Language acquisition age effects and their role in the preservation and change of communication systems

Verhoef, T.; de Boer, B.

Link to publication

Language acquisition age effects and their role in the preservation and change of communication systems[*]

Tessa Verhoef

Amsterdam Center for Language and Communication, University of Amsterdam

Bart de Boer

Amsterdam Center for Language and Communication, University of Amsterdam

In this paper it is investigated what happens in a culturally evolving communication system if children learn faster than adults. An agent based computer model is investigated which contains a population of agents in which vowel systems emerge when individuals play imitation games. These vowel systems are found to be preserved better in the case that there is an age structure in the population. When there is a difference in learning ability between children and adults, the vowel systems are more stable, the mutual intelligibility stays higher and a higher complexity can be maintained. Recent advances in sociolinguistics indicate that changes in language learning ability while growing up can account for the different ways that linguistic changes spread from community to community. They found that the structure of the system is better preserved when speakers of all ages are involved in language contact as opposed to only adult to adult contact. Our findings therefore seem to be in line with observations in real language evolution and we provide additional evidence for the fact that it is important to consider the consequences of language acquisition age effects in the study of language preservation and change.

1   Introduction

Language learning is probably one of the few tasks that children are better at than adults. Adults, who have full-grown cognitive abilities and many years of experience in acquiring all sorts of knowledge and tasks, have the hardest time learning a new language. Children however, who are cognitively still underdeveloped, lack detailed motor control and are not deliberately trained, learn it perfectly. They pick up every aspect of the language without apparent effort. This phenomenon is often referred to as the 'critical period' for language

acquisition (Lenneberg, 1969), or a period of 'heightened sensitivity' (Birdsong, 2005) that gradually declines. In this paper it is investigated what the consequences are for a culturally evolving communication system if children learn faster than adults.

Time frames of heightened sensitivity to environmental input are not uncommon in nature. One well-known example is that of imprinting. Geese, chickens and ducks, soon after they leave the shell of their egg, will follow and keep following the first moving object or animal they perceive. The newborn gets imprinted with this object or animal and will continue to follow it as if it is following its parent (Lorenz, 1937; Hess, 1958; Spalding, 1873). There is an exception though. If the chicks and ducklings are not exposed to any moving object in the first 25 to 30 hours after their birth, this imprinting mechanism discontinues to work and the newborns miss out on their chance to imprint onto their mother (Hess, 1958). Likewise, rhesus monkeys that do not come in contact with other animals in their first year of life will be unable to develop normal social monkey behaviors (Harlow et al., 1965). Several bird species have only limited time to acquire their species specific song (Doupe and Kuhl, 1999) and visual systems of cats need to be exposed to external stimuli and developed within a certain time frame after birth (Hubel and Wiesel, 1970) because their window of opportunity closes.

These effects of age sensitivity are useful in the development of these animals. Sensitive periods have been identified in a more general context as being the result of evolution: "In many species and for different functional systems, strong selection pressures that favor great sensitivity to certain environmental stimuli and a maximum of learning during early stages of life may favor great stability of the results of such early experience, providing some kind or degree of protection against some later possible influences on the individual." (Immelmann and Suomi, 1981). Therefore there might be advantages to a sudden or gradual change in learning abilities for humans as well, for instance in language. Just like it is crucial for the survival of a duck that it knows soon after birth who its mother is, language might be so important for humans that early learning of it is crucial to be able to function in our complex world of social structures.

One important basic design problem that is often encountered in the creation of autonomous machines that can cope with the complexity of the real world, is the stability-plasticity dilemma (Carpenter and Grossberg, 1988): intelligent systems should be plastic enough to be able to adapt quickly to new circumstances and learn new things, while at the same time they should be stable enough to retain the knowledge and skills they already acquired. In language acquisition this dilemma is of equal importance. On the one hand language users should be able to learn a new language when moving to a different community,

but on the other hand they need to keep the ability to identify and communicate with people from their own community as well.

Studying the exact causes and consequences of age sensitivity in language acquisition is not trivial. Language, and the mechanisms for acquiring it are shaped by processes on three different time scales (Kirby and Hurford, 2002): learning, cultural evolution and biological evolution. As Steels (1997b) has pointed out, extreme complexity arises in the interaction between these time scales. As has already been shown by Kirby and Hurford (2002), Steels (1997b) and others, computer models can deal with this complexity, and have already been applied successfully in the study of age sensitivity in language acquisition. Simulations allow researchers to explore hypothetical situations that would be impossible to achieve in real language environments, as these cannot be manipulated experimentally. With computer models it is feasible to explore multiple potential scenarios, a method that can compensate to a certain extent for the information that is inevitably lost in historic processes.

In explaining the cause of the existence of age sensitivity in language acquisition, evolutionary simulations have provided insights. Hurford (1991) created a model in which it is assumed that a selective advantage exists for linguistic abilities. The population goes through many generations of selection and mutation and a critical period effect evolves, caused by the selective pressures on language learning. In Hurford's view, the language acquisition capacity evolves as an adaptation, but the critical period arises as a side effect of "a lack of selective pressure to acquire (more) language (or to acquire it again) once it has been acquired" (Hurford, 1991). In a related model, Hurford and Kirby (1999) show that complex interactions between speed of language acquisition as a biological property and language size as a cultural property are important. The acquisition speed in this model evolves in such a way that full language acquisition is completed before puberty, the age at which agents start to reproduce. The speed will not evolve to become even greater once the complete language of the population can be acquired before puberty. When the language can be fully acquired before puberty, there is no pressure favoring faster acquisition than necessary to be 'ready' before it is needed for reproduction. When an 'innovative potential' is introduced, which gives the agents the ability to expand the language when they have acquired the complete existing language early, the language can grow and then again more speed is needed to acquire it on time, resulting in an arms race with ever increasing size and speed. Among many others, one more example concerns a more mathematical evolutionary model that was investigated by Komarova and Nowak (2001). In their model a critical period exists as an evolutionary stable strategy (Nash equilibrium). Their model assumes that there is a cost to language acquisition, which has a negative effect on the reproductive success

while at the same time a poor language performance reduces fitness as well. The interaction between these two forces yields an evolutionary stable optimal learning period. These examples all show that computer models can be very useful for the discovery of new ideas about language and age structure dynamics.

