



Norwegian University of  
Science and Technology

# Using Evolutionary Algorithms to Investigate the Co-Evolution of Language and Social Networks

**Andreas Lønes**

Master of Science in Informatics

Submission date: July 2017

Supervisor: Björn Gambäck, IDI

Norwegian University of Science and Technology  
Department of Computer Science



## ABSTRACT

Scientists have been researching how human language evolved into the complex language it is today for a long time. This is a difficult task because no historical data has been discovered concerning the earliest forms of human communication. In recent years, computer scientists have attempted to simulate language evolution using computational models.

This thesis will continue the work of using computer simulations to investigate the effects cultural, biological, and social evolution have on language evolution. Four state-of-the-art computational models are discussed: Lipowska (2011), Lekvam (2014), Gong (2011), and Munroe and Cangelosi (2002).

The computational model designed for this thesis has mainly been inspired by Lekvam's work, which used genetic algorithms to simulate the co-evolution of language and social structure. In order to further investigate the effects of cultural and biological evolution, the new model has a redesigned fitness function, an alternative method of calculating the weight of a word, and an original design of the genome. The genome was designed based on the theory of 10 basic human values which was formulated by Schwartz (1992). This theory claims that a person's personality is defined by the values that person has.

The experiments show that parts of the computational model worked as intended. The results indicate that mostly acting introvertly is beneficial compared to an extrovert strategy, that the first years are crucial when it comes to an agent's language, and that language evolution to a certain extent is influenced by the Baldwin effect. However, there are some simplifications that simply cannot be ignored. The naming game is an overly-simplification of human language that might lead the simulations to reach one common language too quickly. Furthermore, the simulated social network does not properly depict the evolution of a social structure because there are no geographical distances between the individuals. This leads to the social network becoming one large network too quickly, which in turn aids the simulation in quickly reaching one common language.

This research has contributed some new ideas that seem to work as intended and brought new issues into the light. With future work being put into researching language games and how social networks affect language evolution, it could bring us one step closer to understanding how human language evolved.



## SAMMENDRAG

Det har lenge blitt forsket på hvordan det menneskelige språk utviklet seg til det komplekse språket det er i dag. Det finnes ikke historisk data fra de tidligste fasene av det menneskelige språk, noe som gjør det veldig vanskelig å forstå hvordan det menneskelige språk oppstod og utviklet seg i begynnelsen.

Denne oppgaven fortsetter arbeidet med å bruke datasimuleringer til å undersøke effekten kulturell, biologisk og sosial evolusjon har på språkutvikling. Fire nylig designede datamodeller vil diskuteres, og teorien som kreves for å forstå modellene vil presenteres i denne oppgaven (Lipowska, 2011; Lekvam, 2014; Gong, 2011; Munroe and Cangelosi, 2002).

Datamodellen som ble designet til denne oppgaven er inspirert av arbeidet til Lekvam. Før å undersøke effekten av kulturell og biologisk evolusjon videre ble modellen redesignet med en ny *fitness function*, en ny metode for å kalkulere vekten av et ord og et nytt genom. Genomet ble designet basert på arbeidet til Schwartz som definerte en teori kalt *the theory of 10 basic values*. Teorien påstår at personligheter kan defineres utifra hvilke verdier en person har.

Eksperimentene viser at deler av modellen fungerte som planlagt. Resultatene indikerer at en introvert strategi er mer lønnsom enn en ekstrovert strategi, at de første årene er avgjørende for en agents' språk og at Baldwin effekten påvirker utviklingen av språk. Resultatene viste også at deler av modellen var for simplifiserte. *Naming gamet* er en grov forenkling av det menneskelige språk og det kan ha ført til at simuleringer nådde et felles språk for raskt. Det sosiale nettverket tok ikke hensyn til geografiske avstander, som kan ha ført til at det sosiale nettverket ble til et stort nettverk for fort. At alle agentene var en del av samme nettverk har også bidratt til at det gikk raskere å nå et felles språk.

Forskningen har bidratt med flere gode ideer som virker å fungere, for eksempel *the fitness function* og genomet. Den har brakt frem flere nye problemstillinger som må studeres videre. Den primære utfordringen er å simulere evolusjonen av språk og sosiale nettverk fordi modellene må inkludere fler av faktorene som påvirker deres evolusjon for å tilstrekkelig simulere utviklingen av språk og sosiale nettverk. Videre arbeid innen språkmodeller og hvordan sosiale nettverk påvirker utviklingen av språk kan bringe oss et steg nærmere å forstå hvordan det menneskelige språk har utviklet seg til det språket det er i dag.



# Preface

This is a Master's thesis in Computer Science with specialisation in Artificial Intelligence at the Norwegian University of Science and Technology (NTNU). The study presented in this thesis was carried out from August 2016 to July 2017.

I would like to thank Professor Björn Gambäck for providing me with the opportunity to work on such an interesting topic. I am grateful for the support and valuable discussions throughout this work. I would also like to thank Per Odlo, who wrote a masters thesis within the same topic this past year, for the discussions and input to this study. Finally, I would like to thank Torvald Lekvam whose work has inspired the work performed in this study.

Andreas Lønes  
Trondheim, Norway  
July 2017



# Table of Contents

- Preface v
  
- List of Figures ix
  
- List of Tables x
  
- 1 Introduction 1**
  - 1.1 Research Questions . . . . . 2
  - 1.2 Contribution . . . . . 2
  - 1.3 Outline of the Thesis . . . . . 2
  
- 2 Background 5**
  - 2.1 Language Evolution . . . . . 5
  - 2.2 Language Games . . . . . 8
  - 2.3 Schwartz' Theory of Basic Values . . . . . 10
  - 2.4 Graph Theory . . . . . 14
  - 2.5 Social Networks . . . . . 15
  - 2.6 Genetic Algorithms . . . . . 16
  
- 3 State-of-the-Art 19**
  - 3.1 Lipowska, 2011 . . . . . 19
  - 3.2 Gong, 2004 . . . . . 21
  - 3.3 Munroe & Cangelosi, 2002 . . . . . 23
  - 3.4 Lekvam, 2014 . . . . . 25
  - 3.5 Comparison of the Four Models . . . . . 27
  
- 4 Methodology 31**
  - 4.1 Genetic Algorithm . . . . . 32

4.2	Genome . . . . .	33
4.3	Fitness Function . . . . .	34
4.4	Social Network and Dialogues . . . . .	34
<b>5</b>	<b>Results</b>	<b>37</b>
5.1	Parameters . . . . .	37
5.2	Experiments . . . . .	37
<b>6</b>	<b>Discussion</b>	<b>57</b>
<b>7</b>	<b>Conclusions and Future Work</b>	<b>69</b>
<b>A</b>	<b>Results Not Included in the Main Report</b>	<b>73</b>
A.1	Experiment 4 . . . . .	74
A.2	Experiment 5 . . . . .	75
A.3	Experiment 6 . . . . .	76
A.4	Experiment 7 . . . . .	77
	<b>References</b>	<b>79</b>

# List of Figures

- 2.1 Illustration of the three forces of language evolution. . . . . 7
- 2.2 The ten values from *the Schwartz Theory of Basic Values*. . . . . 13
- 2.3 The difference between a directed and an undirected graph. . . . . 15
- 2.4 Flowchart of a general genetic algorithm (GA). . . . . 17
- 2.5 How crossover is used to create two offsprings from two parents with a genome represented by a bit-string. . . . . 18
  
- 5.1 Graphs of the results from experiment 2 as a function of number of generations. . . . . 40
- 5.2 Snapshots of the social network at different generations in experiment 1. . . 42
- 5.3 Graphs of the results from experiment 2 as a function of number of generations. . . . . 43
- 5.4 Snapshots of the social network at different generations in experiment 2. . . 45
- 5.5 Graphs of the results from experiment 3 as a function of number of generations. . . . . 46
- 5.6 Snapshots of the social network at different generations in experiment 3. . . 48
- 5.7 Graphs of the results from experiment 4 as a function of number of generations. . . . . 49
- 5.8 Graphs of the results from experiment 5 as a function of number of generations. . . . . 51
- 5.9 Graphs of the results from experiment 6 as a function of number of generations. . . . . 53
- 5.10 Graphs of the results from experiment 7 as a function of number of generations. . . . . 55
  
- 6.1 Comparing the average vocabulary size per generation for experiment number 1, 2, 3, 5, and 7b. . . . . 59

6.2	Comparing the number of unique highest ranked words per generation for experiment 1, 2, 3, 4, 5, 6a, 6b, 7, and 7b. . . . .	60
6.3	Comparing the fitness per generation for experiment 1, 2, 3, 4, 6a, 6b, and 7b. . . . .	61
6.4	Comparing the number of unique highest ranked words per generation for experiment 1, 2, 3, 4, 5, 6a, 6b, and 7b. . . . .	61
6.5	The social network at generation number 10 in experiment 5. . . . .	64
6.6	Comparison of the 5th generation in experiment number 1 and the 100th generation in experiment number 5. . . . .	64
6.7	Comparing the rate of successful dialogues per generation for experiment 1 and 7b. . . . .	66
A.1	Snapshots of the social network of experiment 4 at the following generations: 5, 20, 40, and 100. . . . .	74
A.2	Snapshots of the social network of experiment 5 at the following generations: 5, 20, 100, and 150. . . . .	75
A.3	Snapshots of the social network of experiment 6 at the following generations: 5, 20, and 100, and graphs showing the evolution of the the traits in the genome. . . . .	76
A.4	Snapshots of the social network of experiment 7 at the following generations: 5, 20, and 100. . . . .	77

## List of Tables

5.1	The parameters that is used in the model and their corresponding default values. . . . .	38
-----	--	----

# Chapter 1

## Introduction

The complexity of human language is one of the biggest distinctions between us and other animals (Hauser et al., 2002). Many have tried to explain how human language evolved into what it is today, but no one has managed to do so yet (Chomsky, 1986; Pinker and Bloom, 1990; Müller, 1861). In recent years, computational models have been used to get a greater insight into the field of language evolution. The main contributions have come from the computational models used to test the validity of language evolution theories through computer simulations.

Four computational models which have advanced the field of language evolution are presented and discussed in this thesis:

1. Gong (2011), who simulated agents that attempted to reach a consensus by building up a compositional language using spatial naming games.
2. Munroe and Cangelosi (2002), who made a signalling game where an agent's understanding of the language was represented using an artificial neural network (ANN) and evolved by agents trying to achieve a consensus on the state of the objects in the world.
3. Lipowska (2011), who made a computational simulation using naming games which analysed the cultural transmission over several generations.
4. Lekvam (2014), who used a naming game and an evolutionary algorithm to simulate the co-evolution of language and social networks.

