Self-organisation in language

Bart de Boer

AI-lab
Vrije Universiteit Brussel
Pleinlaan 2
1050 Brussel
Belgium
bartb@arti.vub.ac.be

Abstract. This paper gives an overview of the research into self-organisation in language. Self-organisation is viewed as the emergence of global order through local interactions. It is shown that self-organisation according to this definition can occur in many ways in natural language. It can occur as in linguistic representations in the brain, but also in a population of language users.

Work on the different aspects of self-organisation is reviewed. Such work generally makes use of computer simulations, so a short overview is given about the possible ways in which computer simulations can be used. It is shown that such work has given and continues to give a valuable contribution to the study of language.

Language is possibly one of the most complex aspects of social behaviours. It is crucial for communicating information between humans. Without language, human communication could not be as complex as it is. Language is also a very strong indication of social status. People are very good at inferring from accents (the way words are pronounced, or subtle variations in lexicon and grammar) other people's social and geographical background. Furthermore, language is one of the largest bodies of culturally transmitted information.

At the same time language clearly is the product of evolution. While human infants learn language effortlessly and automatically, even humanity's nearest primate relatives (Chimpanzees and Bonobos) can hardly learn any language at all. This indicates that there must be genetic factors influencing language acquisition. The precise nature of these factors is the subject of intense debate in linguistics. However, it is clear that language must be the result of the interaction between social mechanisms and genetic evolution.

This paper addresses the issue of the role of self-organisation in language. Self-organisation is here defined as *the emergence of order on a global scale through interactions on a local scale*. This is a phenomenon that is found in social systems of much simpler animals, for example in colonies of social insects, so one would expect it to occur in human's most complex social behaviour as well. And indeed, linguists and philosophers have been investigating the role of self-organisation in language for a long time. As there are many definitions of self-organisation, the one used here is explored in some more detail in the next section. Some background on the history of research into self-organisation in language is given in the section after that.

The definition of self-organisation as stated above is very broad and can be applied in different ways to language. One could take the local interactions to be interactions between neurons in the brain and the global order to be linguistic concepts, such as phonemes, words or grammatical classes. But one could also take the view of language as a social system, and take interactions between humans as the local scale and the coherent language of the population as the order on the global scale. Research on both sides of this dichotomy is presented. Also, one can apply the concepts of self-organisation to different aspects of language (sound systems, lexicon, grammar, semantics) and to either a synchronic (how language is) or diachronic (how language changes) view of language. The different ways in which self-organisation can and has been applied to the study of language will be presented in the section on diachronic and synchronic aspects of self-organisation in language, although not all researchers necessarily use the term to describe their work. A selection of this work is presented in the last section of the paper.

About the definition of self-organisation

In fact, many papers that describe processes that resemble self-organisation according to the definition used here do not explicitly use the term self-organisation. Instead terms like "emergent behaviour," "population dynamics," "bifurcations," "catastrophes" and others are used. In this paper, such work will be subsumed under self-organisation. There might be slight differences in the phenomena that are described, but the basic ideas are the same, and an overview of this kind of work would be incomplete if attention was focused on only those papers that contain the term self-organisation.

Another reason why the term self-organisation is not used more frequently might be that the term itself is ill-defined. Different authors use different interpretations of the term. A selection of linguistic papers with "self-organisation" in the title (Lindblom *et al.* 1984; Wildgen 1990; Steels 1995; Ehala 1996; Demolin & Soquet 1999; de Boer 2000; Nicolis *et al.* 2000) all have a slightly different view on what it is and what role it plays. Another reason that the term self-organisation might not be popular among linguists is that it has the negative connotation of being vague.

It is therefore useful and instructive to study in some more depth the definition of self-organisation used here. Self-organisation, according to this definition is "The emergence of order on a global scale through interactions on a local scale". The definition assumes there is a system that has two main components: actors[1] and interactions. There is a population of actors, and the interactions always entail a number of actors that is considerably smaller than the total number of actors in the population. This is what is meant by interaction on a local scale. It has been mentioned above that the local and global scales, as well as the nature of the actors and interactions as well as their time scales can be interpreted in many different ways. The definition is therefore applicable to many aspects of language.

There are a number of notions that are sometimes used in conjunction with self-organising systems. These terms include the notions of chaos, bifurcations, emergence, attractors, catastrophes and (positive) feedback. A number of these terms have strict mathematical definitions and should therefore be used with care; as it is unlikely that all the conditions of the mathematical definition are fulfilled, such terms can only be used metaphorically in the context of language. This is not necessarily a bad thing, but one should not confuse the use of such terms with mathematical rigor.

Attractors, chaos, bifurcations and possibly catastrophes are rigorously defined mathematical terms that can only be used metaphorically when talking about language. In the context of language, attractors can be interpreted as stable states towards which languages tend to evolve. Language universals are sometimes considered to be attractors. Bifurcations and catastrophes can be used to describe linguistic events in which the organisation of a language—either in an individual or a population and either of a language's sound system, lexicon,

[1] By using the term "actors," it is not implied that these can necessarily determine their own actions. Actors can be neurons or even individual molecules, which behave in a completely reactive fashion.

grammar or another aspect—changes relatively suddenly. Chaos is often used to describe any unpredictable behaviour, and although it is true that languages, and social systems in general sometimes show sensitivity to initial conditions, one should be very careful in using the term chaos, as unpredictable behaviour can have other causes as well, such as complex input from outside, real randomness or complex acyclic behaviour.