In the present study, a new simulation is introduced in which the focus is not, contrary to the contributions mentioned above, on what caused the age structure to biologically evolve, but on what it entails on a population scale. It is not our goal to find evolutionary explanations for decreasing language learning abilities but we assume that it exists and investigate what the presence of such an age structure entails for a culturally evolving communication system. We use the model of de Boer and Vogt (1999) and de Boer (2000), because this model is very suitable for the investigation of the interaction between individual level behaviors and consequences at the population level. We argue that the age structure improves stabilization and preservation of complexity in the shared communication system of a changing population. These results are then compared to recent findings in sociolinguistics concerning the role of differences of learning ability in language change.

2   Age sensitivity in language acquisition

Lenneberg (1969) initiated modern research into age effects for language acquisition. He observed that traumatic brain lesions caused permanent aphasia in adults, while children with similar lesions initially had the same problems but were able to recover. Other important observations in this research are cases of language deprivation, such as the well-known case of Genie (Curtiss et al., 1974). When Genie started to acquire her first language, after being raised in almost complete social isolation, she was almost fourteen years old. With a lot of training researchers were able to teach Genie a small vocabulary but her language use stayed far from normal (Meadow, 1978). The case of Genie caused controversy because her situation was so extreme that it is actually hard to determine whether her abnormal linguistic development was due to linguistic stimulus deprivation or to the adverse conditions in which she grew up. Similar problems are connected to the observations that Lenneberg (1969) reported about aphasics. These are all abnormal conditions that make it difficult to draw conclusions about the normal language acquisition pattern.

Cases of delayed sign language exposure form more natural examples (Mayberry, 2007). Newport (1988) describes differences between adult and infant learning of American Sign Language (ASL). The congenitally deaf subjects in this research were raised under normal circumstances. In contrast to the acquisition of spoken language, most ASL learners do not start to learn their

language when they are born but most of the time they are exposed to the sign language for the first time when they start to go to school. However, in the cases where deaf children acquire their language from parents who are also deaf and communicate with ASL, they do get early exposure. Comparing the acquisition of ASL for early and late exposure and the level of ultimate fluency that is achieved, Newport (1988) found age effects. Mayberry (2007) investigated the influence of first (sign) language acquisition on the ability to learn a second language and also found an age of acquisition effect, showing that early exposure to a first language also benefits second language acquisition.

In addition to the observations of development in first language acquisition, age effects are being studied in second language learners. Second language learners and proficient bilinguals provide a source that teach us about the interaction between different languages (Flege et al., 1995), the influence of age of first language acquisition on second language learning (Mayberry, 1993, 2007), the ultimate proficiency that can be reached in the second language (Johnson and Newport, 1989; Oyama, 1976; Flege et al., 1995) and which aspects of the language are the hardest to acquire at an older age (Weber Fox and Neville, 1996; Singleton, 2002).

The exact mechanism behind language acquisition age effects is still unknown. Many different driving forces have been proposed. Some of them assume that the acquisition time frame is biologically determined, for instance coinciding with a decline of brain plasticity during development (Lenneberg, 1969). Brain plasticity has been found to relate to hormonal influences and the changes that occur around the age of puberty (Yun et al., 2004). The link between hormones and acquisition is also found in research done with songbirds. As cited by Doupe and Kuhl (1999), experiments have been conducted in which birds were castrated before they had learned their song. The hormonal changes that accompanied this procedure influence their singing behavior and by the time they would normally learn their song, these birds did not. However, when they would at a later moment receive certain hormones, they were still able to learn their song. This shows that in some cases the special learning ability seems to be extendable. According to the 'less is more' hypothesis (Newport, 1988), the mechanism behind age effects is not driven by special abilities but by cognitive inabilities. Newport suggests that children perform better because they are cognitively more limited when the acquisition takes place. The adult learners are able to store more whole forms and meanings, and may therefore face a more difficult analytic task. So, children have an analytical advantage. Other theories have been proposed that do not assume this direct link with developmental patterns. These theories assume that what we observe as age effects in language acquisition actually has more to do with the quality and quantity of linguistic input. According to Kuhl (2000), the

mechanism involves 'neural commitment', which means that a mental map is being created while learning, adjusting neural structures to the sounds of the native language. Perception is guided by a 'native language magnet'. Newborns are able to distinguish sounds of many different languages, but when they grow up, their perceptual behavior changes so that they are better able to distinguish the sounds of their native language while other, for them irrelevant, contrasts are no longer perceived. Once their neural structure has committed to a language, other languages become harder to learn. Flege (1999) poses a similar hypothesis, which also emphasizes how a first and second language can influence each other. With his 'interaction hypothesis', he proposes a negative correlation between first language and second language proficiency.

3   Age effects in language emergence, change and growth

Several lines of research have shown that children and adults, with their differences in language learning ability, might be involved in language emergence, change and growth in different ways. When people who do not speak a mutually intelligible language come into contact and there is a need for communication, a pidgin language might emerge (Hall, 1962; Sankoff and Laberge, 1980), which contains features of both original languages, but does not contain a lot of structure. When the contact situation is over, the pidgin might disappear again (Hall, 1962) or else it might be transmitted to future generations and children will learn it. In this last situation, it has been argued that native learning might stimulate the emergence of structure (Sankoff and Laberge, 1980). A newly emerging sign language in Nicaragua, that did not originate from language contact but spontaneously emerged when a new school brought together deaf people, showed a comparable pattern in which children played an important role in the formation of the language (Senghas et al., 2004).

   A recent finding by Labov (2007) indicates that the differences in learning ability between adults and children could explain different observations in language change and stability. Two models for explaining linguistic change have dominated for a long time: the 'family tree model' and the 'wave model' (Labov, 2007). The family tree model describes how languages in the world are related and how proto-languages branched into new groups of related languages. One limitation of this model has been said to be that it assumes separate branches are independent. The wave model was introduced as an alternative view in 1872 (as described by Fox (1995)(p.129)) and accounts for the spreading of changes in the case of language contact, across language boundaries. Labov (2007) unifies these two models in one theory in which he accounts for both faithful 'transmission' of changes from generation to generation, resulting in family tree like pattern and the 'diffusion' of changes

through language contact, resulting in a wave like pattern. The difference in learning abilities of young and old language users plays a major role in this theory. Labov (2007) presents a detailed description of two comparisons of sound system development: the first compares the stability of the short-a system within New York City with the spreading of this system to four other cities and the second compares the Northern Cities Shift in the areas of the Inland North with the spreading of this system to the Midland cities. The spreading of the complex short-a system of New York City to other cities had resulted in a loss of structure and regularity, while it had been and remained stable within New York City. According to Labov, this difference is due to the fact that in New York City, there was an unbroken chain of faithful transmission from adults to children, while it was mainly individual adults who diffused the system to other cities. Likewise, the Northern cities shift resulted in a uniform and stable transmission of the system in the Inland North, probably because this case involved the migration of entire communities consisting of families with both adults and children. The changes were therefore steadily transmitted from generation to generation. At the same time, in the Midland cities the result was less stable and involved again a contact situation with single adults traveling around. Labov (2007) proposes therefore that the difference in adult and child learning causes the difference between transmission and diffusion.