The computational model designed for this thesis is based on the work by Lekvam. This thesis will attempt to improve upon his work by adding a new *fitness function*, adding a new method for calculating the weights in the social network, slightly altering how agents decide to conduct a conversation, and introducing a new *genome*. The new genome is based on the work by Schwartz (1992) who formulated a theory which attempts to define personalities through ten basic human values. His theory will be presented in Chapter 2.

## 1.1 Research Questions

The research questions of this thesis are:

1. How are the state-of-art computational models in the field of simulating language designed and what are their strengths and weaknesses?
2. How can one state-of-the-art computational model be extended in order to learn more about how the evolutionary forces affect language evolution?
3. How does social networks have an impact on language evolution?

## 1.2 Contribution

This study will contribute to the field of computational language evolution models from a computer science perspective. Some new ideas will be presented that contribute to enhance our understanding of the effect the evolutionary forces have on language evolution and the co-evolution of language and social networks.

## 1.3 Outline of the Thesis

In order to understand the field of language evolution and how computational models simulating language evolution are designed, necessary theories are presented in Chapter 2. This includes theories on language evolution, graph theory, social networks, genetic algorithms (GAs), and a theory concerning how to model personalities. In Chapter 3, four recently developed computational models designed by Gong, Lipowska, Lekvam, and

Munroe and Cangelosi, respectively, will be presented and discussed. The design of the computational model used in this study is presented in Chapter 4. Chapter 5 presents the results of seven experiments that were performed using the model presented in the previous chapter. Then, in Chapter 6 the results are discussed. Finally, in Chapter 7 a summary and conclusion of the results and recommendations for future work are presented.



# Chapter 2

## Background

To better comprehend the models that are presented in this thesis, an understanding of some central theories within the field of language evolution is necessary. This chapter will present language evolution, language games, Schwartz' theory of basic human values, graph theory, social network theory and genetic algorithms (GAs).

### 2.1 Language Evolution

Biologists, evolutionary linguists, and sociolinguists all have different theories on how the human language evolved. The biologists think that the evolution of human language can be explained through biological evolution, and they argue that the composition of a language, e.g. the sentence structure subject-verb-object in English, can be explained through a language acquisition device which can understand a universal grammar (Chomsky and Halle, 1965). Biologists language has evolved through natural selection and genetic evolution.

The evolutionary linguists believe that language evolved through natural selection in a Darwinian manner. Pinker and Prince (1994) proposed an explanation for the evolution of human language by looking at the evolution of non-human animals' abilities. They suggested that, just like bats evolved the ability to perceive their environment through echolocation, humans evolved the ability to understand and learn cognitive-functional linguistics.


In recent years, there has been a discussion concerning the importance of individual learning and how big of an impact culture and society has on individual learning. Sociolinguists believe that language evolution is affected by society and culture (Tomasello and Rakoczy, 2003). Their theory has been strengthened by the use of computational models of simplified social networks. These models show that social networks affect language evolution.

There are three forces that are believed to affect language evolution: biological evolution, cultural evolution, and social evolution, as seen in Figure 2.1 (Steels, 2011). To fully understand the theories on language evolution, the theories have to be viewed in the light of these three forces.

Biological evolution is the slowest force, working on a *phylogenetic* time scale. It involves how an individual improves its abilities to learn and process language in order to survive and reproduce. Cultural evolution works on a historical time scale, called the *glossogenetic* time scale. In cultural evolution, the change is viewed upon unique languages that exist within a society. It concerns how the features of the language get culturally transferred from one person to another, and from one generation to another. Social evolution is looked at during one person's lifetime, and the evolution is said to work on an *ontogenetic* time scale. It concerns how an individual learns a language. A person's ability to learn the language is an important factor as to how a language is built up. Newborns' ability to quickly go from having no language at all to being able to speak several languages indicates that humans and language have co-evolved so that human language will be easy to learn during infancy (Tomasello and Rakoczy, 2003).

### 2.1.1 Origin of Language

No one knows exactly how or when humans began communicating. In order for language to have emerged, there had to be a need for communication. For example, if one person saw a deer approaching, that person would want to communicate that to the others so that they could go and hunt together. Situations like this, where communication was beneficial, is likely to have arisen often. Most linguists agree upon the theory that the first human form of communication was holistic, i.e. simple utterances and physical gestures, and was used to express a reaction to concepts such as hunt or run (Christiansen and Kirby, 2003). How a holistic language evolved into the syntactical language used today is not known (Bickerton, 2007).



**Figure 2.1:** Illustration of the three forces of language evolution that are believed to affect language evolution: biological evolution, cultural evolution, and social evolution.

### 2.1.2 Baldwin Effect

The Baldwin effect is based on the idea that the traits an individual learns during its lifetime can guide evolution (Baldwin, 1896). Learning a skill comes with a cost, so if that skill is very important, the individuals with an innate ability to learn this skill have an advantage over the others. If the cost of acquiring the skill stays high and the skill remains important over time, individuals that easily master the skill will be favoured in natural selection. For example, the skill of making a fire, if this skill were to become instinctive, the population as a whole would save a lot of time not having to teach everyone how to make fire. The individuals that learn the skill easily have an advantage because they can use the spare time doing other essential tasks, like hunting for food. If the selection pressure is high enough and the cost of learning the art of making a fire is high during several generations, the Baldwin effect states that the population will eventually consist of individuals that easily learn how to make a fire.

Many researchers have experimented with the Baldwin effect during recent years. Lipowska (2011) argues that the results from her model show that the learned traits during an individual's lifetime have an effect on language evolution, meaning that the Baldwin effect has an effect on language evolution. Zollman and Smead (2010) also conclude that their

model indicates that the Baldwin effect is influencing the evolution of language. Chater and Christiansen (2010) argue that, based on their results, the Baldwin effect may have an influence on parts of language evolution. However, the parts of the language that change consistently, such as word order and morphology, i.e. the structure of words, would have a too low selection pressure for the Baldwin effect to have an impact on them.

## 2.2 Language Games

Language games are games where artificial agents interact with each other and try to understand each other. For the agents to understand each other, they need to reach a state where their individual languages are similar. This similarity is reached by the agents conversing and changing their respective vocabularies based on what the other agent utters.

Two agents are picked out, traditionally at random, and one of the agents gets the role of *hearer* and the other agent is the *speaker*. The speaker then utters something about a concept. If the hearer understands the utterance, the conversation is labelled as a success, and if the hearer did not understand, the conversation is labelled as a failure. If the speaking agent does not have any way of describing the concept perceived, it has the ability to make up new words in order to describe the concept.

Language games have become a standard method of simulating language within the field of computer simulations of language evolution. Many variations of language games have been designed: naming games, spatial naming games, and signalling games will be presented in the following subsections.

### 2.2.1 Naming Games

Simple naming games typically have a language without grammar, and all conversations occur without an environment. All agents have a vocabulary containing words they have heard or invented during conversations with other agents. Conversations occur between two agents, they attempt to agree upon an utterance for an object they both have perceived. If both agents can agree upon an utterance, the conversation is redeemed as a success, and that utterance becomes more likely to be used again for both agents. If they cannot agree, the hearing agent adds the new word to its vocabulary and the speaking

agent decreases the probability of that word being used again. Words that have not been used successfully in a long time are eventually removed from the agent's vocabulary.

During the first generations there will be many unsuccessful conversations, leading to words being spread around in the population giving all agents a large vocabulary. When the agents have added enough words to their vocabularies, some conversations will be successful. This will lead to agents removing their least successful words. If the simulation is given enough time, the agents may reach a consensus on almost all words. However, reaching a consensus on the last words seems to require some element of chance (Steels, 1997).

Naming games are often used to analyse cultural transmission between agents of the same generation. Lipowska made a computational simulation using naming games which analysed the cultural transmission over several generations. This model will be reviewed in section 3.1.

### 2.2.2 Spatial Naming Games

Spatial naming games are a form of naming games where all conversations are set in the same environment. The speaking agent describes its understanding of the environment to the listener. If the listener agrees with the description of the environment, the conversation is viewed upon as a success. If the agents do not agree, the conversation is deemed as a failure. As with normal simple games, the goal is that all agents reach a consensus. Gong simulated agents that tried to reach a consensus by building up a compositional language. His work will be examined further in section 3.2.

### 2.2.3 Signalling Games

Signalling games simulate a compositional language using an artificial neural network (ANN). The signal uttered is the output of an agent's ANN. The output will normally be a sequence of numerical values where the order of the numbers represents a sentence in natural language. This abstraction makes it easier to study compositional languages (Munroe and Cangelosi, 2002; Suzuki and Arita, 2008). Munroe and Cangelosi made a signalling game where the language was evolved by agents walking around in a world of mushrooms. The agents' task was to reach a consensus concerning which mushrooms

were edible and which were not. This was done by evolving a language from ANNs. The computational model using signalling games made by Munroe and Cangelosi will be examined further in section 3.3.

## 2.3 Schwartz' Theory of Basic Values

There is no doubt that personality affects a person's social network. Personality is so complex that even making a simplistic model of it might seem impossible. Schwartz (1992) presents a theory which models personalities by figuring out which of the types of values are most important to a person.

*The theory of Basic Values* presents ten values that all are based on one or more of the three "universal requirements of human existence" (Schwartz, 2012). These requirements are:

1. Needs of individuals as biological organisms.
2. Necessity of coordinated social interaction.
3. Survival and welfare needs of groups.

Humans need to express fitting goals to other humans in order to cooperate with others and survive. Values represent these goals and language is how these values are communicated to others. The ten values in Schwartz's theory are:

1. **Self-direction:** A person that values self-direction highly has independent thought and exploring as goals. Self-direction comes from the need for control and independence.
2. **Stimulation:** Goals that are induced from the stimulation value are challenge in life and excitement. Stimulation stems from the need for variation in order to keep a positive activation level.
3. **Hedonism:** Pleasure and appreciation of life are goals people with a priority on Hedonism has. Hedonism comes from the enjoyment that stems from fulfilling needs.
4. **Achievement:** A person that values achievement highly has personal success by displaying competence to others as goal. The achievement values come from the

fact that proficient actions are needed in order to survive and for social groups to reach their goals.

5. **Power:** The goals of someone that values power are social status and control. Social settings seem to require people to have different statuses.
6. **Security:** Safety and stability are the goals of someone that believes security is an important value. Security values work on both a personal level and a social level, e.g. family, nationality.
7. **Conformity:** Goals that are induced from the conformity value are restraint of actions and self-discipline. Conformity stems from that individuals could have inclinations to perform actions that work against the better of the group. A person which highly values conformity does not want to perform actions that work against the better of the group.
8. **Tradition:** The goals of someone with the tradition value are respect and commitment. Groups tend to develop common practises and beliefs that bind them together. These practises become traditions and symbolise that the people in the group are connected and working for the group's survival.
9. **Benevolence:** The goal of someone that values benevolence is caring about the welfare of others. Benevolence comes from requirements of how groups work. In order for a group to function, the people in it need to care about each other. This value creates feelings towards others such as friendship and love.
10. **Universalism:** Understanding and protection for all people and for nature are the goals of someone that believes universalism is an important value. Universalism stands in contrast to benevolence, which focuses on the welfare of the group exclusively. These values activate when people become aware of other groups and how fragile nature is.