The notions of emergence and positive feedback are more useful in the description of language. Emergence and emergent behaviour are generally used to describe behaviour that cannot be predicted directly from the behaviours of individual actors in a system, but that is caused by interaction between actors and/or their environment. When emergent behaviour involves many individuals and results in regular collective behaviour, the behaviour is said to be self-organising. In order for this to happen there must be positive feedback between the behaviour of individuals. Small fluctuations in individuals must be amplified and adopted by other individuals for organisation to spread through the population. Self-amplifying behaviour is said to show positive feedback.

In language, emergence and positive feedback can be illustrated by the emergence of new words for new objects. At first, many words will be coined, but the ones that are most frequently used (or most frequently used by the most prestigious speakers) will be most useful in communication and will eventually be adopted by more and more speakers, until only one word remains.

History of self-organisation in language research

Probably the first time self-organisation was mentioned in the title of a linguistic publication was in the 1984 paper “Self-organizing processes and the explanation of language universals” by Lindblom, MacNeilage and Studdert-Kennedy. Also in the early eighties, ideas very close to self-organisation were being investigated in the field of dynamic linguistics (e.g. Altmann 1985; Ballmer 1985, pp. 1–25 and references therein) and catastrophe theory (Petitot-Cocorda 1985). However, the ideas on which such work is based can be traced as far back as Zipf (1935). Zipf was possibly the first to propose within modern linguistics that linguistic structures emerge from the dynamics of language use. Of course, the term self-organisation had not been invented yet, and so was not applied to this process.

Recently, emergent phenomena and self-organisation in language have received more attention, because user-friendly computing power has become available to language researchers. This makes it possible to test the consequences of complex theories about language in a population. Lindblom *et al.*’s (1984) paper describes computer modelling of a system of syllables, and is probably the first paper to link self-organisation in language and computer modelling. Liljencrants and Lindblom’s (1972) paper already follows a very similar methodology, but it does not talk about self-organisation. However, these models consider self-organisation within the sound system of a language, not in a population of language users. The first computer models to use populations of language users are by Hurford (1987;1989) but these models consider the emergence of linguistic properties more as a result of evolution than of self-organisation. Other early work on population models of emergence of communication was reported by MacLennan (1992) and Werner and Dyer (1992) but these models focused on emergence of communication systems rather than language. Starting approximately mid-nineties, the number of articles on computer modelling of self-organising language systems increased rapidly (Batali 1994; Kirby 1994; Hashimoto and Ikegami 1995; Steels 1995; de Boer 1997). Since then a number of collections of articles have appeared, mostly in the context of evolution of language, e.g. Hurford *et al*. (1998), Knight *et al.* (2000) and Cangelosi and Parisi (2002).

An example of self-organisation in language

It is well-known that when languages use a certain place of articulation in their sound system, (for example the alveolar ridge, used to produce sounds such as /t/, /d/, /n/ etc.) stop consonants tend to occur in both voiced and unvoiced versions. Thus, if a language uses /t/, it is very likely also to use /d/. If it has /p/ it will usually also have /b/ etc. This is not universally so, however. There are languages that do not use the distinction between voiced and unvoiced stop consonants, such as many Australian aboriginal languages, for example Yidiny (Dixon 1977) or Dyirbal (Dixon 1972) . But such languages are rare.

This phenomenon can be explained using the view of language as a self-organising system in a population of language users. Whenever a language looses the distinction between voiced and unvoiced stop consonants, (such as the distinction between /t/ and /d/) there will be no more words that can be changed in meaning by changing a /t/ into a /d/ or vice versa. Then language users are in principle free to use these sounds in free variation (although this might not occur, because doing so might be considered incorrect speech behaviour). When such free variation occurs, it is more likely that a stop consonant that is sandwiched between two vowels becomes voiced, (thus the pronunciation /ada/ will be preferred over the pronunciation /ata/) while a stop consonant at the beginning or end of a word will become unvoiced (thus the pronunciation /ta/ will be preferred over /da/). This is just the consequence of what is easier to pronounce in rapid, casual speech.

Now children learning the language will be exposed to these easy-to-pronounce variants more often, and will come to prefer them in their own speech production. Thus, in actual use, the language will have words with /d/, such as /ada/ and words with /t/, such as /ta/. Still, /d/ and /t/ cannot make a difference in meaning, as there are no words whose meaning changes by only changing a /d/ into a /t/. The use of either /d/ or /t/ can be predicted from the context.

But when the language changes, and through phonetic erosion looses the first syllable of bisyllabic words, words like /ada/ will turn into /da/. However the language still has the word /ta/, and therefore a minimal pair is created, and the language will again have both /t/ and /d/ as speech sounds. As this process is much more rapid and likely than the process of loosing the distinction between voiced and unvoiced sounds, the majority of the world’s languages uses both voiced and unvoiced sounds.

This is an example of a linguistic explanation that uses actions of- and interactions between individuals to explain the structure of language as a whole and is therefore an example of self-organisation in language.