The results of the simulations described below support this finding by Labov (2007). Here, it is also investigated how age effects in learning influence the cultural evolution of a communication system and the results show that an age structure in the population helps stabilization and preservation of structure in emerged vowel systems. We therefore provide additional support for the proposal that it is important to consider the consequences of language acquisition age effects in the study of language preservation and change.

4   Vowel systems in a population of agents

The starting point of this work is the agent-based imitation game model as described by de Boer (1997, 2000). This model follows the language game paradigm, which focuses on relatively large population sizes and communication between individuals of the same generation (de Boer, 2006). Important in this work is the idea of viewing language as a population level phenomenon in which self-organization causes optimization and coherence. In de Boer and Vogt (1999), a version of this model is used which integrates transmission across multiple generations of interacting agents. This allows for the introduction of an age structure, which has also been applied in de Boer and Vogt (1999). Both models are re-implemented and used to investigate the influence of the age structure in detail.

The model consists of a population of individuals that interact with each other by means of imitation games (de Boer, 2000). The agents are equipped with a memory in which they can store articulatory (and acoustic) information about the vowels they have learned. With the use of a realistic articulatory synthesizer they can produce these vowel sounds and they are able to categorize the sounds that others produce with a realistic model of vowel perception. In response to their interactions with other agents they can update their memory and learn new sounds.

4.1   Memory

Following de Boer (2000), the agents have a vowel memory in which they store prototypes of the vowels they have learned. Prototypes are the centers of the vowel categories that the agent has learned. A vowel prototype is described by three properties: tongue position (p), tongue height (h) and lip rounding (r). They have continuous values that can vary between 0 and 1. For each prototype the agent also keeps track of the number of times it has been used and the number of times it has been used successfully.

4.2   Production

Values for position, height and rounding form the input for the articulatory synthesizer. The output of the synthesizer consists of the frequencies of the first four formants of the vowel in Hertz. They are computed with a set of equations that are described by de Boer (2000). Thus the agents can produce a realistic range of sounds. Some noise is added to these formant frequencies to account for the fact that there is variation in the speaker's production of the same vowel. For every ith formant F_i, the frequency F_i' with added acoustic noise is computed by $F_i' = F_i (1 + v_i)$, where v_i is randomly chosen from a uniform distribution such that $-\psi/2 \leq v_i < \psi/2$, where ψ is the maximal noise allowed.

4.3   Perception

The agents perceive a sound as the nearest vowel prototype in their memory. This implements categorical perception. To do this, they need to determine the distance between two vowels in acoustic space. This is done in a two-dimensional space: the first dimension is the first formant and the second is the effective second formant (both on a Bark frequency scale). The effective second formant is computed from the three higher formants (Schroeder et al., 1979). A graphical representation of this acoustic space is shown in figure 1.

Figure 1: Graphical representation of the acoustic space (the blue area), with a repertoire of vowel prototypes (the colored dots).

To compute the effective second formant F_2', the formant values are first converted from the Hertz to the Bark scale using equation (1) (Schwartz et al., 1997).

$$(1) \quad F_{Bark} = 7\sinh^{-1}(F/650)$$

Then F_2' is computed for the two vowels (vowel a and vowel b) that are compared, following the equations in de Boer (2000) and these two values are used to compute a weighted Euclidian distance D, like in de Boer (2000) and as shown in the following equation:

$$(2) \quad D = \sqrt{(F_1^a - F_1^b)^2 + 0.3(F_2'^a - F_2'^b)^2}$$

4.4 Interactions

The agents in the simulation interact with each other by playing imitation games (de Boer, 2000). In every game, two agents are randomly selected from the population. One of them is the initiator and starts the interaction by selecting a random vowel from its prototype repertoire and producing this sound. The other agent is the imitator and perceives the vowel that the initiator has just produced: it finds its prototype that is closest to the sound the initiator produced. Next, the imitator produces the selected prototype and the initiator perceives this sound. The initiator then determines which prototype in its memory is closest to this sound. If this prototype is the same as the one that the initiator initially produced, the game is a success, if not, it is a failure. This information is communicated to the imitator non-verbally and they both update their memory in response to the game. When an agent is selected for the first time, its repertoire is still empty. In the role of initiator it then creates a random sound

from the articulatorily possible range and adds this prototype to its memory. An imitator with an empty repertoire adds a new prototype. In order to turn this new prototype into a close imitation of the heard sound, the agent adopts a rehearsal strategy in which the agent repeats the sound for itself and improves its pronunciation with a hill-climbing heuristic. In figure 2 both a successful (2(a)) and an unsuccessful (2(b)) game are illustrated.

(a) Imitation game success.

(b) Imitation game fail.

Figure 2: Procedure of an imitation game: 1. selecting a random vowel prototype, 2. producing this vowel, 3. perception of this vowel in the acoustic space of the other agent, 4. finding the closest vowel prototype, 5. producing this vowel, 6. perception of this vowel in the acoustic space of the first agent, 7. finding the closest vowel prototype. In 2(a) the game results in a success because the recognized vowel in step 7 is the same as the one produced in step 1. In 2(b) it is a failure because the recognized vowel in step 7 is not the same as the one produced in step 1.

4.5 Memory update steps

In response to an imitation game both players update their vowel memory. After each game the use and success counters for the vowel prototypes are updated for both the initiator and the imitator. Whenever the success/use ratio of a prototype moves below a predefined threshold (0.7) and it has been used often enough (a minimum of 5 times), it is removed from the vowel repertoire. Prototypes that have become too close to each other in either acoustic or articulatory space will be merged. In articulatory space they are too close when the Euclidian distance is smaller than 0.17, corresponding to a minimal difference of 0.1 for each articulatory parameter. In acoustic space the threshold θ depends on the relative difference in Bark $\Delta_{relBark}$ and the maximal acoustic noise ψ. For a multiplication by a factor of x in Hertz, the relative difference in Bark is given by:

$$(3) \quad \Delta_{relBark}(x) = 7 \ln(x)$$

The threshold θ is then given by:

$$(4) \quad \theta = \Delta_{relBark}(1 + \psi) - \Delta_{relBark}(1 - \psi)$$

This threshold defines a noise-dependent just noticeable difference in acoustic space. There is no meaning in this model, which means that the signals do not refer to anything. In order to simulate a pressure for having more different signals in a repertoire, improving the expressivity, the agents can add a random vowel to their repertoire with a small probability, thus causing the repertoire of signals to grow. In addition, there are adjustments that only the imitator makes. These steps depend on the outcome of the game. If it was a success, the imitator will move the prototype that it used closer to the sound that was heard. If the game failed, there are two alternatives: if the prototype that was used is not a very successful one (its success/use ratio is lower than 0.5), it is improved by shifting it closer to the heard signal with a predefined step size; if, on the other hand, the used prototype is a successful one, it remains unchanged (it has contributed to successful games with other agents) and the initiator adds a new prototype which is determined with the previously described rehearsal strategy. The details of how these steps were implemented exactly can be found in the original article by de Boer (2000).