According to Schwartz, values have six features that make it possible to figure out which goals a person has based on how a person prioritises the ten values. By figuring out a person's goals and assuming that a rational person is trying to fulfil those goals, it is possible to anticipate the actions a person decides to perform. When one knows a person's goals and can anticipate someone's actions, it is fair to claim that one knows the person's personality. The six features are:

1. When a value is activated, a feeling associated with the value and situation emerges.
2. A person's values indicate which goals the person has.
3. Values is a general guideline that indicates how a person will act in many different situations.
4. Values are used to evaluate possible actions. Which actions are good or bad is decided through someone's values. These evaluations are usually performed sub-consciously.
5. A person's values are ordered after importance.
6. The importance of the values guides a person's actions. Almost all situations trigger several values, so the importance of the values compared to each other affects which action is chosen.

All values have these features. What separates the values are the goals they are connected to.


To display the relations between the values, Schwartz made a circular model of them, see Figure 2.2. The figure is divided into four categories that display the values that work towards similar goals. The categories are openness to change, self-enhancement, self-transcendence, and conservation. The values of the openness to change category contradict the values of the conservation category, while the self-enhancement category contradicts the self-transcendence category. Hedonism is split between openness to change and self-enhancement because it relates to both. The farther the values are apart from each other, the more contradictory they are. Values that are opposite to each other are therefore completely contradictory values. E.g. Conformity and tradition both completely contradict stimulation and hence both are placed opposite to stimulation. Schwartz thought that conformity is slightly more contradictory to stimulation than tradition is, so conformity was placed innermost of the two.

### 2.3.1 The Schwartz Value Survey

In order to figure out how people prioritised their values, a survey was made by Schwartz (1992). The survey made personality related statements and the participants answered how well the statement fitted their personality. Based on the results each of the ten

values are assigned with a numerical value indicating how important that value is to the participant. The exact numerical value of the values is not what is important, it is their importance in relation to one another that matters. Normalisation is used on the ten values to scale them.

The survey was answered by people from all over the world with many different backgrounds. A pattern in the prioritising of the values was found. It was found that benevolence, universalism, and self-direction were the most important values, while power, stimulation, and tradition were the least important ones. Almost all nations in which the study was conducted had these results. According to Schwartz, this indicates that the way people work and the way people are affected by society is quite similar in many different cultures.



**Figure 2.2:** The ten values from *the Schwartz Theory of Basic Values*: self-direction, stimulation, hedonism, achievement, power, security, conformity, tradition, benevolence, and universalism. The circle is divided into four categories: openness to change and conservation, which represent the contradictories independence and obedience respectively, and self-transcendence and self-enhancement, which represent the contradictories interest in welfare of others and interest in welfare of oneself respectively. Values that work towards similar goals are placed in the same category (Schwartz, 2012).

## 2.4 Graph Theory

Language game models with structured networks have, in recent years, been presented in a few papers. For instance, Lipowska (2011) used a lattice to simulate a social graph, while others have incorporated social structures into their models (Lekvam, 2014; Gong et al., 2004). This section will present discrete graphs, some of their features and how they are used in social networks.

A discrete graph represents values at specific points along the number line, for example integers. This makes the graph consist of points, and not a line, which represents a function that is continuous along its entire domain. These points are usually referred to as vertices in graph theory. The vertices in a discrete graph are connected with edges. These edges can be associated with a weight representing the cost between two vertices. The edges could also be not weighted, which essentially is the same as weighting all edges with the same value. In a discrete graph cities may be represented as vertices, and the roads between them as edges. The weight of the edges between cities could represent the travel time or the distance between the cities. If it is possible to travel both ways along an edge, the edge is called an *undirected edge*. Otherwise, if it is only possible to travel in one direction along the edge, like a one-way driven street, it is called a *directed edge*. A directed graph can have up to two edges between a vertex pair, while an undirected graph can have up to one edge between a vertex pair. An example of a directed and an undirected graph can be seen in Figure 2.3. Unless otherwise specified, the graphs presented in this thesis will be undirected.


Two vertices are called neighbours if they are connected to each other by an edge. The number of neighbours a vertex has, is called the degree of the vertex. The density of a graph,  $\text{DENSITY}(G)$ , is the number of edges in the graph  $G$  divided by the theoretical maximum number of edges  $G$  can have,  $\text{MAX}_{\text{edges}}(G)$ . In an undirected graph  $G_{\text{undirected}}$ , with  $N_V$  number of vertices, the maximum number of edges are

$$\text{MAX}_{\text{edges}}(G_{\text{undirected}}) = \frac{N_V(N_V - 1)}{2} \quad (2.1)$$

Using the definition of the density of a graph from above together with Eq. (2.1), the graph's density becomes

$$\text{DENSITY}(G_{\text{undirected}}) = \frac{2N_E}{N_V(N_V - 1)}, \quad (2.2)$$

where  $N_E$  is the number of edges. Figure 2.2(b) shows an undirected graph with three vertices: A, B and C. Vertices A and B have a degree of 1 since they have one neighbour each, while vertex C has a degree of 2 since it has two neighbours. By using Eq. (2.2), the density of the graph becomes  $2 \cdot 2 / (3 \cdot 2) = \frac{2}{3}$ .



**Figure 2.3:** An illustration of the difference between (a) a directed and (b) an undirected graph. The circles represent vertices and the straight lines, both with and without and arrow, represent edges.

## 2.5 Social Networks

A social network can be represented by a discrete graph where each vertice represents a person and the edges represent how strongly two people are connected. There are a lot of factors that affect social networks and how they evolve over time. Some of these factors are language, personalities, and the environment of the social network. Several computational models that simulate the evolution of social networks have been designed.

Lipowska (2011) used a lattice where each agent was restricted to only communicating with its neighbours except during its first dialogues which it had with its parents. This model obviously has many limitations, such as agents cannot gain or lose connections and the connections are not weighted.

Lekvam (2014) used a social network where the edges were weighted based on the communication between the agents. It was also possible to loose and gain connections. This

social network is a better representation of the real world than Lipowska's lattice, but there are other factors than just having the same language that affect a social network, like the personalities of the people in the network.

## 2.6 Genetic Algorithms

A genetic algorithm (GA) is a search algorithm based on the Darwinian principle of biological evolution. Imagine that you got the job of optimising an exam schedule at a university so that every student could take their desired combination of courses and be pleased with the exam dates. This would be impossible for humans due to the large search space. There probably does not even exist a solution where every student is pleased. In problems without a guaranteed optimal solution and with a large search space, GAs perform well.

The general structure of a GA can be seen in Figure 2.4. First, a population of solutions is made, and all solutions are tested through a fitness function, which evaluates a solution and returns a fitness value that reflects how fit the solution is. Then survival selection is performed by having the fittest individuals of the population, i.e. the best solutions, survive while the rest of the population dies. The solution that each individual represents is defined by a set of parameters, often referred to as an individual's *genome*. The genome is often represented as a bit-string. New solutions are then created by having the surviving solutions combine their genomes into new genomes by crossover, see Figure 2.5. When a new solution is made, it has a small chance of mutation, i.e. a small alteration of the genome of the solution. Then these solutions are tested, and once again the fittest individuals survive to the next generation and get to breed.

The flow of a general GA looks like this (Michalski et al., 2013):

1. **Initialisation:**

Initialise a population of individuals with random solutions.


2. **Evaluation:**

Calculate the fitness of all individuals.

3. **Termination clause:**

While the highest fitness value in the population is smaller than the desired fitness value or the maximum number of iterations has not been reached, the following steps are performed:

- (a) **Parent selection:** Select which individuals to bring into the next generation based upon their fitness.
- (b) **Recombination:** Probabilistically select two individuals based on their fitness and combine their solutions into two new solutions by using crossover.
- (c) **Mutation:** Randomly select a few of the new individuals and let their genome mutate.
- (d) **Evaluation:** Calculate the fitness of the new solutions.



**Figure 2.4:** Flowchart of a general GA. First a population is initialised with random genomes and the fitness of all individuals is calculated. As long as the highest fitness value in the population is smaller than the desired fitness value, the GA loop is run: 1) the parents for the next generation are selected; 2) the genomes from the parents are combined by crossover to constitute the genomes of the offspring; 3) a mutation of the genome can occur; 4) the fitness of all individuals is calculated; and 5) if the highest fitness value in the population is higher than the desired fitness value, the loop is terminated, otherwise, the loop is run again.

A *fitness function* is a function that returns a value which gives information about how good a solution is. The fitness function of a GA that tries to find the fastest route from one city to another would simply be the travel time of the required solution. Many problems do not have so clear objectives, or have several objectives that need to be weighted in a more complex manner.

A common issue when searching with a GA is that if the problem at hand contains many local maxima and few global maxima, the algorithm is likely to get stuck in a local maximum because one solution is more likely to discover a local maximum than a global maximum. A solution which is stuck in a local maximum will then often be the fittest in the population, and spread its genes to the solutions of the following generations. It is


called *exploitation* when a good solution is spread to many agents. Exploitation allows the algorithm to perfect one solution, but it stops exploring other possibly better solutions. On the other hand, *Exploration* is searching for new types of solutions, which is done through mutation and probabilistically choosing parents.

One of the most common methods used in the parent selection phase is called *tournament selection*. This method combines exploration, by having some randomness, and exploitation, by prioritising the most fit individuals. First,  $N$  individuals are chosen at random from all individuals. Then the individuals are sorted by their fitness, and given a rank,  $r$ , accordingly, i.e. the fittest individual is given  $r = 1$ , the second fittest  $r = 2$ , etc. Then each individual is chosen with a probability, given by

$$p(1 - p)^r, \quad (2.3)$$

where  $p$  is a user-given probability,  $0 < p < 1$ , for choosing the fittest individual in a tournament.

When two parents have been selected to mate, crossover is performed and two new individuals are created based on the parents' genomes. The most common form of crossover is performed by randomly choosing one point on the bit-string, and splitting each parent's genomes into two parts, A and B. One child receives part A from one parent and part B from the other parent. The other child receives the opposite parts. See Figure 2.5 for a simple graphic explanation of crossover. Mutation is performed on an individual with a user-given probability. Mutation is normally performed by simply flipping one bit in the bit-string.