Self-organisation on different levels

As mentioned above, self-organisation in language can occur in many different forms. It can occur in the organisation of language in the brain or in a population of language users. It can be invoked to explain language universals, but also to explain how language changes and how language originated. Finally, self-organisation can occur in all aspects of language, from its sound system to meanings of words.

Self-organisation in a language user’s brain

Self-organisation in the brain is probably best known from the emergence of ocular dominance columns (see *e.g.* Erwin *et al*. 1995). However, it is quite possible that it plays a role in the way language is organised in the brain as well. Lindblom *et al*. (1984) have proposed that self-organisation causes speech sounds to be organised in a phonemic (combinatorial) way instead of in a holistic way. They optimised systems of consonant-vowel syllables for acoustic distinctiveness and articulatory ease and found that this caused a limited number (of the available possible ones) of consonants and vowels to be used in a combinatorial way. They also found co-articulatory effects that resembled those that are found in human languages.

It seems that this work has remained relatively isolated. However, recently work on computer simulations has been published (Oudeyer 2001) that investigates emergent structure and self-organisation in systems of speech sounds with a focus on what happens in an individual language user. But this work takes interactions between language users into account as well. The main problem with modelling self-organisation of language in the brain is that extremely little is known about how language is learned and stored on the neural level.

Self-organisation in a population of language users

Most of the recent work on self-organisation and emergent phenomena in language has focused on populations of language users. Such work has been on almost all aspects of language, including sound systems, grammar, lexicon formation, semantics and language change. A lot of this work focuses on evolution of language, but such work has relevance to the study of the role of self-organisation in language as well.

The basic idea behind this kind of work is that properties of language can be explained as a result of interactions between the individual language users in addition to the properties and capacities of each individual language user. In this respect this approach differs substantially from the approach to the study of language as proposed by De Saussure (reprinted, 1987) and later Chomsky (1965) where language is seen as abstract knowledge of an individual. The form in which language actually appears (called *parole* by De Saussure and *performance* by Chomsky) is seen as secondary. In the study of language as a self-organising system in a population, the use and appearance of language in daily use is seen as equally important to understanding what language is and what form it can take as the linguistic representations in the brain (see e.g. Steels 1998a). Also, the distinction between *diachronic* linguistics (the study of how language changes) and *synchronic* linguistics (the study of the grammar of an individual language as well as the study of the human capacities for language) becomes less meaningful. Traditionally, these two branches have been separate and from the point of view of certain formal theories language change was even seen as problematic. However, when one views language as a self-organising system in a population, it becomes clear that one cannot study the one without the other. Of course, most linguists have realised this as well, but because of the complexity of the interactions and mechanisms involved it has not been possible to study it properly. More recently, linguistic work has started using elements of self-organisation and dynamics in the population (see e.g. Silberman 2000).

Self-organisation of language in a population can be studied on many levels of detail. Some work focuses mainly on the interactions between transmission and learning of language (e.g. Kirby 2001), some work uses larger populations of more realistic agents (e.g. Steels 1998b) and even robots (Vogt 2000). On the other hand, some work uses simpler mathematical models in order to understand the basic mechanisms (Nicolis 2000). Other work takes the point of view that language in a population is subject to a kind of cultural evolution towards near-optimal solutions to the problem of communication under constraints (Redford *et al.* 2001).

What all these models have in common is that they view language (and therefore the universal properties of language) as the solution to the problem of communicating information and transferring the language from one generation to the next under constraints of speech production, perception and learning. They are therefore inherently dynamic: linguistic variations between individuals and populations as well as language change are considered to be inevitable outcomes. In this respect this view of language is more “complete” than the view of language as knowledge of an individual.

Of course, taking the population dynamics into account in a theory of language makes it much more complicated and makes it a lot harder to make predictions on the basis of the theory and to test these predictions. For this reason, computer models are almost universally used to investigate these models.

Self-organisation and diachronic/synchronic linguistics

Systems of self-organisation in language are inherently a synthesis of synchronic and diachronic views of language. Some research, however, has focused more on one aspect than on the other.

Diachronic work

As models of self-organisation in language are inherently dynamic, they are not only very well suited to investigate language change, but also to investigate the origins of language. Models of language as an emergent property of a population of agents that needs to communicate can be used to find answers to such questions as: under what conditions will language emerge? What form will such a language take? Will only one language emerge, or will linguistic diversity be an emergent property of the system? Most recent research that sees language as a self-organising system has looked at the origins of language (e.g. see contributions in Hurford *et al.* 1998; Knight *et al.* 2000).

There is a small body of work, however that looks mostly at the explanation of language change and linguistic diversity (Ehala 1996; Livingstone 1999; Nettle 1999). Such models investigate the questions of how language can change over time and how a homogeneous language can split into multiple distinct languages. Such questions cannot be answered, and might even appear problematic when only looking at language as the capacity of an individual. From the perspective of language as a self-organising system in a population, such questions can be investigated. Again, because of the complexity of the models, computer simulations are used.