4.6 Population dynamics

In the original model (de Boer, 2000) the population does not change during an experiment. De Boer and Vogt (1999) introduced a version in which the population does change. Agents can die and new agents with empty vowel

repertoires are born. The new agents have to learn the existing vowel system from the older agents. This introduces a component of vertical transmission, which makes it possible to investigate what happens to a vowel system when it is passed from generation to generation. If the system is run for long enough, with a high enough replacement rate, at a certain point in the simulation none of the initial speakers are left in the population.

A necessary adjustment by de Boer and Vogt (1999) to the original model concerns the rehearsal strategy. In the original model agents were able to repeat the sounds for themselves an indefinite number of times but in reality this is impossible. The number of times that agents can rehearse is therefore limited to 10 steps per prototype. With this model it was shown that the vowel system is better preserved over time if the population has an age structure. In this case, agents that have been in the population for a shorter period can learn faster than the ones that have been around for longer. The older agents therefore make smaller changes to their vowel repertoires and these provide a more stable target for the new agents. The age structure is implemented by a variable step size with which the agents can shift the vowels in their repertoire. This step size decreases with age according to equation (5), where ε_t is the step size at time t, α is the speed of aging and ε_∞ is the ultimate step size. In our case the step size decreases from $\varepsilon_0 = 0.03$ to $\varepsilon_\infty = 0.01$.

$$(5) \quad \varepsilon_t \leftarrow \varepsilon_{t-1} + \alpha(\varepsilon_\infty - \varepsilon_{t-1})$$

Equation (6) defines an exponential decay following:

$$(6) \quad \varepsilon_\infty + (\varepsilon_0 - \varepsilon_\infty)e^{\ln(1-\alpha)t}$$

This equation was used to determine the value for α used in the experiments, based on the approximate number of steps that we wanted the 'sensitive period' of the agents to last.

5   Simulations

In the previous section, the original imitation game model of de Boer (2000) and the extension including population dynamics of de Boer and Vogt (1999) were described. These models have been re-implemented, with the addition of comparing two different types of the age structure: a gradual decline of learning ability such as described above and a critical period with a strict cut-off moment. In the following the simulated experiments we did with this model are described as well as the results of these simulations.

5.1   Experiments

The experiments described in this section all start off with the same emerged vowel system that was the outcome of a run of 200 000 games of the original model with 50 agents. This emerged vowel system is shown in figure 3.
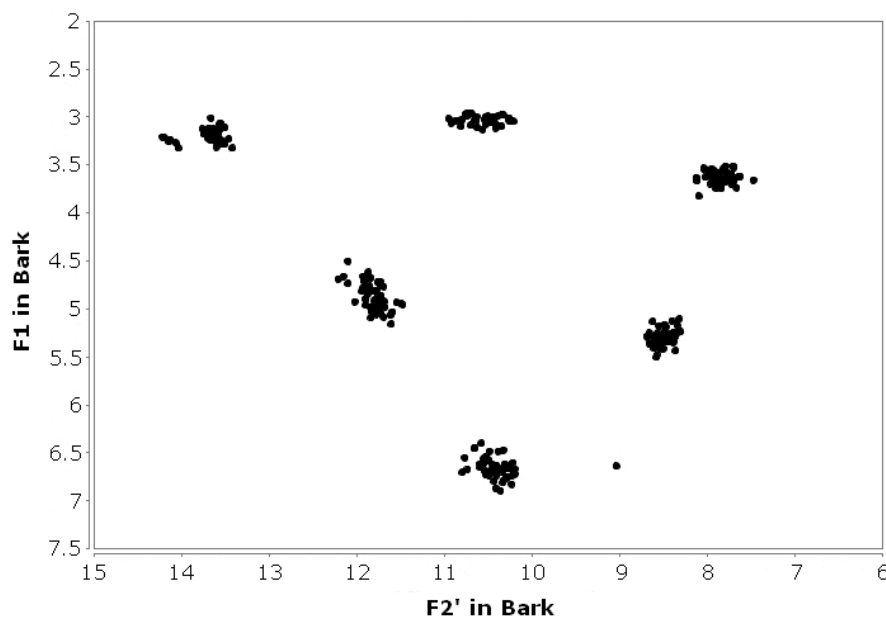

Figure 3: Emerged vowel system after 200 000 games in a population with 50 agents. Each dot represents a vowel prototype in an agent's memory. This system was used as the starting point of the experiments with changing populations.

The games were continued with a replacement probability of 0.003 per game. With this probability old agents were replaced by a new, empty agent. This replacement rate was chosen such that the agents had enough time to acquire the vowel system of the population before their replacement. Simulations were run under four different conditions. In two conditions there was no age structure and in one of these the step size, which determines the speed of learning, was large (0.03) and in the other small (0.01). In the other two conditions there was an age structure in which the learning step size decreased across the lifespan from $\varepsilon_0 = 0.03$ to $\varepsilon_\infty = 0.01$ according to equation (5). The speed of aging, α differed for these two conditions. One population aged quickly ($\alpha = 0.02$) such that the sensitive phase was approximately 1/20th of the expected lifetime. The other population aged slowly ($\alpha = 0.005$) so that agents were more sensitive for approximately 1/5th of their lifetime.

For each of the conditions 30 000 games were played so that it was unlikely that agents from the first 'generation' would still be in the final population. All runs were repeated 100 times.

5.2   Measures

At the end of each run, measures were computed for a comparison of the results: the success, the size and the energy. *Success* is the average imitation game success over all games in a run, *size* is the average over all agents of the number of vowels in their repertoire at the end of a run and *energy* is the average over all agents of the energy of their vowel systems at the end of a run. The energy of a vowel system is calculated using the method from Liljencrants and Lindblom (1972), which measures perceptual contrast (a lower energy means a better contrast). In addition, two measures were used to compute the extent to which the resulting vowel systems resembled the initial system. A *similarity* measure is used which is based on the average communicative success between agents from the initial population and those of the final population. In addition a *distance* measure is used which clusters the vowel prototypes of all agents in a population using Leader-Follower clustering (Duda et al., 2001). Then, the average is computed of all Euclidian distances between each cluster center in the final population, and the nearest cluster centers in the original population.