**Figure 2.5:** How crossover is used to create two offsprings from two parents with a genome represented by a bit-string. First, the genomes of the parents are divided into two parts at the same spot. Then the first part of the genome of the first parent is combined with the second part of the genome of the other parent to form the genome of the first offspring, while the first part of the genome of the second parent is combined with the second part of the genome of the first parent to form the genome of the second offspring.

# Chapter 3

## State-of-the-Art

In this chapter, a description of the computer simulations of language evolution, from Lipowska (2011); Gong et al. (2004); Munroe and Cangelosi (2002); Lekvam (2014) will be presented. These four were chosen because they have contributed with new ideas and represent different approaches in the field of using computational models to simulate language evolution.

### 3.1 Lipowska, 2011

Lipowska made a computational simulation, which used a naming game to model a non-structured language. A non-structured language is a language that has no grammar or compositionality. The objective of Lipowska's model was to illustrate the Baldwin effect.

#### 3.1.1 Model

Lipowska incorporates the naming game in her model, which means that every agent is equipped with its own vocabulary, which contains all words it has heard of. Each word in the vocabulary is associated with a weight,  $w_i$ . The weight of a word represents how successful the word has been in conversations for an agent. All agents in the model are placed in a *lattice*. A random agent is chosen as the speaker, and with probability  $p \in [0, 1]$  it chooses one of its neighbours in the lattice as the listener. If a neighbour is not chosen as a listener, the agent dies. If a neighbour is chosen as a listener, the speaker

then utters a word from its vocabulary, and the listener receives that word. Which word the speaking agent utters is decided by the weight of each word relative to the sum of all weights in its vocabulary,  $w_i / \sum_j w_j$ , where  $w_i$  is the weight of word  $i$  and  $\sum_j w_j$  is the sum of all weights of the words in the agent's vocabulary. If the agent has no words in its vocabulary, it makes one up at random. If the vocabulary of the listening agent contains the spoken word, the conversation is considered a success. If the word is not in the listener's vocabulary, the conversation has failed.

Both the listening and the speaking agent adjust the weight of their word after a conversation. If the conversation was a success, both agents increase the weight of the word according to their respective learnability variables. The learnability variable  $l \in [0, 1]$  is an alignment strategy. Alignment is a strategy that is supposed to bring their languages closer to each other (Steels, 2012). At the end of each generation, which consists of several conversations between agents, some agents die based on a probability set by the average weight in the agent's vocabulary and its age. A high average weight in an agent's vocabulary increases its fitness, whilst an agent loses fitness when it gets older. For example, if the fitness of two agents with different ages, but with the exact same average weight over their vocabularies, is compared, the youngest of the two will have the highest fitness. A young agent with a high average weight in its vocabulary is an optimal agent. A surviving agent may breed, and if it does, the offspring inherits the learnability of its parent with a certain probability. If the offspring does not inherit its parent's learnability value, it is randomly set.

### 3.1.2 Results and Discussion

The experiments that were performed with a small  $p$  resulted in agents that did not evolve a common language. Only small clusters of agents created common languages. The value of  $p$  was slightly increased for every simulation. When  $p$  increased, the size of the clusters of agents slightly increased. When  $p$  reached a certain threshold, almost all agents became a part of the same cluster. For each simulation, the communication success rate,  $s$ , was also calculated.  $s$  is defined as the fraction of all successful communications over the total number of communication attempts. When  $p$  got the value of about 0.23, the communication success rate increased rapidly.

The results of the experiment show a correlation between communication success rate and average learnability. In Lipowska's model, the new individuals require learning to

incorporate the communal language that the generation before used. By communicating with the parents, the new agents learn the language the parents used. Children with a high learnability will learn more quickly than the others. According to Lipowska, the model shows that learning can direct the evolution, which indicates that the Baldwin effect has an influence on language evolution (Lipowska, 2011, Section 4).

## 3.2 Gong, 2004

This section will present the word order regularity model, its results, and what was concluded in the paper by Gong et al. (2004).

### 3.2.1 Model

Gong made a simulation of the evolution of compositional language where agents converse about “integrated events such as tiger is running” (Gong et al., 2004, Section 3). The speaking agent makes utterances trying to explain what it perceives, while the listening agent tries to understand the utterance in accordance with partial information about the environment they are in. The language evolves from a simple holistic language into a compositional language, thereby replicating how the human language is thought to have evolved. A holistic language is a language where one utterance can be mapped to a concept, such as hunt, storm, or fire. Holistic languages are primitive and complicated sentences such as *a storm is coming from the east soon and, therefore, we have to find shelter*, can be hard to communicate using a holistic language. A compositional language is what humans speak today. This type of language allows multiple concepts to be combined easily. In Gong’s simulation, a holistic rule is a rule that can be mapped to a specific concept, while compositional rules can be combined into different meanings depending on the order they are presented.

The language in Gong’s model is a set of mappings between meanings and utterances, so called M-U mapping. Meanings are represented as predicate-argument structures. The predicates are actions, such as *run*, and the arguments are the objects that the actions are performed upon. When several arguments are used together with a predicate, the order of the arguments decide their role. An agent’s language is stored through three rules: lexical rules which are M-U mappings; syntactic rules which define the compositionality

by describing the order to use lexical rules; and syntactic categories which contain sets of lexical and syntactic rules that are linked to each other. All syntactic and lexical rules have a strength which describes how likely they are to use their M-U mapping successfully. The lexical rules also have an association weight which describes how likely the rule is to be linked to the category it contains. For a more detailed explanation of Gong's compositional language see (Gong, 2011, section 3.2).

Gong's model uses a *random communication framework* (Gong, 2011, section 3.4). Two randomly chosen agents perform many transactions of utterances. The speaker begins by randomly selecting a meaning to produce. The speaker then goes on to activate the lexical rules and the syntactic categories that regulate the rules to form a sentence. The speaker then calculates the set of winning rules, and builds up the sentence using these rules. If the speaker does not have enough knowledge to express the meaning, then random creation of rules occur.

The listener receives the utterance from the speaker together with an *environmental cue*. The cue contains partial information about the environment. The lexical rules that fully or partially match the sentence received are activated.

A candidate set consisting of possible ways to understand the input gotten from the speaker is created. The listener calculates the strength of each candidate set. The rules of the strongest set is then used to interpret the received sentence. If the strength of the strongest set exceeds a certain threshold, the M-U mapping is added to the listener's buffer and positive feedback is given to the speaker. Then both agents reward or retract strength from their rules according to how the communication went. It is never checked whether the agents actually ended up with the same meaning.

### 3.2.2 Results and Discussion

In total, 20 simulations with 6000 communication rounds were performed during each simulation. The simulations were performed with a population size of 10 and with 20 utterances per communication. During the first 100 rounds of communication, many holistic rules were created, but almost no compositional rules were formed. After 200 rounds, the number of holistic rules dropped, while many compositional rules were created. Throughout the first rounds of communication, holistic rules are the main resource for understanding utterances and cues. Very few agents share holistic rules, so a lot of

the communication results in the agents not understanding each other. When compositional rules start appearing, a clear increase in the number of successful communications is observed. Once some compositional rules become shared among the agents, almost all communications are successful. It is also observed that while the rate of successful communications rise, the average amount of meanings that each agent can make increases. That fact that the agents can express many meanings and that the agents understand each other almost every time, indicate that a compositional language has appeared. Gong concludes that “given some general learning abilities, such as pattern extraction and sequential learning, a communal language showing a certain degree of systematicity can emerge in a population of individuals” (Gong, 2011, chapter 6).

### **3.3 Munroe & Cangelosi, 2002**

Munroe and Cangelosi’s simulation (Munroe and Cangelosi, 2002) studies how learning during the lifetime of an individual affects language evolution.

#### **3.3.1 Model**

In their simulation an agent is supposed to identify whether a mushroom is edible or not based on an 18-bit representation of the mushroom’s features. The world contains three sorts of edible mushrooms and three non-edible ones. The three edible mushrooms require their own type of preparation in order for the mushroom to be edible (wash, cut, squash). An agent walks around in a grid-based world for 50 steps. In order to see that an agent could perform in different environments, each agent completed 20 different worlds during one generation. The fitness of an agent is calculated by awarding 1 point every time an agent performs the correct preparation of an edible mushroom. On some of the simulations, it was intended to simulate a learning cost. The learning cost would subtract 1 point if a non-edible mushroom was eaten. A feed forward neural network (FFNN) was used by the agents to classify the mushrooms. An FFNN is the simplest form of artificial neural network (ANN), where all signals are distributed in one direction, forward. Other ANNs, such as recurrent neural networks can send signals backwards in the network creating cycles (Jain et al., 1996).

The simulation has two stages. The first stage lasts for 300 generations with a population

of 100 agents. All agents try to eat as many edible mushrooms by analysing the mushroom's features. After all agents have completed their 20 worlds, the 20 best performing agents become parents of the next generation. The 20 agents make 5 copies of themselves where 10 % of their weights are mutated. The model was set up so that cultural variation could be simulated by applying noise before the 20 best agents were selected.

The second stage lasts for 100 generations and now the agents are allowed to communicate with each other. The 20 best performing agents are carried over to the next generation and act as teachers. 90 % of the time the agents will not have access to the features of the mushroom, but they will receive an input from their parent telling them what kind of mushroom they are observing. The other 10 % of the time the agents have access to both the features of the mushroom and the input from their parents before deciding what action to take. After this, the child is provided with the features of the mushroom and generates its own description of the mushroom. Using backpropagation, the child corrects its output based on the parent's description of the mushroom. Finally, the child imitates the description of the parent by only taking the parent's description as input, trying to reproduce the parent's description as output, and performs backpropagation.

### 3.3.2 Results and Discussion

Ten experiments were ran with different population sizes. At the end of the first stage (300 generations), eight of the experiments solved the game perfectly, by avoiding all non-edible mushrooms and preparing all the edible ones correctly. The agents reached the optimal fitness of 70 at approximately generation 150. Those eight experiments went on to the second stage of the simulation. Seven of these managed to solve the game by combining the occasional input from the environment with the linguistic input from the parents. During the second stage, the agents reached a fitness level of 70 after only 90 generations. Four of these experiments created compositional languages. The language created was compositional because the first symbols were always related to the action and the last symbols were associated with the type of mushroom.

Experiments with the cultural variation and learning cost parameters at different values were performed. The idea behind varying the cultural variation and learning cost was that it is assumed that the Baldwin effect will be strongest when the cultural variation is low and the costs of learning outweigh the benefits.

It was concluded that a Baldwin effect was observed in these experiments. Munroe and Cangelosi found that with a learning cost and a changing environment, some individuals learned the language quicker than others. It was also found that when the learning environment was fixed, specific behaviours got stored in the genome. When the cultural variation was set to 0, i.e. no noise when selecting the fittest agents, even the language structure itself could be built into the agent's genome.