Synchronic work

Most of the work that has looked at self-organisation in language has focussed on grammar (e.g. Batali 1994; Kirby 1994; Hashimoto & Ikegami 1995; Steels 1998b). However, all aspects of language have been investigated. The earliest work mentioning self-organisation has focused on sound systems (Lindblom *et al.* 1984) and more recently new work has appeared on this subject (Berrah *et al.* 1996; Berrah and Laboissière 1999; de Boer 1997, 2000; Demolin and Soquet 1999; Nicolis *et al*. 2000; Oudeyer 2001; Redford *et al*. 2001). Lexicon formation has received a lot of attention as well (e.g. Oliphant 1996; Steels 1995) while formation of concepts (meanings) has been a topic of investigation, too (e.g. Hurford 1989; Steels 1995; Vogt 2000). Morphology has received rather little attention, but there are some papers discussing this topic (e.g. Batali 1998; Kirby 2001).

Most of this work focuses on the role of self-organisation in a population of language users; how it causes language to become coherent, and how it causes language to show certain universal tendencies as the result of interactions between individual language users under constraints of perception, production and learning. These aspects can be investigated for each aspect of language separately, and this causes models to remain relatively simple. However, more realistic models must clearly take into account the interactions between different aspects of language, such as the influence of articulatory simplification on erosion of morphology and thus syntax. So far, little work has taken this into account, although especially Steels (e.g. 1998a; 1998b) work stresses that the interactions between the different levels of language are extremely important.

Computer modelling of self-organisation in language

As has been stressed several times, computer modelling is crucial for investigating self-organisation in language. It is only because lots of computing power have recently come to the desktop of the average researcher, that research into self-organisation in language has taken such a flight. Nevertheless, existing computers are not sufficiently powerful to model all complexities of language. It is therefore crucial for successful computer simulations to use the right kinds of simplifications. Finding such simplifications is the biggest challenge for designing a good computer simulation.

It is, however, impossible to give guidelines about which simplifications to make. It all depends on which aspect of language is to be investigated, and whether one wants to explain abstract properties of language (e.g. if compositionality is a necessary outcome of language evolution) or whether one wants to investigate a more specific question about language (e.g. why certain vowels occur more frequently than others or why more frequently occurring verbs tend to be irregular). In any case, simplifications must fulfil two criteria: they must not qualitatively change the problem under investigation and there must be a mapping from the simplified system to real language. Qualitative changes could occur, for example, if the computational complexity of a learning model is exponential in the length of utterances, or if the simulation depends on agents transferring or sharing information in an unspecified (telepathic) way.

Computer models modelling populations of language users can be constructed in different ways. It is possible to make extremely simplified, almost mathematical models in order to study general behaviour of language in a population. In such models, there is no need to model a real population, but only certain variables that describe general properties (such as the fraction of agents that have a certain linguistic trait) of the language are followed. These models are simple to build and investigate, and it is often possible to prove mathematical properties of their behaviour. However, such models can of necessity provide only a limited insight in the dynamics of self-organisation in language.

More sophisticated models model a real population of “agents” (abstract models of language users) that each have the ability to produce, perceive and learn certain aspects (i.e. sounds, words, grammatical rules etc.) of language. The way in which these agents operate and learn does not necessarily have to be modelled directly on the way the human brain works. Some models are based on neural network implementations (e.g. Batali 1998) but most of the work cited here uses higher-level, symbolic approaches. This is all right, as the focus this work is to investigate the dynamics of language in a population. One should just be careful not to equip the agents with capabilities that are biologically impossible.

Interactions between agents are implemented directly, and the system is iterated through a large number of interactions. Often this approach allows for a dynamic population: agents can be inserted into and removed from the population in order to implement realistic population dynamics. Also, spatial structure can be given to the population: agents have a spatial location, and agents that are closer together have a higher probability of interacting than agents further apart. Furthermore, the world in which the agents operate and about which they communicate can be made more complicated. Many models have only a fixed number of meanings, or only abstract meanings about which the agents can communicate, but other models have the agents communicate about a more elaborately modelled world, or even about the real world as observed through video cameras (Steels 1998b).

Another way of implementing populations in a computer model is in the form of a genetic algorithm (Batali 1994; Hashimoto & Ikegami 1995; Zuidema & Hogeweg 2000). Such a model works like biological evolution in nature. Individuals get assigned a fitness on the basis of their behaviour and their likelihood of creating offspring is proportional to their fitness. Offspring will be like their parents, but there is a possibility of mutation and random recombination of genetic information from the parents. Fitness in linguistic models can for example be evaluated by the ability of the individuals to parse sentences, to produce or perceive sounds or to communicate about the environment. Although self-organisation can occur in such models, and although such work is closely related to research into emergent behaviour and self-organisation, in general there is an element of global control in such evolutionary models. Unless the function by which individuals are evaluated also evolves (this would be co-evolution) it has a global influence that falls outside the definition of a self-organising system. However, the results of these simulations do have relevance for the study of language as a self-organising system, and so they cannot be excluded in this discussion.

The most sophisticated models combine evolution and self-organisation in a population. The individuals are evaluated by how well they learn the language, but as the language changes because of self-organisation in the population, the evaluation function is no longer static, but changes with the changing system. However, such systems would be computationally extremely complex, and it appears that so far only plans exist to build such systems, but none have been realised.