These two measures complement each other. The average communicative success reflects more realistically how one would assess the difference between natural languages because it uses mutual intelligibility. However, communicative success is remarkably robust to small changes in the vowel systems. In addition, communicative success depends on the size of the vowel system. With smaller systems the expected amount of mistakes is smaller. The distance measure is able to detect small changes more easily, but a disadvantage is that the results of Leader-Follower clustering are sensitive to the order of data presentation and might not work desirably in cases of confusion such as big differences in individual realizations of the same vowel. Using both measures therefore seemed like a good combination.

5.3   Results

The results of the experiments are presented in table 1, showing the computed measures for each of the four situations and averaged over the 100 runs.

From table 1 it is clear that when there is an age structure the similarity is higher than when there is none. Independent samples t-tests reveal that this difference is significant with $p < 10^{-29}$ for all four cases: comparing both aging speeds with both the low ε_0 case and the high ε_0 case. The effect sizes (computed with Cohen's *d*), which indicate the discriminability or non-overlap between the two distributions, are for all four comparisons larger than 1.5, which counts as a very strong effect according to Cohen's (1992) scale. For all four comparisons the exact effect sizes and corresponding confidence intervals (computed using the

bootstrapping method following Kelley (2005)) are shown in table 2. The distance measures of the four conditions also show a clear difference, with a lower distance in conditions with an age structure. The values for both aging speeds are significantly different (with all four $p < 10^{-15}$) from both non-aging conditions. The effect sizes were all larger than or almost equal to 1. The exact effect sizes and corresponding confidence intervals for the distance measure are shown in table 3.

	$\varepsilon_0 = 0.01$ $\varepsilon_\infty = 0.01$ $\alpha = 0$ (no aging)	$\varepsilon_0 = 0.03$ $\varepsilon_\infty = 0.01$ $\alpha = 0.02$ (fast aging)	$\varepsilon_0 = 0.03$ $\varepsilon_\infty = 0.01$ $\alpha = 0.005$ (slow aging)	$\varepsilon_0 = 0.03$ $\varepsilon_\infty = 0.03$ $\alpha = 0$ (no aging)
Success:	0.910 ± 0.012	0.879 ± 0.009	0.878 ± 0.011	0.914 ± 0.012
Energy:	0.79 ± 0.26	1.02 ± 0.23	0.98 ± 0.26	0.65 ± 0.34
Size:	3.05 ± 0.30	3.43 ± 0.30	3.35 ± 0.31	2.74 ± 0.51
Similarity:	0.719 ± 0.033	0.774 ± 0.031	0.768 ± 0.031	0.709 ± 0.037
Distance:	0.792 ± 0.104	0.686 ± 0.085	0.692 ± 0.108	0.852 ± 0.126

Table 1: Results of the experiments with and without age structure, showing the averages and standard deviations of the measures. ε is the learning step size and α the speed of aging. Note that the similarity is higher and the distance lower for runs with the age structure: the vowel systems are preserved better. The average success is smaller in the two cases with aging because vowel systems here are larger.

In addition to this, the size measures show that the age structure helps to preserve the original size of the vowel system. The size in conditions with age structure remains significantly higher than in conditions without age structure (with all four $p < 10^{-16}$ and effect sizes larger than 1, as shown in table 4).

No aging			Fast aging	Slow aging
low ε_0		d	1.7118	1.5506
		CI	$1.3694 - 2.1668$	$1.2086 - 1.9849$
high ε_0		d	1.8840	1.7337
		CI	$1.5407 - 2.3287$	$1.3842 - 2.1872$

Table 2: Similarity measure: Effect sizes (based on Cohen's *d*) and the corresponding 95% confidence intervals (CI).

It seems that there still exists some disagreement on the exact slope of the critical period for language acquisition and about whether it is a strict cut-off moment or a more gradual decline (Birdsong, 2005). We do not take a position in this discussion since it is not important for our simulation how exactly learning changes over age, as long as younger learners learn faster than older learners. The critical factor is that there is "a temporal span during which an

No aging		Fast aging	Slow aging
low ε_0	d	1.1174	0.94697
	CI	0.83322 − 1.48669	0.65448 − 1.27165
high ε_0	d	1.5433	1.3646
	CI	1.2742 − 1.8750	1.0840 − 1.6629

Table 3: Distance measure: Effect sizes and confidence intervals.

No aging		Fast aging	Slow aging
low ε_0	d	1.2915	1.0052
	CI	0.96376 − 1.73162	0.7016 − 1.3856
high ε_0	d	1.6632	1.4560
	CI	1.3408 − 2.0171	1.1734 − 1.7888

Table 4: Size measure: Effect sizes and confidence intervals.

organism displays heightened sensitivity to certain environmental stimuli" (Birdsong, 2005), in our case to language. To illustrate that the exact shape of the age structure does not influence the current results, we repeated the experiments that were described above, but this time with a strict cut-off moment instead of a gradual decline in learning ability. All parameter values are the same except for the ones concerning aging. We considered both an early critical period (strict cut-off at 1/5[th] of the expected life time) and a late critical period (at 1/2 of the expected life time) at which the step size decreased at once to 0. Table 5 shows the results of these experiments. For both the similarity and the distance measure as well as the size, the differences between the conditions with and without critical period are significant (with every p $< 10^{-30}$ and effect sizes above 2).