This experiment indicates that the theory that individuals can inherit capabilities that will help learn features more easily might be valid. In order to understand exactly how these mechanisms work more research is needed.

## 3.4 Lekvam, 2014

Lekvam's simulation (Lekvam, 2014) was based on the genetic algorithm (GA) framework, that was explained in section 2.6, and inspired by the model by Lipowska (2011) and the work on social networks done by Quillinan (2006). Lekvam also had a goal that it should be easy to add extensions to his model.

### 3.4.1 Model

Lekvam's model has two processes evolving at the same time, a social structure and a language. The agent's goal is to acquire a social network, which is accomplished by having successful conversations with other agents. The fitness of the agent is a combination of its connections and its age using the formula

$$\text{fitness} = [\exp(0.02 \cdot N_{\text{relations}}) - 1] \exp(-0.05 \cdot t), \quad (3.1)$$

where  $N_{\text{relations}}$  is the number of connections an agent has and  $t$  is the age of the agent. The more relations an agent has, the higher fitness it has, while the agent's age negatively influences its fitness. Lekvam argues that this is a good measurement of fitness because how well an agent can communicate will affect its ability to reproduce.

When an agent acquires  $z$  connections, it will stop reaching out to others, but focus on keeping the connections it has. Other agents with less than  $z$  connections can still contact it, meaning that having more than  $z$  connections is possible. All connections in the social

network are weighted. A successful conversation increases the weight by 1.0, while an unsuccessful one subtracts 0.5 from the weight.

The genome of the agents consists of four genes:

1. **Extraversion** is the probability of searching one layer out in the network for new friends, given that the agent has less than  $z$  connections.
2. **Teach child** is the probability that the parent will be the first to speak to its child. If both of the parents have high values, it is very likely that the parents will speak to the child first.
3. **Lexicon limit** is the maximum size of an agent's vocabulary. If the agent learns a new word, but its vocabulary is full, the lowest weighted word is removed.
4. **Speech ability** is the probability that an agent will not randomly invent a word even though it has chosen one from its vocabulary to utter. If the probability of speech ability is high, there is a high chance that the chosen word is uttered.

When deciding which agents that will go through to the next generation, tournament selection is used. When the set of surviving agents has been selected, they breed. Each agent chooses a partner close to it in the network and then crossover is used to combine their genes. The child has a small probability to mutate.

### 3.4.2 Results and Discussion

The simulations were conducted with 200 generations per simulation, a population of 1225 agents per simulation, and 5 conversations per agent per generation. Seven versions of the simulation were ran. This summary will only present the results from simulation 1, which was the main simulation of the thesis.

It was observed that the first relationship almost all newborn agents had, was with their parents. This is an easy way to make connections for the agents. The agents quickly found out that reaching far out into the social network was beneficial. Since the agents' method of acquiring more than  $z$  connections was to have other agents contact them, reaching far out will increase the chances of other agents reaching out to you. The population very rarely reached full consensus. Most of the time two or three languages remained when the simulation was completed. The final social network reflects this result, as the agents

usually create two or three groups.

Lekvam discusses that his model might have too many flaws to conclude anything certain. Using a naming game might not be the best way of modelling signals being mapped to meaning. He also mentions that the fitness function used in the model might not be realistic enough. However, he does conclude that he believes that the methodology used has “a great potential”.

### 3.5 Comparison of the Four Models

All four language evolution models presented utilise language games. These four language game models are simplifications of how humans communicate while maintaining the essence of why humans communicate. The essence being that in order for the human species to survive, humans need to communicate and understand each other. Language games replicate how an individual is able to learn a language through having each person evolve a personal language based on communication with other people. The language game models made by Lipowska (2011) and Lekvam (2014) both used a naming game. The two other language games presented in the literature study were a spatial naming game and a signalling game.

The spatial naming game and the signalling game incorporated word order regularity because the models studied how a compositional language could evolve, and how a population may reach a consensus over a compositional language. These models became complicated because they studied something as complex as the evolution of a compositional language.

The two naming game models presented were simpler than the spatial naming game and the signalling game because both naming games used a holistic language. The naming game was used to study how individuals in a population could reach a consensus concerning what utterance to use to name an object. These models studied how cultural learning has an effect on language evolution, and to do that, compositional models were not needed.

An agent’s language changed based on conversations in all the models, so how an agent decided whom to converse with had a substantial effect on how a language evolved. Gong’s model (Gong, 2011) used randomly chosen speakers and listeners, and hence, the agents

did not get to decide their listener, which could have made the language evolve unnaturally. However, in the model by Munroe and Cangelosi (2002), the parents spoke to their children, which was an improvement compared to Gong's model, but the agents did not get to choose conversational partners by themselves.

In the model by Lipowska (2011), the agents were set in a lattice world, with newborn agents being placed between their parents and two other agents in the lattice. This meant that every agent was directly connected to four other agents, and like in Munroe and Cangelosi's model an agent would conduct most of its conversations with its parents while they were still alive. It also meant that an agent had a set amount of edges, and lacked the ability to establish a connection to others. The edges were not weighted, meaning that all relationships between the agents were equally strong in this model. Lipowska gave all agents the ability to choose their own conversational partner, but the way a social network was represented had its limitations.

In the model by Lekvam (2014), which was inspired by Lipowska (2011) and Quillinan (2006), each agent could acquire new connections by contacting other agents. In contrast to Lipowska's model, agents could lose edges if several conversations between two agents were unsuccessful in. Each edge was weighted so that not all relationships were equal in this model. Lekvam made a social network in which the agents decided who to converse with, unlike in Gong's model which randomly chose speakers and listeners. It could be someone the agent knew or it could contact a new agent. What was missing from this model was how agents decided whom to converse with, which was done randomly, and that both agents had the same degree of connection towards one another.

In all the models, the fitness of an agent was supposed to be an image of how well that agent performed in a specific environment. In both Lekvam's and Lipowska's model, the fitness of an agent was supposed to reflect the agent's ability to be understood by others. Lipowska used a fitness based on the average weight of the words in an agent's vocabulary and its age. The average weight over the vocabulary described how well an agent was understood, but it did not describe how many agents that understood the agent.

Lekvam used a fitness based on the number of edges of an agent and the agent's age. The number of edges described how many other agents that understood the agent, unlike Lipowska's fitness function. However, the fitness function did not describe to what degree the agents understood each other. In order to do that, the fitness function would have required the weights of the social network to be used. Lekvam's fitness function did not

reward having strong edges. If one agent had ten edges, with weights less than 0.1, while another agent had eight edges with weights more than 0.8, the agent with the strongest weights should be compensated for having strong connections, in terms of an increased fitness. Both models used the age of an agent as well, which represented that the older an agent became, the more likely it was to die.

Munroe and Cangelosi (2002), on the other hand, used a very different fitness function which was based on how well an agent acted in an environment. How well an agent performed was directly correlated to how well the agent understood the input it got from its parents, and so the fitness function reflected how an agent was able to understand other agents. The language was not used as a part of the fitness directly, a separate world was created where an agent's ability to understand language was tested. This was a clever method which tested how an agent learned the language of its parents.



# Chapter 4

## Methodology

The computational model developed during this study is based on Lekvam's model. As opposed to his model, this model add personalities based on Schwartz' 10 values model (Lekvam, 2014; Schwartz, 2012) in order to further investigate the understanding of the effect the evolutionary forces have on language evolution. The model will also present an alternative method for choosing a listener, calculating the weights in the social network, and evaluating the fitness of the agents.

The social networks used in the models presented in the previous chapter were relatively simple. For example, the fitness functions in the previous models highly reward an extrovert strategy, i.e. attempting to meet new people. However, in reality, the majority of conversations in most peoples' lives are conducted with people they already know. A fitness function should ideally allow both strategies to be viable and allow the evolutionary process to display which traits are favourable.

The previously discussed models also assumed that if two agents with the same language conduct a conversation, the conversation was always viewed as a success, and the agents would develop a strong connection over time. However, many people with the same language can not connect with each other due to personal differences, even though they are able to communicate. If an agent does not fit into a group, the agent has two main strategies to improve its situation. Either, the agent can look for another group that is more compatible with the agent's personality, or the agent can slowly alter its personality to fit its original group.

Cultural evolution is about the transferring of language between individuals between indi-

viduals and between generations. In what way language is transferred between individuals is affected by whom one chooses to converse with. How someone decides to interact with others might be affected by that someone’s personality. A social network can be used to visualise how and with whom someone chooses to interact. Personality might have an effect on a person’s social network, and incorporating individual personality into a language game based model and visualising it through a social network might enhance the understanding of the effects cultural learning has on language evolution.

Among the models studied, Lekvam’s model was best suited for these extensions because it already had a social network implemented. The model also use a naming game, which is a language game that is often used when the effects of cultural learning are studied (Lekvam, 2014; Lipowska, 2011; Steels, 2011). Lipowska’s model used a lattice, which does not suit the proposed extension because it does not allow the number of connections an agent has to change. The models made by Gong and Munroe and Cangelosi primarily studied compositional language and not the effects of social networks. Also, Lekvam argued that his model was explicitly made in such a manner that it was suited for extensions.

## 4.1 Genetic Algorithm

The model uses a genetic algorithm which simulates the evolution of language and social networks over time. Parent selection is performed using tournament selection. First, a pool of agents is randomly chosen. Then, these agents are sorted from highest to lowest fitness and given a value corresponding to their ranking, i.e. the agent with the highest fitness has rank 0. Finally, the probability of an agent being chosen to become a parent is

$$P_{\text{parent}} = p(p - 1)^i \quad (4.1)$$

where  $i$  is the agent’s rank in the pool and  $p$  equals the probability of the fittest agent in the pool being chosen as the parent. A newborn agent’s genotype is made from the genotypes of its two parents using crossover. The newborn agent’s genotype has a small probability of mutation, which is conducted by randomly altering one value in the genotype.

At the end of each generation, turnover selection is performed by randomly choosing  $k\%$  of the population into a pool. Then, the  $n\%$  fittest agents of that pool survive to the next generation. One generation in the simulation does not equal one generation in real life.