An example of self-organisation

The example that will be presented here is from my own work on self-organisation in vowel systems, and has been published more completely in de Boer (1997, 2000, 2001). It fits in the tradition of investigating vowel systems with computer models (*e.g.* Liljencrants & Lindblom 1972; Schwartz *et al*. 1997b) and is more directly based on two artificial life models of vowel system transfer (Glotin and Laboissièrre 1996; Berrah *et al*. 1996). At its basis lays the observation that vowel systems in human languages show remarkable regularities (*e.g.* Crothers 1978; Schwartz *et al*. 1997a).

Although humans can produce and distinguish many different vowel sounds, human languages tend to use only a limited subset of these. At least 45 different qualities are recognised by phoneticians (Ladefoged & Maddieson, 1996, ch. 9) and languages can make use of up to at least 15 of these at the same time (Maddieson, 1984; Schwartz *et al*. 1997a). Some languages are said to have more vowels, but they make use of secondary distinctions, such as vowel length, nasalisation etc. However, in the large majority of human languages, vowel systems tend to be quite small and tend to be constructed among relatively simple rules. The most frequently occurring number of vowels in human languages is five, and these vowels are almost always /i/, /e/, /a/, /o/ and /u/. An example of a language with such a vowel system is Spanish. Moreover, most of the world's languages contain the vowels /i/, /a/ and /u/, and there is a small repertoire of (about eight to ten) vowels that accounts for most of the vowel sounds occurring in the world's languages. Only when vowel repertoires become really large are more exotic speech sounds used, such as the rounded front vowels /y/, /ø/ and /œ/ that are found in French, German and the Scandinavian languages. Coincidentally, Western and Northern Europe are areas where languages with extraordinarily many vowel sounds are spoken, and English is one of these. In this respect English is quite unusual.

But there is yet another remarkable property of human vowel systems. They tend to be symmetrical. If a language contains, for example the vowel /o/ it is more likely than expected on the ground of the a-priori probability, to also contain the vowel /e/. This vowel corresponds in tongue height, but is articulated in the front part of the mouth, instead of in the back. Similar symmetries are found for many different pairs of vowels.

Traditionally (Jakobson & Halle 1956), and within the tradition of generative grammar (Chomsky & Halle 1968), these regularities have been explained as the result of innate properties of the human brain, and more specifically the human language faculty. But one can ask whether it is really necessary to postulate innate feature detectors and preferences to explain the regularities. Liljencrants and Lindblom (1972) have shown that one would expect to find such regularities in systems that are optimised for acoustic distance between the individual vowels. They built a computer simulation that optimised the distance between a given number of points representing vowels in an acoustic space. From this computer simulation emerged vowel systems that resembled the most frequent systems found in human languages. Subsequent refinements of this model (e.g. Schwartz *et al*. 1997b) have shown even better correspondence with human language data.

The question remains, however, how vowel systems become optimised. No individual language learner optimises. In fact, infants are very good at learning the subtlest distinctions in vowel quality. Also, there are languages with vowel systems that are far from optimal, but these are learned just fine. Apparently the cause of optimisation is not to be found in the individual infant's learning behaviour.

Within the framework of self-organisation, an alternative hypothesis can be formulated. Perhaps the optimisation is the result of repeated interactions between agents that learn and use vowels under constraints of speech production, perception and learning? This hypothesis has been investigated using a computer model (de Boer 1997, 2000, 2001). This computer model can only be described very briefly here, for details the reader is referred to the original references.

The model is based on a population of agents that can each produce, perceive and learn vowels in a realistic way. Each agent has a simple speech synthesiser that can generate all basic vowels, based on three inputs: the tongue height, tongue position and lip rounding needed to articulate the vowel. Thus for an /i/, the tongue needs to be high, to the front and the lips need to be spread, while for an /o/, the tongue needs to be somewhere between high and low, but to the back of the mouth, and the lips need to be rounded. These three parameters are sufficient for generating all basic vowel qualities. Vowels are stored as "prototypes". For each vowel an agent knows, a point in both acoustic and articulatory space is stored that is most representative of that particular vowel. Perception is based on a cognitively plausible distance function that is based on properties of the sound spectrum of the vowels. For a given signal, its distance is calculated to all acoustic prototypes, and the one with the shortest distance is defined to be the vowel recognised. Agents start out with empty vowel repertoires, and add and remove vowels on the basis of the interactions with other agents. In order to get the games started, and in order to put pressure on the agents to extend their repertoires, agents can add random new vowels with low probability (1% per game).

The agents interact in so-called *imitation games*. In each imitation game, two agents are chosen randomly from the population. One agent chooses a random vowel from its repertoire and produces this, while adding noise. The other agent analyses this sound in terms of its prototypes, and picks the one that is closest to the signal. It then produces the corresponding sound in turn, also adding noise. The first agent then analyses this sound in terms of its prototypes, and checks whether the prototype it recognises is the one it originally used for producing the sound. If this is the case, the game is said to be successful. If not, it is a failure. This is communicated to the other agents through "non-verbal feedback". Of course, in reality children do not get feedback (non-verbal or other) about the sounds they produce. However, one could imagine infants to derive information from things such as emotional state of the caretaker or failure to achieve a certain communicative goal.