	$\varepsilon_0 = 0.00$ $\varepsilon_\infty = 0.00$ No CP	$\varepsilon_0 = 0.03$ $\varepsilon_\infty = 0.00$ CP at $\frac{1}{5}$ of lifetime	$\varepsilon_0 = 0.03$ $\varepsilon_\infty = 0.00$ CP at $\frac{1}{2}$ of lifetime	$\varepsilon_0 = 0.03$ $\varepsilon_\infty = 0.03$ No CP
Success:	0.922 ± 0.009	0.846 ± 0.010	0.895 ± 0.018	0.919 ± 0.010
Energy:	0.65 ± 0.16	1.29 ± 0.23	0.95 ± 0.25	0.59 ± 0.27
Size:	2.90 ± 0.11	3.79 ± 0.25	3.45 ± 0.33	2.67 ± 0.44
Similarity:	0.698 ± 0.020	0.804 ± 0.021	0.780 ± 0.029	0.688 ± 0.038
Distance:	0.835 ± 0.039	0.394 ± 0.068	0.567 ± 0.116	0.858 ± 0.105

Table 5: Results of the experiments with a strict critical period (CP), going to step size (ε_0) of 0, showing the averages and standard deviations. As in the case with the gradual decline in learning ability, the similarity is higher and the distance lower for runs with the age structure: the vowel systems are preserved better.
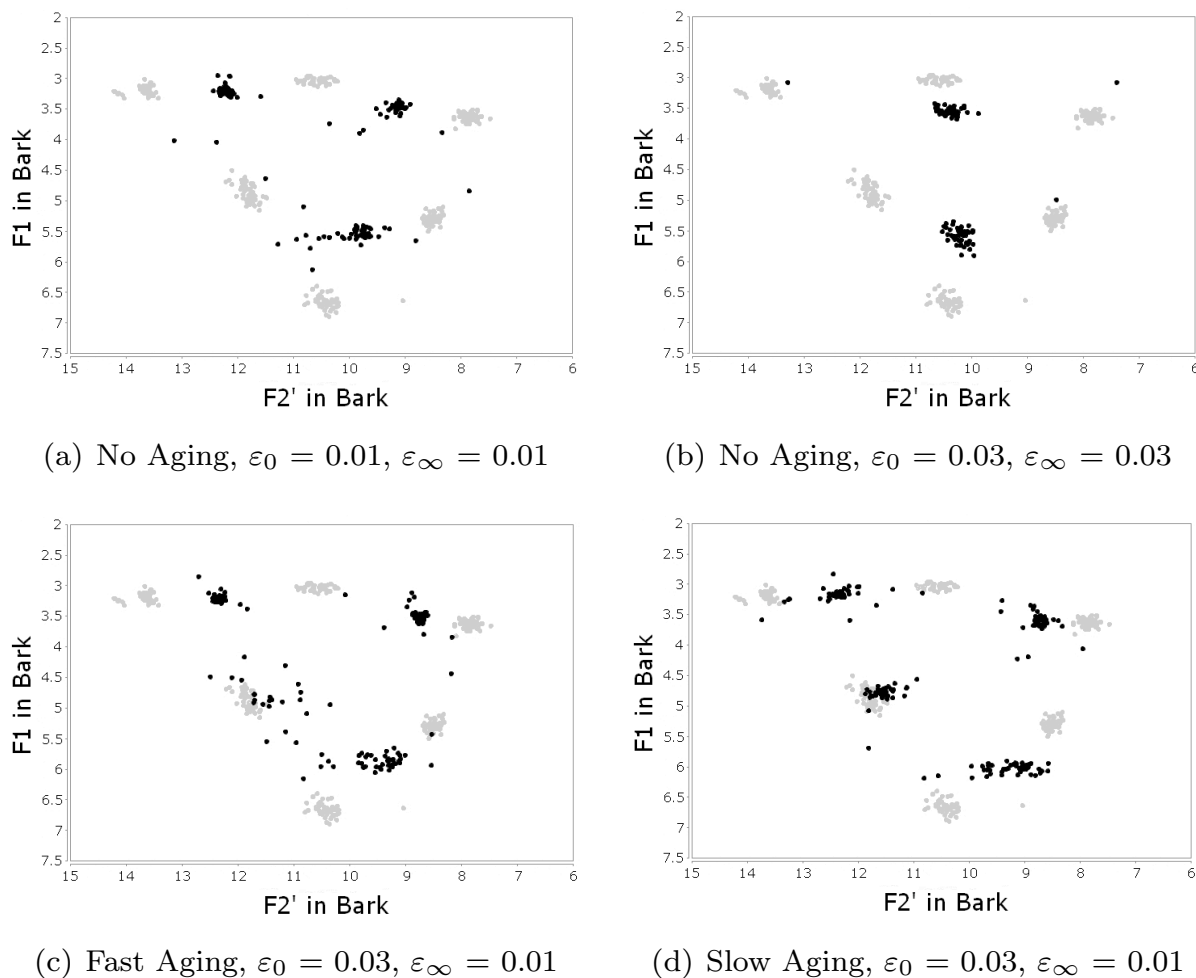
(a) No Aging, $\varepsilon_0 = 0.01$, $\varepsilon_\infty = 0.01$

(b) No Aging, $\varepsilon_0 = 0.03$, $\varepsilon_\infty = 0.03$

(c) Fast Aging, $\varepsilon_0 = 0.03$, $\varepsilon_\infty = 0.01$

(d) Slow Aging, $\varepsilon_0 = 0.03$, $\varepsilon_\infty = 0.01$

Figure 4: Influence of the age structure on the preservation of vowel systems. Note that although in all cases changes occur, they are smaller for the systems with the age structure.

6   Discussion

A simulation of the emergence of vowel systems was presented, using the model of de Boer (2000) and de Boer and Vogt (1999), in which the population went through several generations and the influence of the presence of an age structure on the culturally evolving vowel system could be measured. This model shows very clearly that vowel systems in a changing population remain more complex and are preserved better if agents learn faster when they are younger. The results presented here reproduce the findings in de Boer and Vogt (1999) and show that the same counts in the case of a differently shaped age structure. Vowel systems in a population with age structure stay closer to the initial vowel system in terms of size and shape, and the mutual intelligibility between initial and final generations is higher.

The stability-plasticity dilemma, which is an important problem in the design of intelligent autonomous systems (Carpenter and Grossberg, 1988), suggests that mechanisms for self-stabilization are needed to be able to function in the real world: "non-self-stabilizing learning systems are not capable of functioning autonomously in ill-controlled environments" (Carpenter and Grossberg, 1988). One learning system that possesses such self-organizing properties, Hebbian learning (Hebb, 1949), has been applied to the problem of language acquisition (McClelland et al., 1999; McCandliss et al., 2002), showing how it can be easier for such a model to learn new representations in an earlier phase than in a later phase. For language acquisition therefore, it might be that our general selfstabilizing learning system results in observable age sensitivity effects, causing enough plasticity to be able to adjust in foreign environments while being able to identify and communicate with people from home as well. Our findings suggest that the stabilization reaches further than the level of the individual, and that it has additional advantageous consequences on the population level. We found that vowel systems remain larger in populations with an age structure than in populations without an age structure. The fact that older agents do not adjust as fast as younger agents provides new agents with a more stable target. This is expected to facilitate acquisition because it causes less confusion and ambiguity in the interactions. In consequence, this makes it possible for the agents to learn more complex systems. Apparently, an age structure can contribute to the (cultural) preservation of the complexity of the communication system.