## 4.2 Genome

Personalities are added into the model in the form of an agent’s genome. The genotype of an agent consists of ten numbers, in the interval between 0 and 100, each representing one of the ten values from Schwartz’ model of basic human values. The values cannot be randomly initiated because some of the values depend on each other as some of the values describe relatively equal goals. If two contradicting values both have a high value, the goals of that person would contradict each other. So, in order to ensure the validity of the relations between the ten values, not all values can be created randomly. Schwartz categorises his ten values into four groups, as seen in Figure 2.2. Values within a group are related and should have numerical values close to each other, while values on the opposite side of the circle should have opposite numerical values. In this model, the validity of the relations between the values is kept by randomly setting the Self-direction and Achievement value. Values in the same quadrant are set by using the following formula

$$V_i = \frac{\text{random}(0, 20)}{100} \cdot V_{i-1}, \quad (4.2)$$

where  $V_i$  is the number associated with value  $i$ ,  $V_{i-1}$  is the number associated with the value counter clockwise of  $V_i$  in the circle, and  $\text{random}(0, 20)$  is a random integer  $i$ ,  $0 \leq r \leq 20$ . Security and Benevolence are set based on the opposite value using the formula

$$V_i = \frac{\text{random}(1, 20)}{100} \cdot (100 - V_{\text{opposite}}). \quad (4.3)$$

The other values in the quadrants with Security or Benevolence are set based on Equation (4.2). Hedonism is split between two quadrants and it is therefore calculated based on the two values next to it in Figure 2.2 through this formula

$$V_i = \frac{\text{random}(1, 20)}{100} \cdot \frac{V_{i-1} + V_{i+1}}{2}. \quad (4.4)$$

The values of the genome is then normalised. The resulting genome is representing the individual’s personality and phenotype. Each agent has a vocabulary similar to the vocabulary used in Lekvam’s model. The vocabulary consists of the utterances an agent has heard or created. Each utterance has a corresponding weight describing how successful the utterance has performed in past conversations.

### 4.3 Fitness Function

The fitness function is designed to both represent how well an agent communicates with the members of the community, and how the agent's personality fits in with the group. An agent's ability to communicate is projected into the social network because the weights are adjusted based on an agent's ability to be understood. The personality of the agents determines how good it is at learning a language, how extrovert it is, and whether it cares about its children. The fitness function also takes into account the age,  $t$ , of the agent. The fitness function is

$$\text{Fitness}(j) = \left[ \exp \left( \alpha \sum_{i=1}^N W_{ij} \right) - 1 \right] \left[ \exp \left( \beta N_{W_{ij} \geq \Theta} \right) - 1 \right] \gamma \exp(-0.05t) \quad (4.5)$$

where  $W_{ij}$  is the weight of the connection between agent  $i$  and agent  $j$ ,  $t$  is the number of generations it has lived,  $N$  is the total number of connections of agent  $j$ ,  $N_{W_{ij} \geq \Theta}$  is the total number of connections that agent  $j$  has that are stronger than the threshold  $\Theta$ , and  $\alpha$ ,  $\beta$ , and  $\gamma$  are constants adjusting the three factors of the fitness function. The first factor takes into account how strongly an agent connects to others by calculating the sum of all its weights. The second factor calculates how many agents with a connection strength stronger than the threshold  $\Theta$  agent  $j$  has. The third factor provides an agent with less fitness the older it becomes. These three factors combined display the agent's fitness in the society. They take into account the number of connections the agent has, the strength of the connections, and the age of the agent.

### 4.4 Social Network and Dialogues

The social network in the model is represented by a weighted discrete graph, where the vertices represent the agents, and the edges represent the connections between the agents. If a dialogue is successful, i.e. the listeners has the uttered word in its vocabulary, the weight of the edge between the two agents is increased. If the dialogue was unsuccessful, i.e. the listener did not have the uttered word in its vocabulary, the weight is decreased. An edge gets added a weight of 1 if the dialogue is successful and 0.5 is subtracted if the dialogue was unsuccessful. This method of calculating the weight allows the language of the agent to affect the weight of the edge.

When an agent chooses a listener, the agent takes into account both its social network and its personality. An agent can either be extrovert or introvert, which is reflected in its genome. The extrovert trait is connected to a high Achievement value, which is found in people who seek recognition for their achievements. Many people enjoy sharing achievements with others. For some, the achievement has greater value if many knows about it. The introvert trait is connected to a high Conformity value, which is found in people who are very aware of the consequences of their actions and might be careful in social settings because of it. Each of these values are taken from an agent's genome and is therefore between 0 and 100. Based on these values, the following formula describes the probability of an agent reaching out to agents that it does not already have a connection with

$$P = \left( \frac{\sum \text{extrovert values} + (100 - \sum \text{introvert values})}{200} C \right) / 100 \quad (4.6)$$

where  $C$  is a constant in the range  $0 < C \leq 1$ , which forces all individuals to have a tendency towards introvertedness. People tend to have more conversations with people they already know, and hence skewing the probability towards introvertedness is reasonable. The first conversations a newborn agent conducts, might be with its parents. The probability of having dialogues with parents is equal to the normalised value of the Security value in the genome. Each parent will conduct three dialogues each and the newborn will function as the listener in all these conversations. If the speaking agent  $i$  decides to act introvertly, the probability of contacting one specific agent  $j$  is related to the weight between the two agents by the following relation

$$P_{i,j} = \frac{W_{ij}}{\sum_{k=\{\text{connected to } i\}} W_{ik}}, \quad (4.7)$$

where  $W_{ij}$  is the weight of an edge and the summation is performed on all agents connected to agent  $i$ . If the agent decides to be extrovert, it randomly chooses an agent to contact from the population it is not already connected to. If the parents do not conduct conversations with their newborn agent, the agent will have its for first dialogue with whomever that contacts it first from the population.

Agents also alter their personality based on who they communicate with. An agent's personality is updated by calculating the personality change rate (pcr) for all the values in the genome. Then each value is updated through the formula  $\text{new}_{\text{personality}} = \text{old}_{\text{personality}} + \text{old}_{\text{personality}} \cdot \text{pcr}$ . After a conversation, both agents alter their values using the following

formula

$$\text{Personality change rate} = C \frac{\Delta V_i S}{t}, \quad (4.8)$$

where  $\Delta V_i$  is the difference of a specific value  $i$  in the two agents' genomes,  $S$  equals one if the conversation was successful and zero if it was unsuccessful, which ensures that the personality is not changed if a conversation is unsuccessful,  $t$  is the number of generations the agent has lived, and  $C$  is a constant,  $0 < C \leq 1$ .  $C$  makes the pcr have less of an impact the smaller  $C$  is. After the pcr is calculated for all values in the genome, the genome is normalised again.

A conversation is conducted by the speaker uttering a word  $i$  from its vocabulary. The probability of a specific word from the vocabulary being uttered is given by

$$P_i = \frac{W_i}{\sum_{j=1}^N W_j} \quad (4.9)$$

where  $W_i$  is the weight of word  $i$  in the speaking agent's vocabulary, and  $\sum_{j=1}^N W_j$  is the sum of the weights of all words in the agent's vocabulary. If the speaking agent has an empty vocabulary, it invents a new word by randomly deciding the length of the new word and randomly adding that many letters to the word.

If the listener has the utterance in its vocabulary, the conversation is deemed as a success, and the weight of the edge between the agents in the social network is increased, the personalities of both the listener and speaker are altered, and the weight of the word in the vocabularies of both agents are increased. If the listener did not have the word in its vocabulary, the conversation is deemed a failure, and the weight of the edge between the agents in the social network decreases. The speaking agent decreases the weight of the word in its vocabulary, while the listener adds the word into its vocabulary. None of the agents' personalities are altered when the conversation is unsuccessful.

Schwartz found that Self-direction, Benevolence, and Universalism is the three values most people value the highest, while Power, Stimulation, and Tradition are the least important (Schwartz, 1992). The weight of a word in an agent's vocabulary is adjusted after a conversation based on the *learning rate*, which is a value between 0 and 1. The learning rate is calculated by adding the normalised values of the three positive values and subtracting the normalised values of three negative ones. If the sum is less than 0, the learning rate is set to 0. This should allow agents with a strong set of values to learn a language through fewer dialogues than agents with a weak set of values.

# Chapter 5

## Results

This chapter will present the results gathered through seven experiments using the computational model presented in Chapter 4.

### 5.1 Parameters

The parameters used in the main experiment can be seen in Table 5.1 in addition to the corresponding base value and a short explanation. In a given experiment, the value of a parameter can be assumed to be the base value associated with it, unless no other value is specified. Unless otherwise stated, the value of the parameters in the experiments are set to the associated base value.

### 5.2 Experiments

This section will present seven experiments. The experiments were conducted in the order in which they are presented. The first experiment represents the main results and serves as the foundation for the following experiments. The second and third experiments tested what happens when the population decreases and increases respectively. The fourth experiment was performed with an increased amount of dialogues per generation. The results from the third experiment led to a hypothesis which was tested in the fifth experiment through lowering the probability of the occurrence of parent-child dialogues.

**Table 5.1:** The parameters that is used in the model and their corresponding default values.

Parameter	Value	Description
a	1000	Number of agents in the population.
g	100	Total number of generations the simulation will last for.
d	1	Number of dialogues per agent per generation.
k	0.9	$(k * 100)\%$ of the population is randomly chosen into the survival selection pool.
n	0.5	The $(n * 100)\%$ fittest agents in the pool survive to the next generation.
m	0.02	Probability of mutation.
c	0.5	The constant which affects how extrovert agents can act.
$\delta$	3	Constant related to the probability of a child conducting its first dialogues with its parents.
$\epsilon$	0.2	Probability in the parent selection pool is $(1 - \epsilon)$ .
$\alpha$	0.2	Constant related to the fitness function part which concerns an agent's weights.
$\beta$	0.12	Constant related to the fitness function part which concerns an agent's edges.
$\gamma$	-0.05	Constant related to the fitness function part which concerns an agent's age.
$\Theta$	0.5	The minimal strength of the edge for it to be counted in the part of the fitness function that concerns the agent's edges.

The hypothesis will be presented in the discussion of the results. Experiment six tested how altering the turnover rate affects the evolutionary processes in this model. The last experiment added a probability for agents to invent a new word during a dialogue even though their vocabulary was not empty.

All results for experiment 1, 2, and 3 are included. For the other experiments, only the results relevant results are included in this chapter. The omitted results can be found in Appendix A.

### 5.2.1 Experiment 1 – Main Results

The main experiment was conducted with the default parameters shown in Table 5.1. The default parameters mean that the population is of average size, an amount of generations that ensures convergence of the results, few dialogues per generation, a high turnover with low variance, and a high probability of conducting parent-child dialogues. See Figures 5.1 and 5.2 for the graphs related to this experiment.


The average fitness of the agents greatly increased during the first 40 generations, and then the average fitness stabilised at about 0.85, as seen in Figure 5.1(a). The average degree of the agents also increased rapidly during the first 40 generations, eventually stabilising at an average of about 10 edges per agent, as seen in Figure 5.1(b).

