sounds have been found.

The speech synthesiser can be controlled by four articulatory parameters: lip rounding, tongue position, tongue height and voicing. These parameters are transformed in three formant frequencies using a simple linear transformation based on some observations by Ladefoged [12:ch. 8]. The relative strengths of these formant frequencies are calculated using another simple heuristic, where formants with higher frequency are

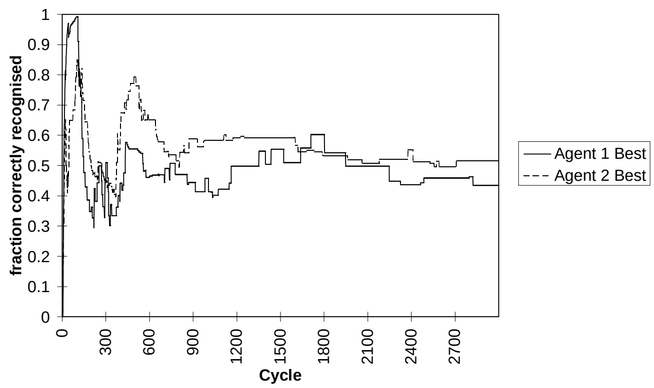

Figure 9: Best score of phonemes in the complex system.

weaker, especially if the lips are rounded.

The synthesiser is able to produce vowels in a recognisable way, and also to produce an approximation of (glottalised) stops. However, in further experiments a more realistic model would be needed, as our present model is not able to produce a large enough variation of speech sounds.

The articulation of speech sounds is dynamic (based partly on the work of Browman and Goldstein [2] and Saltzman [20]). Phonemes are described by a number of goal values of the articulators. If the phoneme has to be articulated, the articulators start moving towards this goal position. If the next phoneme starts, the articulators start moving towards this new goal position, even if the first goal position was not already

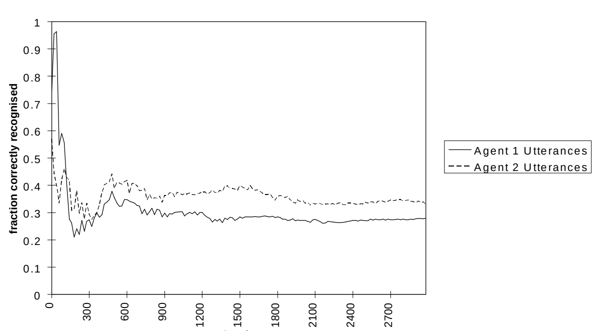

Figure 10: Average score of phonemes in the complex system.

reached. In this way three effects are captured: the incomplete articulation of phonemes, the influence of phonemes on each other and the transitions between phonemes that can be observed in human speech. However, no attempt was made to capture the idea of crucial features of the previous experiment.

The processing of the speech signal by the listener is done by partitioning it using a very simple heuristic, and by trying to recognise the segments using a neural network. The partitioning is necessary, in order to slice up the continuous input signal in snapshots that the neural network can handle. Such a snapshot is taken if the change of the frequency spectrum is less than a certain threshold (this would indicate a vowel in our system) or when a signal appears after silence or silence appears after the signal (the first case would indicate the end of a stop consonant, the second case the beginning of a stop consonant). The frequency spectrum is then fed into a neural network.

The neural network is based on the perceptron [10:ch. 5], and is organised as follows: for each *frequency bin* in the frequency spectrum (32 in our system) there is an input node in the neural network. For each phoneme there is an output node. The network is trained in such a way that the output nodes produce a maximum signal if the frequency spectrum is prototypical for the phoneme that is associated with the node. The network functions in a winner-take-all way, so that the node with the maximum strength wins, and the associated phoneme is considered recognised.

The agent trains its networks by talking to itself (the self-training arrow in Figure 2). It says words to itself, and listens to the sounds it produces. It then reinforces the connections between the frequency bins and the node of the phoneme that it is currently saying.