As mentioned before, our findings are corroborated by findings from the field of sociolinguistics, where it has been suggested that an age structure plays a role in the preservation of structure in linguistic change. Like our results, these findings indicate that an unbroken chain of transmission from adults to children enhances stability and regularity. Similar stability is not found in adult-to-adult contact only (Labov, 2007).

Although both the sociolinguistic finding by Labov (2007) and the simulation that was presented here point in the same direction and both clearly indicate that it is important to consider the consequences of language acquisition age effects in the study of language preservation and change, there are also points of difference which makes it harder to make a more in-depth comparison at this moment.

First, the examples of real-life linguistic change that were used in Labov (2007) to illustrate the difference between transmission and diffusion both involved language contact. In the computer simulations, our populations did not come in contact with other populations, but the changes in the vowel systems happened through cultural transmission only. It might be that language contact introduces other dynamics that would change the results we found in the current

implementation. Therefore, it might be necessary to model the language contact situation as well. With a type of language game that is related to the imitation games, the 'Naming Game', similar experiments have been designed in which a spatial structure exists and language contact is modeled (Steels, 1997a). It remains to be discovered whether the integration of so many different aspects (several open populations, a spatial structure, an age structure and contact) in one simulation run will be suitable to provide much more clarification. This is because it becomes harder and harder to analyze simulations when several levels of complexity interact, but an integration is definitely worth the try.

Second, the linguistic changes in Labov (2007) all involve structures and constraints that are more complex than the level of vowel shifts. They deal with the borrowing of complex rules determining for instance differences in vowel quality connected to lexical constraints. Therefore, to be able to model the same phenomena of linguistic contact, we might need a model with more complex repertoires of shared structures. We expect that the found effect of stabilization is generalizable to more complex sound systems and other areas of linguistics. In future work, it would therefore be interesting to find out what would happen if the individuals in the population were able to learn more complex sound patterns, such as syllables. This could be done, for example, by combining the experiments of de Boer and Vogt (1999) with syllable production and perception models from Oudeyer (2001, 2005).

Third, the literature indicates that the role children play in language systems is very important in the formation of structure and that their learning is not only faster but also of a different nature. Newport (1988) describes differences between adult and infant learning of American Sign Language and finds that adult learners use different generalization strategies. The signs produced by adult learners have a more holistic relation to their meanings while infant learners make more use of generalization and analyze words into smaller structural elements. Hudson Kam and Newport (2005) investigated this difference in more detail in artificial language learning experiments. Both adults and children in these experiments were taught an artificial language, which contained 'unpredictable variation': inconsistencies in the language. It turned out that both groups could learn the language, but that children as opposed to adults regularized the inconsistencies. Adults showed probability-matching behavior, exactly reproducing the variability in the input, while children found (whether they were actually there or not) more general patterns in the input and removed the inconsistencies (Hudson Kam and Newport, 2005). Although for language acquisition, differences in adult and child learning, whatever the underlying mechanisms, are observed to result in faster acquisition in childhood than in adulthood, a model in which adults and children only differ in their speed of acquisition might fail to grasp the exact nature of the actual differences.

In short, more complete models would be needed to simulate Labov's proposal more exactly because the current model might lack some important ingredients. Spatial structure, social structure, meaning and language contact, for instance, are not taken into consideration in the current simulations, which has the advantage of making it easier to analyze the behavior of the model, but a disadvantage is the loss of realism. However, it is certainly interesting that this model, even with many abstractions, still points in the same direction and indicates that the age structure in language learning ability plays an important role at the population level of language change and stability.

By using only one specific computer model to show the consequences of an age structure on a culturally evolving communication system, we run the risk of drawing conclusions about findings that are actually the result of behaviors of this particular model instead of results that have to do with the phenomenon that is investigated. This is another reason why it is important to replicate these results with other models, not only with more complete or complex versions of the current simulations, but also with different approaches. Currently a second manuscript is in preparation in which we describe the use of a different, linguistically grounded, model for vowel category learning. We have found our results to be reproducible with this different model.

We investigated the consequences that an age structure has on a culturally evolving communication system. We did not show, however, where such an age structure might have originated from. The age structures in our model were built-in and therefore it cannot explain what caused such a structure to come into being. Although we show that a change of language learning ability has a clear population level benefit, this does not prove that it evolved for this reason. Many more plausible explanations have already been proposed. It could have been a by-product of evolution under the assumption that language provides a selective advantage (Hurford, 1991), it may be the result of general self-stabilizing learning mechanisms (McClelland et al., 1999; McCandliss et al., 2002) or it might have originated from an interaction of several evolutionary and developmental influences.

In summary, we conclude that the results presented in the preceding sections indicate how the presence of an age structure can have an effect on the cultural evolution of speech. It has been shown that an age structure causes preservation and stabilization of sound systems in an open population. On a coarse scale, our findings are in line with the observations in real language evolution, namely the recent sociolinguistic findings that changes in language learning ability can account for the different ways that linguistic changes spread. However, to make a more legitimate comparison, more complete simulations are desirable.

7 References

Birdsong, D. (2005). Interpreting age effects in second language acquisition. In J. Kroll, & A. De Groot (Eds.) *Handbook of Bilingualism: Psycholinguistic Perspectives*, New York: Oxford University Press, 109–127.

Carpenter, G. A., & Grossberg, S. (1988). The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*, 21(3): 77–88.

Cohen, J. (1992). Statistical power analysis. *Current Directions in Psychological Science*, 1(3): 98–101.

Curtiss, S., Fromkin, V., Krashen, S., Rigler, D., & Rigler, M. (1974). The linguistic development of genie. *Language*, 50 (3): 528–554.

de Boer, B. (1997). Self-organisation in vowel systems through imitation. In J. Coleman (Ed.) *Computational Phonology, Third Meeting of the ACL Special Interest Group in Computational Phonology*. 19–25.

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

de Boer, B. (2006). Computer modelling as a tool for understanding language evolution. In N. Gontier, J. P. Van Bendegem, & D. Aerts (Eds.) *Evolutionary Epistemology, Language and Culture*, Springer, 381–406.

de Boer, B., & Vogt, P. (1999). Emergence of speech sounds in changing populations. In D. Floreano, J. D. Nicoud, & F. Mondada (Eds.) *Advances in Artificial Life, Lecture Notes in Artificial Intelligence*, Berlin Springer Verlag, 664–673.

Doupe, A. J., & Kuhl, P. K. (1999). Birdsong and human speech: common themes and mechanisms. *Annual Review of Neuroscience*, 22:567–631.

Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern Recognition*. New York: Wiley-Interscience.