As a reaction on the imitation game, agents update their vowel system. They can add a vowel if necessary—this happens especially often in the beginning of the game when agent's repertoires are empty and new vowels are added as close approximations of heard signals. Also, agents sometimes add random vowels, in order to create pressure to increase the size of their repertoires or to get imitation started when an agent's repertoire is empty and it has to produce a sound nevertheless. Vowels can also be discarded if it turns out that they are not successful for imitating other agents' vowels. This is evaluated on the basis of their past suc-

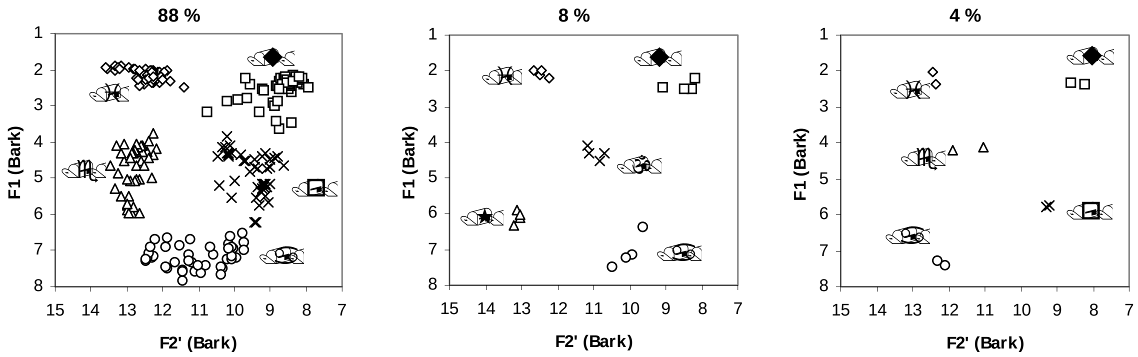

Figure 1: Emergent five-vowel systems, classified for configuration type. The acoustic space is based on the frequencies of the first two main peaks in the vowel's frequency spectrum (called F_1 and F_2') and these are plotted in the logarithmic 'Bark' frequency scale. Distance in the Bark frequency scale corresponds better to perceptual distance than in the Hertz frequency scale. The plots are structured in such a way that vowels appear in positions in which phoneticians usually plot them: high front vowels to the upper left, low back vowels to the lower right.

cess or failure in imitation games. Vowels can also be merged if they come too close together in either acoustic or articulatory space. Finally, agents can shift vowels in their repertoire over a small distance in order to approximate more closely the signal heard in the imitation game. These updates to an agent's vowel repertoire are not directly inspired by the way children learn. However, they are biologically plausible in principle, as they only involve simple manipulations and local information. Human infants probably use more powerful statistical learning techniques, but for simplicity's sake, a learning model that was as simple and transparent as possible was used.

All these actions and interactions lead to the emergence of realistic vowel systems. It turns out that these vowel systems are not only remarkably like the vowel systems found in human languages, but that the frequency with which different types of vowel systems occur agrees remarkably well with the frequency in which they occur in human languages. An example of this is given in figure 1. In this figure, five-vowel systems that emerged are classified according to the configuration of vowels in the systems. In 49 out of 100 populations that had the same parameter settings, five-vowel systems emerged (four- and six-vowel systems also emerged). Systems with similar configurations (determined on the basis of the number of front-, back- and central vowels, as well as their symmetry) are grouped together in the same figure. Each vowel system emerged in a population of twenty agents. From each of the populations, one agent was selected randomly, and this agent was used to represent that population's vowel system. The vowel systems of all agents with the same configuration were plotted together in acoustic space. Thus, each data point per cluster belongs to a different agent in a different population, while for every agent five data points (each indicated with a different symbol) appear in the figure.

It can be observed that the frequencies of the emerged vowel systems correspond quite well with those found in human languages. In Schwartz *et al.*'s (1997a) survey, 89% of the vowel systems were like the ones in the leftmost frame, and both other types of system occurred in 5% of the languages in their sample.

As has been defined above, self-organisation is the emergence of global order through local interactions. The local interactions in this simulation consist of imitation games. In each game only two agents participate, and only local updates are made to an agent's vowel system (no global optimisation is performed). The emerging global order consists of concentrated clusters of vowel prototypes that are shared in a population of agents. The possible vowel systems that emerge are restricted to a reasonably small class of realistic vowel systems. Note that the emerged vowel systems are not necessarily completely optimal. The fact that the agents need to have a vowel system that is similar to the systems used by the other agents in the population may cause the system to stay stuck in a local optimum. This is reminiscent of the situation in human languages, where sub-optimal vowel systems also occur from time to time. Here too, a vowel system can become sub-optimal due to the history of the language. However, single language users cannot change the system too much towards optimality, as this would make him or her incomprehensible for the other speakers. On the basis of these and similar results, it can be concluded that self-organisation in a population under constraints of perception, production and learning can explain the universal tendencies of human vowel systems; no specific innate tendencies are necessary.

Conclusion

This paper has presented an overview of work on language as a self-organising system. It has been shown that viewing language in this way is extremely useful. Self-organisation is the emergence of global order through interactions on a local scale. It can happen (and has been investigated) in the human brain, but most work into self-organisation in language has focused on linguistic phenomena in a population of language users.