The speaking to parents probability, learning rate, and extrovert probability evolved, can be seen in Figure 5.1(c). The easiest method for a newborn agent to acquire connections was through dialogues with its parents, and the probability of that happening was slowly increasing during the entire simulation. The learning rate increased during the first 20 generations. After the first 20 generations, the learning rate slowly increased during the remainder of the simulation. The extrovert probability increased during the first three generations, and then decreased to about 5%, indicating that agents who focused on gaining strong relationships gained a better fitness than extrovert agents with many weak relationships.


The graphs describing the number of unique highest ranked words in the vocabularies, as seen in Figure 5.1(d), the average number of words in a vocabulary, as seen in Figure 5.1(e), and the percentage of successful dialogues, as seen in Figure 5.1(f), all tell the same story: At the beginning of the simulation there existed many words in the population, all agents had large vocabularies, and few dialogues were successful. The number of

unique words decreased substantially during the first 10 generations, then at generation 62, the population reached consensus. Consensus means that all agents had the same word as their highest weighted word. The average vocabulary size increased during the first generations, and then decreased until generation 60, ending at an average of slightly more than one. The graph which displays the percentage of successful dialogues reached 70% at generation 13, then the percentage rose with 0.23 percentage points per generation for the remainder of the simulation, ending at 90%.

The graphical representation of the social network only include the agents with a connection, as seen in Figures 5.2(a)–(d). This explains why the number of nodes is not equal in all four figures, and why there is a discrepancy between the social network and the graph displaying the average degree per agent. The farther an agent is away from the centre of the social network, the fewer edges it has. The nodes closest to the centre are the ones with the most edges. At generation 5, most of the nodes had two or three connections each, as seen in Figure 5.2(a). By generation 20, the number of unique words had gone down and many of the agents had acquired several connections, as seen in Figure 5.2(b). The network at generation 20 was more polarised than it was at generation 5. Some agents had many connections, while a lot of agents had only one or two connections. At genera-




**Figure 5.1:** Results from experiment 1: (a) the average fitness; (b) the average degree; (c) three graphs that show the evolution of the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly; (d) the number of unique highest ranked words in the population; (e) the average vocabulary size; and (f) successful dialogues divided by total number of dialogues.


**Figure 5.1:** *(continued)*

tion 40, the social network had become even more polarised, as seen in Figure 5.2(c). The agents in the centre had gotten more connections than at the beginning of the simulation, while the agents outside of the centre still had one or two edges. At generation 100 most of the population with connections had several connections, as can be seen by the high density of edges at the centre of the figure, as seen in Figure 5.2(d). There were still a select few agents with only one or two edges, but those were by far in a minority.


(a)




(b)



(c)



(d)




**Figure 5.2:** Snapshots of the social network of experiment 1 at generation (a) 5, (b) 20, (c) 40, and (d) 100.

### 5.2.2 Experiment 2 – Small Population


This experiment was conducted with a population of 100, which is 10 times smaller than in experiment 1, otherwise the parameters were the same. See Figures 5.3 and 5.4 for the figures related to this experiment. The variance seems to be higher with fewer agents, causing the average fitness and degree to vary more from one generation to another, as seen in Figures 5.3(a) and 5.3(b). The average fitness seems to stabilise at about 0.85 just like in experiment 1, but after generation 30 which was earlier than in experiment 1.

The number of unique highest weighted words reached 1 after 15 generations, as seen in Figure 5.3(d), and the percentage of successful dialogues reached 90% after 15 generations, as seen in Figure 5.3(f). Both converged on the same values as in experiment 1 but they converged at an earlier stage.


At generation 5, the social network consisted of a few agents with two or three edges per agent, as seen in Figure 5.4(a). Up until generation 40, the social network became more and more dense for each generation that went by, as seen in Figure 5.4(c). Between generation 40 and 100, there seemed to be little change in the structure of the social network, as seen in Figure 5.4(d).




**Figure 5.3:** Results from experiment 2: (a) the average fitness; (b) the average degree; (c) three graphs that show the evolution of the the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly; (d) the number of unique highest ranked words in the population; (e) the average vocabulary size; and (f) successful dialogues divided by total number of dialogues.

**Figure 5.3:** *(continued)*


(a)




(b)



(c)



(d)




**Figure 5.4:** Snapshots of the social network of experiment 2 at generation (a) 5, (b) 20, (c) 40, and (d) 100.

### 5.2.3 Experiment 3 – Large Population


This experiment was performed with a population of 10000, which is 100 times larger than the population in experiment 1. The figures presenting the results from this experiment can be seen in Figures 5.5 and 5.6.

The results show that the average fitness and degree of the agents was lower than in the results from experiment 1, as seen in Figures 5.5(a) and 5.5(b). When there were many agents, the simulation required more time to reach consensus, as seen in Figure 5.5(d). In fact, after 100 generations, the agents had not reached an agreement upon one word. Because of the computational cost of the large population size the simulation was not ran for more than 100 generations. When the population struggled to reach consensus, communication was more often being conducted unsuccessfully, as seen in Figure 5.5(f).


The difference in the social network compared to experiment 1 is that the agents were divided into one main group and several other groups of two to three agents for the first generations, as seen in Figure 5.6. After generation 20 almost all of the small groups had disappeared and after 100 generations there none left.




**Figure 5.5:** Results from experiment 3: (a) the average fitness; (b) the average degree; (c) three graphs that show the evolution of the the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly; (d) the number of unique highest ranked words in the population; (e) the average vocabulary size; and (f) successful dialogues divided by total number of dialogues.

**Figure 5.5:** *(continued)*


(a)




(b)



(c)



(d)




**Figure 5.6:** Snapshots of the social network of experiment 3 at generation (a) 5, (b) 10, (c) 40, and (d) 100.

### 5.2.4 Experiment 4 – Dialogues


This experiment was conducted with  $d = 5$ , meaning that each generation conducted five times as many dialogues as in experiment 1. The results from this experiment can be seen in Figure 5.7.

The increase in number of dialogues per agent per generation ensures that each agent has more and stronger connections leading to a higher average fitness, as seen in Figure 5.7(a). The average degree naturally increased when five times as many dialogues were conducted, as seen in Figure 5.7(b). The probability of acting extrovertly stagnated at about 16%, which is almost three times more than experiment 1, as seen in Figure 5.7(c). The simulation reached consensus at about the same time as the main experiment, as seen in Figure 5.7(d).

The social network evolved similarly as in the main experiment, except that the agents on average had more edges. Since the social network figures for this experiment are not very informative, they are only included in Appendix A.



**Figure 5.7:** Results from experiment 4: (a) the average fitness; (b) the average degree; (c) Three graphs that show the evolution of the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly; (d) the number of unique highest ranked words in the population; (e) the average vocabulary size; and (f) successful dialogues divided by total number of dialogues.


**Figure 5.7:** *(continued)*

### 5.2.5 Experiment 5 – Less Parent-Child Dialogues


Experiment 5 was conducted with the constant  $\delta = 1$ .  $\delta$  is related to the probability of a child having its first dialogues with its parents. The result from experiment 5 can be seen in Figure 5.8. The simulation was ran for 150 generations, i.e.,  $g = 150$ . The reason for running 50 more generations than in the other experiments was that it did not reach consensus after 100 generations and knowing how this simulation evolved up until consensus might be interesting.

The average fitness increased slowly during the entire simulation, as seen in Figure 5.8(a). After 10 generations, the average fitness was 0.23 in experiment 5 compared to 0.45 in experiment 1. Experiment 5 converged on an average fitness of about the same as experiment 1 but it required more generations to reach that convergence. The average degree reached 4 at generation 70, and during the next 80 generations the average only increased by 0.5, as seen in Figure 5.8(b).

The simulation reached consensus at generation 127, as seen in Figure 5.8(d). The average vocabulary size increased during the first 25 generations, and then decreased and ended



**Figure 5.8:** Results from experiment 5: (a) the average fitness; (b) the average degree; (c) three graphs that show the evolution of the the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly; (d) the number of unique highest ranked words in the population; (e) the average vocabulary size; and (f) successful dialogues divided by total number of dialogues.



**Figure 5.8:** (continued)


at 1.04, as seen in Figure 5.8(e). The percentage of successful dialogues reached 60% at generation 6, and then gradually increased during the rest of the simulation ending at 93%, as seen in Figure 5.8(f).

### 5.2.6 Experiment 6 – Turnover


Experiment 6 was performed in two parts, part 6a with  $k = 0.95$  and  $n = 0.7$  and part 6b with  $k = 0.95$  and  $n = 0.85$ , compared to  $k = 0.9$  and  $n = 0.5$  in the other experiments. This means that a larger fraction ( $n$ ) of the population was chosen for survival and from a larger selection pool ( $k$ ). The results from experiment 6 can be seen in Figure 5.9.

Experiment 6a was run for 150 generations because the population was far away from reaching consensus after 100 generations but it still did not reach consensus. The figures that display the genes and the evolution of the social network have not been included because they were similar to the ones from experiment 1, but they can be found in Appendix A. Both the average fitness and the average degree stabilised after about 75 generations, as can be seen in the Figures 5.8(a) and 5.8(b). The final value of the average degree stabilised at 11 and the average fitness at 0.88. The number of unique words decreased a lot during the first 20 generations, as can be seen in Figure 5.9(e). The simulation reached two unique words at generation 125, and also ended with two unique words after all 150 generations.

Experiment 6b was conducted with  $k = 0.95$  and  $n = 0.85$  which allowed a larger portion of the fittest agents to survive each generation, giving the simulation a lower selection pressure. Experiment 6b was conducted in order to see what happened when the selection



**Figure 5.9:** Results from experiment 6: (a) the average fitness; (b) the average degree; (c) the average vocabulary size; (d) successful dialogues divided by total number of dialogues; and (e) the number of unique highest ranked words in the population of experiment 6a and 6b.



**Figure 5.9:** (continued)


pressure was lower than in experiment 6a. How experiment 6a and 6b evolved towards consensus can be seen in Figure 5.9(e). Experiment 6b was ran for 200 generations in order for it to reach consensus. Consensus was reached at generation 178. At generation 125, when experiment 6a had reached two words, simulation 6b had six unique words.

### 5.2.7 Experiment 7 – Making up New Words


Experiment 7 was conducted in two parts, both with an added probability for agents to invent a new word even though they had at least one word in their vocabulary. For the other experiments new words was only invented if the agent did not have any words in its vocabulary, i.e. zero probability of inventing a new word when its vocabulary was not empty. The results from experiment 7 can be seen in Figure 5.10.

Experiment 7a was performed with the probability of inventing a new word set to 25%. The average fitness stabilised at 0.37, and the average degree stabilised at about 6.5, as seen in Figures 5.9(a) and 5.9(b). Both of these values were considerably lower than in experiment 1. The average vocabulary size was 6 and the percentage of successful dialogues quickly reached 0.5, but it did not slowly increase like in the previous experiments, as seen in Figures 5.10(d) and 5.10(e). The number of unique highest ranked words evolved similarly to experiment 1, as seen in Figure 5.10(f). The population reached consensus at generation 68, only six generations after experiment 1.