Some variations on the speech recognition have been tried out, such as using feature detectors, that react to

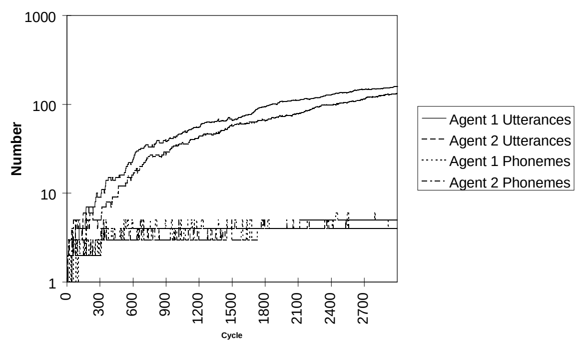

Figure 11: Number of phonemes and number of words in complex system.

certain events in the speech signal, or the use of special detectors for the transitions between phonemes. Both of these approaches were not sufficiently successful for separate presentation. Both the feature detectors and the transition detectors grew to fast in number. This does not mean that the approaches are without value. Just that the main focus of this research is on the selectionist, self-organising aspects of the system. When these are understood more fully, we can work on better signal processing.

The results have been less impressive than in the previous experiment. The rate of success of phonemes was much lower than in the previous experiment. The results in Figure 9 and Figure 10 and can be compared with those of Figure 5 or Figure 4. Although these results are much worse, they are still better than random performance, if one compares the number of phonemes in Figure 11. In random performance one would expect the fraction correctly recognised phonemes to be the reciprocal of the number of phonemes. In fact the fraction is higher.

To give an idea what kind of conversations take place in the complex system, an utterance and the attempt at imitation are presented in Figure 12. Note that in fact the imitation is quite like the original utterance. Although this example was selected because it showed a good imitation, and there are a lot of conversations in which the agents do not succeed in properly imitating each other, there is still quite a large number of conversations in which the imitation is as good as this one. The average performance is quite low, because there also is a large number of conversations that fail totally.

A definite problem of the complex system is that the number of phonemes stays very low, so that rather long words have to be formed, which hinders recognition. This indicates that the sound system of the agents petrifies too quickly, probably because the neural networks are overtrained on the first few phonemes.

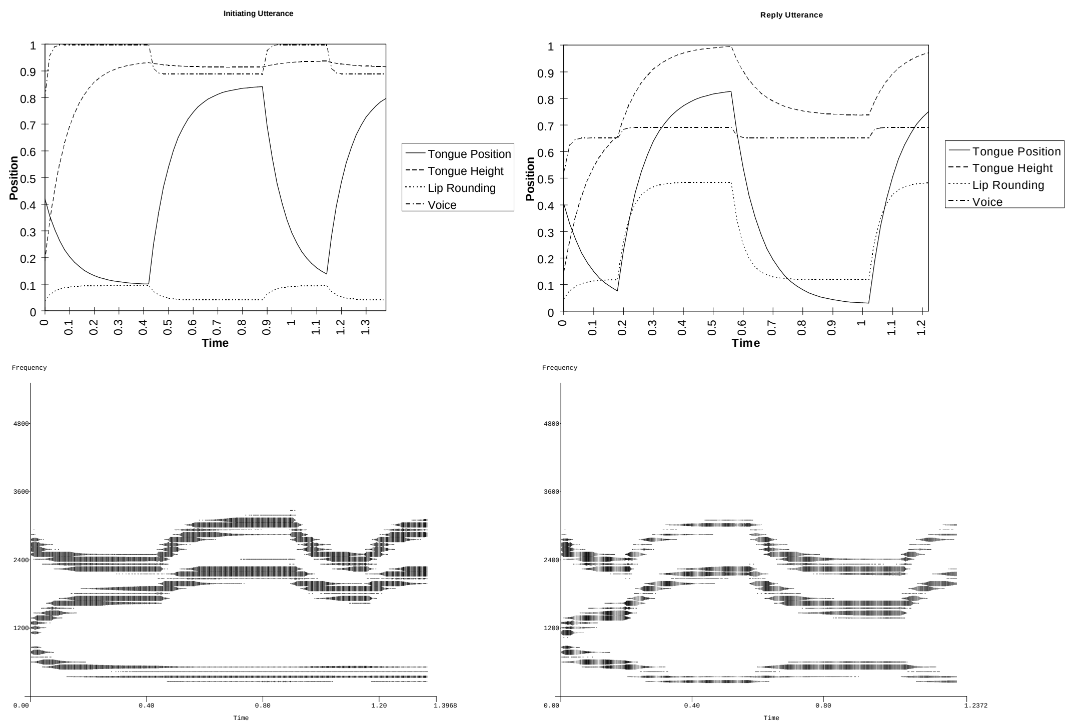

Figure 12: A sample conversation of the complex system. Top left is the gestural score of the initial utterance, bottom left the acoustic image, top right is the gestural score of the imitating reply, bottom right is its acoustic image. Note that apparently the *timing* does not matter, only the position of the formant frequencies.