Self-organisation provides a means by which diachronic linguistics (the description of how language changes) can be unified with synchronic linguistics (the description of grammars of human languages and the capacities that humans bring to bear on the tasks of learning and understanding language). Self-organisation can be used to gain insight in such diverse aspects of language as phonological universals of sound systems, the emergence of grammar, linguistic change or the way a population of language users adopts new words. All these approaches have in common that they view language as a dynamic system in which interactions between language users is as important as the knowledge and capacities of those language users. In this respect models of self-organisation attach equal importance to both De Saussure's *langue* and *parole*, and Chomsky's *performance* and *competence*.

As theories of self-organisation in language involve extremely complex dynamics, most of the work uses computer simulations. However, this reliance on the computer makes it more difficult for traditional linguists to understand and contribute to this work. Perhaps because of this, a lot of work on language and self-organisation has been published in non-linguistic journals and conference proceedings.

Much work has already been done on language as a self-organising system, and this area of study is rapidly expanding. However, a lot of extremely interesting experiments still remains to be done. An important task for researchers of self-organisation in language remains to make their work accessible to the large community of linguists that are less literate in computers.

Acknowledgement

This article was written at the AI-lab of the Vrije Universiteit Brussel. The author wishes to thank Willem Zuidema for suggestions about literature.

References

Altmann, G. (1985) On the Dynamic Approach to Language, In: Ballmer (1985) pp. 181–189

Ballmer, Th. T. (1985) *Linguistic Dynamics*, Berlin: Walter de Gruyter.

Batali, J. (1994) Innate biases and critical periods: combining evolution and learning in the acquisition of syntax. In R. Brooks & P. Maes (eds.) *Proceedings of the Fourth Artificial Life Workshop*, Cambridge (MA): MIT Press, pp. 160–171

Batali, J. (1998) Computational simulations of the emergence of grammar, In: J. R. Hurford, M. Studdert-Kennedy & C. Knight (eds.) *Approaches to the evolution of language*, Cambridge: Cambridge University Press. pp. 405–426

Berrah, Ahmed-Réda, Hervé Glotin, Rafael Laboissière, Pierre Bessière & Louis-Jean Boë (1996) From Form to Formation of Phonetic Structures: An evolutionary computing perspective. In Terry Fogarty & Gilles Venturini (eds.) *ICML '96 workshop on Evolutionary Computing and Machine Learning*, Bari 1996, pp. 23–29.

Berrah, A.-R. & Laboissière, R. (1999) SPECIES : An evolutionary model for the emergence of phonetic structures in an artificial society of speech agents, In: Floreano, D., Nicoud, J.-D. & Mondada, F. (Eds.), *Advances in Artificial Life, Lecture Notes in Artificial Intelligence*, Volume **1674**, Springer, 1999 pp. 674–678.

Cangelosi, A. & Parisi, D. (2002) *Simulating the evolution of language*, London : Springer

Chomsky, Noam (1965) *Aspects of the Theory of Syntax*, Cambridge (MA): MIT Press.

Chomsky, Noam & Morris Halle (1968) *The sound pattern of English*, Cambridge (MA): MIT Press.

Crothers, John (1978) Typology and Universals of Vowel systems. In Joseph H. Greenberg, Charles A. Ferguson & Edith A. Moravcsik (eds.) *Universals of Human Language*, Volume 2 Phonology, Stanford: Stanford University Press pp. 93–152.

de Boer, B. (1997) Generating vowels in a population of agents. In P. Husbands & I. Harvey (eds.) *Proceedings of the Fourth European Conference on Artificial Life*, MIT Press, pp. 503-510

de Boer, B. (2000) Self organization in vowel systems, *Journal of Phonetics* **28** (4), pp. 441–465

de Boer, B. (2001) The origins of vowel systems, Oxford: Oxford University Press.

Demolin, D. & Soquet, A. (1999) The Role of Self-Organization in the Emergence of Phonological Systems. *Evolution of Communication* **3**(1) pp. 21–48

De Saussure, F. (1987*) Cours de linguistique générale, édition préparée par Tullio de Mauro*, Paris: Payot.

Dixon, R. M. W. 1972. *The Dyirbal Language of North Queensland*. Cambridge: Cambridge University Press.

Dixon, R. M. W. 1977. *A Grammar of Yidiny*. Cambridge: Cambridge University Press

Ehala, M. (1996) Self-organisation and language change, *Diachronica* **13** pp. 1–28

Erwin E., K. Obermayer, and K. Schulten (1995). Models of Orientation and Ocular Dominance Columns in the Visual Cortex: A Critical Comparison. *Neural Computation* **7**, 425-468.

Glotin, Hervé & Rafaël Laboissière (1996) Emergence du code phonétique dans une societe de robots parlants. *Actes de la Conférence de Rochebrune 1996 : du Collectif au social*, Ecole Nationale Supérieure des Télécommunications – Paris.

Hashimoto, T. & Ikegami, T. (1995) Evolution of Symbolic Grammar Systems In: F. Morán, A. Moreno, J. J. Merelo & P. Chacón (eds.) *Advances in Artificial Life, Lecture Notes in Artificial Intelligence* **929**, Berlin: Springer pp. 812–823

Hurford, J. R. (1987) *Language and number*. Oxford: Blackwell.