Flege, J. E. (1999). Age of learning and second language speech. In D. P. Birdsong (Ed.) *Second language acquisition and the critical period hypothesis*, Lawrence Erlbaum, 101–131.

Flege, J. E., Munro, M. J., & MacKay, I. R. (1995). Factors affecting strength of perceived foreign accent in a second language. *The Journal of the Acoustical Society of America*, 97:3125–3134.

Fox, A. (1995). *Linguistic reconstruction: an introduction to theory and method*. Oxford University Press, USA.

Hall, R. (1962). The life cycle of pidgin languages. *Lingua*, 11(1): 151–156.

Harlow, H. F., Dodsworth, R. O., & Harlow, M. K. (1965). Total social isolation in monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, 54(1): 90–97.

Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.

Hess, E. H. (1958). Imprinting in animals. *Scientific American*, 198(3): 81–90.

Hubel, D. H., & Wiesel, T. N. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *Journal of Physiology*, 206(2): 419–436.

Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language Learning and Development*, 1 (2): 151–195.

Hurford, J. R. (1991). The evolution of the critical period for language acquisition. *Cognition*, 40(3): 159–201.

Hurford, J. R., & Kirby, S. (1999). Co-evolution of language-size and the critical period. In D. P. Birdsong (Ed.) *Second language acquisition and the critical period hypothesis*, Lawrence Erlbaum, 39–63.

Immelmann, K., & Suomi, S. J. (1981). Sensitive phases in development. In K. Immelmann, G. W. Barlow, L. Petrinovich, & M. Main (Eds.) *Behavioral Development: The Bielefeld Interdisciplinary Project*, Cambridge University Press, 395–431.

Johnson, J. S., & Newport, E. L. (1989). Critical period effects in second language learning: the influence of maturational state on the acquisition of english as a second language. *Cognitive Psychology*, 21(1): 60–99.

Kelley, K. (2005). The effects of nonnormal distributions on confidence intervals around the standardized mean difference: Bootstrap and parametric confidence intervals. *Educational and Psychological Measurement*, 65(1): 51–59.

Kirby, S., & Hurford, J. R. (2002). The emergence of linguistic structure: an overview of the iterated learning model. In A. Cangelosi, & D. Parisi (Eds.) *Simulating the Evolution of Language*, New York: Springer Verlag, 121–148.

Komarova, N. L., & Nowak, M. A. (2001). Natural selection of the critical period for language acquisition. In Proceedings: Biological Sciences, The Royal Society, 268: 1189–1196.

Kuhl, P. K. (2000). A new view of language acquisition. *Proceedings of the National Academy of Sciences*, 97(22): 11850–11857.

Labov, W. (2007). Transmission and diffusion. *Language*, 83(2): 344–387.

Lenneberg, E. H. (1969). On explaining language. *Science*, 164(880): 635–643.

Liljencrants, J., & Lindblom, B. (1972). Numerical simulation of vowel quality systems: The role of perceptual contrast. *Language*, 48(4): 839–862.

Lorenz, K. Z. (1937). The companion in the bird's world. *The Auk*, 54(3): 245–273.

Mayberry, R. I. (1993). First-language acquisition after childhood differs from second-language acquisition: the case of american sign language. *Journal of speech and hearing research*, 36(6): 1258–1270.

Mayberry, R. I. (2007). When timing is everything: Age of first-language acquisition effects on second-language learning. *Applied Psycholinguistics*, 28: 537–549.

McCandliss, B. D., Fiez, J. A., Protopapas, A., Conway, M., & McClelland, J. L. (2002). Success and failure in teaching the [r]–[l] contrast to Japanese adults: Tests of a Hebbian model of plasticity and stabilization in spoken language perception. *Cognitive, Affective, and Behavioral Neuroscience*, 2: 89–108.

McClelland, J. L., Thomas, A. G., McCandliss, B. D., & Fiez, J. A. (1999). Understanding failures of learning: Hebbian learning, competition for representational space, and some preliminary experimental data. *Progress in brain research*, 121: 75–80.

Meadow, S. G. (1978). Review: A study in human capacities. *Science*, 200(4342): 649–651.

Newport, E. (1988). Constraints on learning and their role in language acquisition: Studies of the acquisition of american sign language. *Language Sciences*, 10(1): 147–172.

Oudeyer, P. (2001). The origins of syllable systems: an operational model. *In proceedings of the International Conference on Cognitive Science*, Lawrence Erlbaum, 744–749.

Oudeyer, P. Y. (2005). How phonological structures can be culturally selected for learnability. *Adaptive Behavior*, 13(4): 269.

Oyama, S. (1976). A sensitive period for the acquisition of a nonnative phonological system. *Journal of Psycholinguistic Research*, 5(3): 261–283.

Sankoff, G., & Laberge, S. (1980). *The acquisition of native speakers by a language*, ch. 10, University of Pennsylvania press, 195–209.

Schroeder, M. R., Atal, B. S., & Hall, J. L. (1979). Objective measure of certain speech signal degradations based on masking properties of human auditory perception. In B. Lindblom, & S. Ohman (Eds.) Frontiers of speech communication research, London: Academic Press, 217–229.

Schwartz, J., Boë, L., Vallée, N., & Abry, C. (1997). The dispersion-focalization theory of vowel systems. *Journal of Phonetics*, 25(3): 255–286.

Senghas, A., Kita, S., & Ozyurek, A. (2004). Children creating core properties of language: Evidence from an emerging sign language in nicaragua. *Science*, 305(5691): 1779–1782.

Singleton, D. (2002). Age and second language acquisition. *Annual Review of Applied Linguistics*, 21: 77–89.

Spalding, D. A. (1873). Instinct, with original observations on young animals. *Macmillan's Magazine*, 27: 282–293.

Steels, L. (1997a). Language learning and language contact. In W. Daelemans, A. Van den Bosch, & A. Weijters (Eds.) *Workshop Notes of the ECML/MLnet Familiarization Workshop on Empirical Learning of Natural Language Processing Tasks*, Prague, 11–24.

Steels, L. (1997b). The synthetic modeling of language origins. *Evolution of communication*, 1(1): 1–34.

Weber Fox, C. M., & Neville, H. J. (1996). Maturational constraints on functional specializations for language processing: Erp and behavioral evidence in bilingual speakers. *Journal of Cognitive Neuroscience*, 8(3): 231–256.

Yun, A. J., Bazar, K. A., & Lee, P. Y. (2004). Pineal attrition, loss of cognitive plasticity, and onset of puberty during the teen years: is it a modern maladaptation exposed by evolutionary displacement? *Medical hypotheses*, 63(6): 939–950.