5. Discussion and Conclusion

The above experiments have shown that it is possible for agents that do not have any linguistic knowledge in common, but that do have the proper speech production and -processing capabilities, to develop a common set of speech sounds with which to communicate. Although the experiments so far have not been very realistic (but there is a number of ways in which they could be improved, see section 6), we still might want to consider their implications for linguistics in general and the study of phonology in particular.

The question the experiments were supposed to answer was whether self organisation and evolution[4] could be the basis of a coherent set of speech sounds for communication between two agents. This was based on Luc Steels' theory [23] that these processes might be responsible for the origin of language. The results of the experiments are positive in the sense that a higher than random communicative success and an increase in communicative complexity was found.

If the phonologies of languages are the result of an evolutionary process and of self organisation, then phonological universals must also be the result of this process, and not of innate processes. The contrary would be true: phonological universals, caused by evolution and self-organisation might be the (indirect) cause of genetic change through the Baldwin-effect. The universals will then have a functional cause: the combination of evolution and self-organisation picks out certain solutions to the problem of communication through speech more easily than others, because their communicative success is higher.

Another conclusion is that there is no such thing as an ideal and complete description of the phonology of a language. The phonology is an emergent phenomenon of the interactions between the speakers. No speaker has a complete and accurate picture of the phonology. Change and variation are inherent in such a system, as everybody probably has a slightly different knowledge of the phonology. Change is caused by amplifica-

[4] Meaning evolution of the language itself, not of the genetic makeup of the speakers. We are therefore talking about cultural evolution, not about genetic evolution. The language can evolve even when the speakers themselves do not change.

tion of the variations, which in turn is dependent on external, mostly social factors. In this way rapid cultural evolution can take place.

The author realises that these ideas differ radically from those that are current in linguistic research. They put to question the central position of Chomsky's [4] idealised language competence or de Saussure's *langue* [6]. Not only do they imply that nobody has knowledge of the ideal language, they also imply that it does not exist at all, and is not even necessary. Communicative success is assured by constant adaptation and self-organisation, not by a standard to which everybody has to confirm.

Even if one does not accept these conclusions, the approach to phonology presented in this paper might be an interesting new way of investigation. Theories about the origins of phonology and of phonological change can now be tested in another way than by linguistic observation and reconstruction. Implementing them as computer models using techniques from artificial life can bring to light their strengths and weak points and makes comparison of different theories easier.

6. Future work

The systems described in this paper have been rather simplistic. Three ways of extending the research should be explored. First of all the self-organising and evolutionary mechanisms should be extended. The creation of new phonemes is random in our system, but probably needs to be guided by knowledge about phonemes that are already present and by observations of the other speaker's phonemes. Also the detection and recognition of phonemes should evolve from general to specific. If an agent only has a few phonemes, its recognition mechanism can be crude. This mechanism should become more and more specific if more phonemes are created.

Secondly the speech production and speech recognition models should be made more realistic. An articulatory synthesiser based on a physical model of the human vocal tract should be added, together with a more complete model of the dynamics that underlie articulator movements, so that more and more complex sounds can be made. When this is done, the processing of speech sounds should be made more powerful accordingly.

A third way of extending the research is by making variations on the imitation games. At the moment here is one type of interaction: that between two agents that start without any knowledge. Variations could be: interactions between multiple agents, between agents with different levels of knowledge (for example one without any knowledge and one or more with a developed phonology), between agents that already have knowledge of phonology, but of different phonologies, and experiments where a large number of agents with a spatial distribution communicate, and where nearby agents have a higher probability of communicating with each other than agents that are further apart.

7. Acknowledgements

This work is part of an ongoing research into the origins of intelligence and language, financed by the Belgian federal government FKFO project on emergent functionality (FKFO contract no. G.0014.95) and the IUAP 'Construct' project (no. 20). It would not have been possible without the suggestions and advice of Prof. Luc Steels. I would also like to thank Egbert Boers of the computer science department of Leiden University, the Netherlands for advice on the Baldwin effect and dr. George van Driem of the Himalayan Languages project of Leiden University, the Netherlands for advice on the Tibetan language.