Hurford, J. R. (1989) Biological evolution of the Sausurean sigan as a component of the language acquisition device, *Lingua* **77** pp. 187–222

Hurford, J. R., Studdert-Kennedy, M. & Knight, C. (1998) *Approaches to the evolution of language*, Cambridge: Cambridge University Press.

Jakobson, Roman & Morris Halle (1956) *Fundamentals of Language*, the Hague: Mouton & Co.

Kirby, S. (1994) Adaptive explanations for language universals: a model of Hawkins' performance theory. *Sprachtypologie und Universalienforschung* **47**, pp. 186–210

Kirby, Simon (2001) Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation* **5**(2) pp. 102-110

Knight, C., Studdert-Kennedy, M. & Hurford, J. R. (2000) *The evolutionary emergence of language*, Cambridge: Cambridge University Press

Ladefoged, Peter & Ian Maddieson (1996) *The Sounds of the World's Languages*, Oxford: Blackwell.

Liljencrants, L. & Björn Lindblom (1972) Numerical simulations of vowel quality systems: The role of perceptual contrast. *Language* **48** pp. 839–862.

Lindblom, B., MacNeilage, P. & Studdert-Kennedy, M. (1984) Self-organizing processes and the explanation of language universals, In: B. Butterworth, B. Comrie & Ö. Dahl (eds.) Explanations for Language Universals, Berlin: Walter de Gruyter & Co. pp. 181–203

Livingstone, D. & Fyfe, C.Modelling the evolution of linguistic diversity. In: D. Floreano, J.-D. Nicoud & F. Mondada (eds.) *Advances in artificial life. Lecture notes in artificial intelligence* **1674**, Berlin: Springer, pp. 704–708

MacLennan, B. (1992) synthetic Ecology: An approach to the study of communication, in: C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen (eds.) *Artificial Life II*, Redwood city (CA): Addison-Wesley Publishing Company pp. 631–658

Maddieson, Ian (1984) *Patterns of sounds*, Cambridge University Press.

Nettle, D. (1999) *Linguistic Diversity*, Oxford: Oxford University Press.

Nicolis, S. C., Deneubourg, J. L., Soquet, A., & Demolin. D. (2000). Fluctuation induced self- organization of a phonological system *International Conference on Complex Systems*, Nashua.

Oliphant, M. (1996) The dilemma of Saussurean Communication, *Biosystems* **37**, pp. 31–38

Oudeyer, P-Y (2001) Coupled Neural Maps for the Origins of Vowel Systems, In: G. Dorffner, H. Bischof & K. Hornol (eds.) *Proceedings of the International Conference on Artificial Neural Networks 2001, Lecture Notes in Computer Science* **2130**, pp. 1171–1176

Petitot-Cocorda, J. (1985) *Les catastrophes de la parole*. Paris: Maloine.

Prigogine, I., & Stengers, I. (1988) *Order out of Chaos: Man's New Dialogue with Nature*, London: Fontana Paperbacks

Redford, Melissa A., Chun Chi Chen and Risto Miikulainen (2001) Constrained Emergence of Universals and Variation in Syllable Systems, *Language and Speech*, **44** (1) pp. 27–56

Silverman, D. (2000) Hypotheses concerning the phonetic and functional origins of tone displacement in Zulu, *Studies in African Linguistics* **29**(2)

Schwartz, J. L., Boë, L. J., Vallée N. & Abry, C. (1997a), Major trends in vowel system inventories. *Journal of Phonetics* **25**, pp. 233–253

Schwartz, Jean-Luc, Louis-Jean Boë, Nathalie Vallée & Christian Abry (1997b), The Dispersion-Focalization Theory of vowel systems. *Journal of Phonetics* **25**, pp. 255-286.

Steels, L. (1995) A Self-Organizing Spatial Vocabulary, *Artificial Life* **2**(3) pp. 319–332

Steels, Luc (1998a) Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation. In James R. Hurford, Michael Studdert-Kennedy & Chris Knight (eds.) *Approaches to the Evolution of Language*, Cambridge: Cambridge University Press pp. 384–404.

Steels, Luc (1998b) The origins of syntax in visually grounded robotic agents, *Artificial Intelligence* **103**, pp. 1–24

Vogt, P. (2000) Bootstrapping grounded symbols by minimal autonomous robots. *Evolution of Communication* **4**(1) pp. 89-118

Werner, G. M. & Dyer, M. G. (1992) Evolution of communication in artificial organisms, in: C. g. Langton, C. Taylor, J. D. Farmer & S. Rasmussen (eds.) *Artificial Life II*, Redwood city (CA): Addison-Wesley Publishing Company pp. 659–687

Wildgen, W. (1990) Basic principles of self-organisation in language, In: H. Haken & M. Stadler (eds.) *Synergetics of Cognition*. Berlin: Springer, pp. 415–426.

Zipf, G. K. (1935) *The psychobiology of language*, Cambridge (MS): MIT Press

Zuidema, W. & Hogeweg, P. (2000) Selective advantages of syntactic language - a model study, in: Gleitman and Joshi (eds.), *Proceedings of the 22nd Annual Meeting of the Cognitive Science Society*, Lawrence Erlbaum Associates, Hillsdale, USA, pp. 577 - 582