8. Literature

1. Baldwin, J. Mark (1896) A new Factor in Evolution, *The American Naturalist* **30** (June 1896) pp. 441–451, 536–553, reprinted in: R.K. Belew and M. Mitchell (eds.) *Adaptive Individuals in Evolving Populations: Models and Algorithms*, SFI Studies in the Sciences of Complexity, Proc. Vol. XXVI, Addison Wesley, Reading, MA, 1996.
2. Browman, Catherine P. and Louis Goldstein (1995) Dynamics and Articulatory Phonology, in: robert F. Port and Timothy van Gelder (eds.) *Mind as Motion*, MIT Press, Cambridge Mass. pp. 175–194.
3. Carré, René, Marc Bordeau and Jean-Pierre Tubach (1995) Vowel-Vowel Production: The Distinctive Region Model (DRM) and Vowel Harmony, *Phonetica* **52**, pp. 205–214
4. Chomsky, Noam (1965) *Aspects of the Theory of Syntax*, MIT Press, Cambridge, Mass.
5. Chomsky, Noam and Morris Halle (1968) *The sound pattern of English*, MIT Press, Cambridge, Mass.

6. De Saussure, Ferdinand (1987) *Cours de linguistique générale, édition préparée par Tullio de Mauro*, Paris, Payot.
7. Greenberg J.H. (1966) Synchronic and Diachronic universals in phonology, *Language* **42**, number 2 pp. 508–517
8. Hahn, Michael (1985) *Lehrbuch der klassischen tibetischen Schriftsprache*, Bonn: Indica et Tibetica Verlag.
9. Hawkins, Peter (1995) *Introducing Phonology*, Routledge.
10. Hertz, John, Anders Krogh and Richard G. Palmer (1991), *Introduction to the Theory of Neural Computation*, Addison Wesley.
11. Jakobson, Roman and Morris Halle (1956) *Fundamentals of Language*, the Hague: Mouton & Co.
12. Ladefoged, Peter (1981), *Preliminaries to Linguistic Phonetics*, Midway Reprint, The University of Chicago Press.
13. Ladefoged, Peter and Ian Maddieson (1996) *The Sounds of the World's Languages*, Blackwell.
14. Langton, Christopher G.(ed.) (1989) *Artificial Life*, Addison Wesley
15. Langton, Christopher G, Charles Taylor, J. Doyne Farmer, Steen Rasmussen, (eds.)(1990) *Artificial Life II*, Addison Wesley.
16. Lindblom, Björn and Ian Maddieson (1988), Phonetic Universals in Consonant Systems, in: Hyman, Larry M. and Charles N. Li (eds.) *Language, Speech and Mind*, pp. 62–78.
17. Maynard-Smith, John and Szathmary Eörs (1994), *The major transitions in evolution*, Oxford: Oxford University Press.
18. Öhman, S. (1966) Coarticulation in VCV utterances: Spectrographic measurements, *J. acoust. Soc. Am.* **39** pp. 151–168.
19. Pinker, S. and P. Bloom (1990) Natural Language and Natural Selection. *Behav. Brain Sci*. **13**, pp. 707–784.
20. Saltzman, Elliot L. (1995) Dynamics and Coordinate Systems in Skilled Sensorimotor Activity, in: Robert F. Port and Timothy van Gelder (eds.) Mind as Motion, MIT Press, Cambridge, Mass. Pp. 149–174.
21. Skousen, R. (1973) Evidence in phonology, in: C. W. Kisseberth (ed.) Studies in generative phonology, *Papers in Linguistics*, monograph 3
22. Steels, Luc (1996) The Spontaneous Self-organization of an Adaptive Language, in: S. Muggleton (ed.) *Machine Intelligence* **15**.
23. Steels, Luc (1997) Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation, in: J. Hurford (ed.) *Evolution of Human Language*, Edinburgh: Edinburgh University Press.
24. Trudgill, Peter (1995) *Sociolinguistic: An Introduction to Language and Society*, Penguin books.
25. Wittgenstein, Ludwig (1967) *Philosophische Untersuchungen*, Frankfurt: Suhrkamp.