Experiment 7b was performed with the probability of inventing a word increased to 45%.



**Figure 5.10:** Results from experiment 7: (a) the average fitness; (b) the average degree; (c) three graphs that show the evolution of the the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly; (d) the average vocabulary size; (e) successful dialogues divided by total number of dialogues; and (f) the number of unique highest ranked words in the population of experiment 7a and 7b.



**Figure 5.10:** (continued)

See figure 5.10(f) for the figure related to this experiment. The other graphs have not been included because the purpose of this experiment was to see how the evolution towards consensus changed with a higher probability of inventing a word. The number of unique words quickly dropped to about five, but from five it struggled to actually reach consensus. Experiment 7b reached consensus at generation 73, but that only lasted a few generations. For the remainder of the simulation, the number of unique words varied from one to four.

# Chapter 6

## Discussion

This chapter will compare and discuss the results of all the experiments reported in Chapter 5. Some of the graphs from different experiments in relation to each other can be seen in Figures 6.1–6.7 on Pages 59 to 66.

In experiment 1, the evolution favoured agents that acted extrovertly during the first generations. This happened because most dialogues were unsuccessful, and hence, having many weak connections was better than having few weak connections. As the number of unique words decreased, dialogues became more likely to be successful, and a successful dialogue increased the involved agents' fitness. An extrovert agent had many edges, which increased its fitness. On the other hand, those edges had a lower average weight than an agent with fewer edges. A low average weight would decrease the agent's fitness. So there was a trade-off between acting introvertly and having fewer edges with high weights, and acting extrovertly and having more edges with lower weights. The population evolved towards favouring introvert characteristics. The new fitness function was designed so that both extrovert and introvert strategies have advantages and disadvantages. Considering that an extrovert strategy is optimal in the early stages and later on an introvert strategy is favourable, it seems that the fitness function gives room for both strategies to be viable. Which strategy is optimal depends on the environment and not the fitness function itself.


All agents in the population evolved into one personality because one agent ended up in a local maximum and spread its genes to the rest of the population through children and changing personalities based on dialogues. The most optimal personality found favoured an introvert strategy to gain a high weight on its edges, speaking to its newborns to gain new edges with a relatively high weights, and a relatively high learning ability to

learn a language more quickly. It is clear that learning a language becomes easy when few languages exist. Therefore the focus shifts towards gaining edges and strong weights on those edges because most dialogues will be successful anyway. The personality based genome was successful when it comes to depicting the important values to have in this model, but it did not depict the large variance between personalities that exists in real life. In reality, a population has some traits that will always be viewed as favourable, and several traits that have both positive and negative sides. Because evolutionary algorithms are intended to search for optimal solutions in a complex search space, it is natural that the variance in personality was not represented properly.

Agents evolved towards higher learning rate, most likely because a high learning rate made the language easier to learn. When the population had almost reached consensus, and the average vocabulary size was close to 1.0, the average learning rate stopped increasing. This might have happened because the pressure on being able to learn one language in a pool of many was gone. When only a few languages were being used, the pressure on having a high learning rate seemed to drop.

During the first ten generations of experiment 1, the number of unique highest ranked words decreased. The agents that had edges with the highest weights survived the first generations. Those agents had a high fitness, and hence, they were probable to be chosen as parents several times during the parent selection phase. The agents that got selected as parents several times and had a high probability of speaking to their children acquired more edges with strong weights, which increased their fitness and likelihood of surviving to the next generation. This cycle kept a select few agents alive for the first generations, allowing them to spread their language to many others. This is how the number of unique words dropped quickly during the first generations of the simulation.

The average vocabulary size increased during the first seven generations of experiment 1, then it quickly decreased during the next three generations, and for the remainder of the simulation the average vocabulary length decreased slowly. During the first generations most agents did not have one word with a much higher weight than the others in the vocabulary, as seen in Figure 6.1 on Page 59. Because the words had similar weights, the probability of choosing the same word in each conversation was small, leading to words being spread around. In order to gain edges during the first generation, the agents favoured acting extrovertly. The combination of acting extrovertly and not choosing the same word in every conversation lead to the agents learning and teaching new words, which made the average vocabulary size increase. When a few words in the vocabularies




**Figure 6.1:** Comparing the average vocabulary size per generation for experiment number 1, 2, 3, 5, and 7b.

had gotten a significantly higher weight than the others, those words were the only words that the newborn agents learned. The process where newborns only picked up the most successful words of their parents quickly reduced the number of unique words and the average vocabulary length. After the first ten generations, the evolution towards an average vocabulary length of 1.0 continued slowly because most words were highly weighted in someone’s vocabulary.

The social network of experiment 1 became denser by each generation that passed up until generation 40 which was when the average degree of the agents stabilised, as shown in Figure 5.2 on Page 42. Even though the agents mainly acted introvertly, the social network was not divided into groups, like in Lekvam’s experiment (Lekvam, 2014, Section 5.1). The reason why this happened probably was because an agent could contact any other agent when it decided to act extrovertly. This meant that the entire population existed without any physical distance from one another. A community of individuals living so close to each other might have a social network similar to the one from this simulation. The network consisted of several circular layers, where the innermost layer had agents with the most edges, and the outer layer had the agents with the fewest edges.


Experiment 2 reached consensus after fewer generations than experiment 1, as seen in Figure 6.2 on Page 60. This happened because there were fewer words to begin with




**Figure 6.2:** Comparing the number of unique highest ranked words per generation for experiment 1, 2, 3, 4, 5, 6a, 6b, 7, and 7b.

and fewer agents to teach one word to. An agent is more likely to become a parent in a small population, which leads to more newborns learning the same word from an earlier generation. The tendencies concerning how the language evolved were similar, only that it all happened faster with higher variance. In experiment 2 the average degree and fitness had especially big variations from one generation to another compared to experiment 1. See Figure 6.3 on Page 61 for the fitness and Figure 6.4 on Page 61 for the degree. This is natural considering that there were randomness involved in the form of choosing the agents involved in dialogues, and making children. When fewer dialogues and fewer children were made per generation, the variance increased.

Experiment 3 should have less variance and require more generations to reach consensus than experiment 1 if this reasoning from experiment 2 was correct, and it did indeed have less variance and reached consensus later. The evolution began similarly to the two previous experiments, where the population's average degree, fitness and unique words evolved quickly. See Figure 6.3 on Page 61 for the fitness, Figure 6.4 on Page 61 for the degree, and Figure 6.2 on Page 60 for the unique highest words. The population stabilised on a high number of unique words and a fitness of 0.5. At generation 60, the parents taught their children a few words and the average vocabulary size started decreasing, while the average fitness increased, as seen in Figure 6.1 on Page 59. This is exactly what happened in experiment 1, only that the population in experiment 3 required more time to reach



**Figure 6.3:** Comparing the fitness per generation for experiment 1, 2, 3, 4, 6a, 6b, and 7b.



**Figure 6.4:** Comparing the number of unique highest ranked words per generation for experiment 1, 2, 3, 4, 5, 6a, 6b, and 7b.

a state where some words had a significantly higher weight than the rest of the words because of the population's size.

The social network of experiment 3 had small groups of two agents in addition to one big main group for many generations. Those small groups eventually died out or joined the main group because the agents in the small groups were unable to get a large enough fitness on their own. The evolution towards one social network happened more slowly than in the two previous experiments because the newborns in the main group learned many different words from their parents, giving the agents in the biggest group a lower average fitness than the biggest group in the two previous experiments. This trend will probably be seen if the population size was to be increased even more.

The number of dialogues per generation was quintupled in experiment 4. The average fitness was higher than in experiment 1, as seen in Figure 6.3 on Page 61, because the average weight of an edge increased a lot due to the high number of introvert dialogues. Although four times as many dialogues were performed, most of them were still introvert, so the average degree only increased by 50% compared to experiment 1, as seen in Figure 6.4 on Page 61.

More newborn agents died in experiment 4 compared to three previous experiments. The newborn agents died because their parents had gotten a high fitness, which the newborn agents could not compete with in the selection process. The parents' fitness was high because they had conducted so many dialogues during their lifetime that they had several highly weighted connections.

It would be natural to assume that more conversations per generation would allow the population to reach consensus faster than in experiment 1, but it did not, as seen Figure 6.2 on Page 60. This simulation required the same number of generations to reach consensus as experiment 1. This result leads to the conclusion that the dialogues has a small impact on the agents' vocabulary. The only other method for a newborn agent to learn a language in this model is through initial conversations with its parents, and experiment 4 indicates that parent-child dialogues might be important for the model to reach consensus. Parent-child dialogues could also be an important factor on the *ontogenetic* time scale, which concerns how an individual learns language.

The fifth experiment tested the hypothesis from experiment 4 by decreasing the probability of parents speaking to their offsprings to three times less than in experiment 1. If the hypothesis was to be true, this experiment should have required more generations than


experiment 1 to reach consensus, and it did. Throughout the experiments performed thus far, the average vocabulary length decreasing seemed to be an indicator of that the population was approaching consensus. In experiment 1, the average vocabulary size started decreasing at generation 8, but in experiment 5, the average vocabulary size started decreasing at generation 25, as seen in Figure 6.1 on Page 59. Experiment 1 had reached consensus at generation 61, while this experiment reached consensus at generation 127, as seen in Figure 6.2 on Page 60. This result supports the hypothesis that this computational model's most important factor for reaching consensus is the parents speaking to their children.

The result also indicates that the first dialogues have a substantial impact on the *ontogenic* time scale and are essential for the newborn agents' language in this computational model. The language of the children in the model is defined through their first dialogues, and after that it is very hard to learn a new language as fluently as its native language, just like humans who tend to struggle when attempting to learn a language at an adult age.

The average degree in experiment 5 was significantly lower than in experiment 1 because a lot of newborns did not get the initial dialogues with their parents, as seen in Figure 6.4 on Page 61. The parent-child dialogues were reduced, but the experiment had the same number of dialogues, and therefore the same number of extrovert dialogues. This indicates that the average degree should be roughly 1/3 of the main experiment. The main experiment stabilised on an average degree of about 10, and this experiment stabilised on an average of about 4, which is 2.5 times fewer edges as experiment 1.


The social network in experiment 5 evolved through a select few agents that had a much higher fitness and a much higher probability of speaking to their children than the other agents in the population. Those parents connected with many children and created the fan-like shapes, as seen in Figure 6.5 on Page 64.

The social network in experiment 5 was significantly less dense than experiment 1 during the entire simulation. For example, at generation 100, the social network in experiment 5 had evolved into a network similar to the social network at generation 5 from experiment 1, as seen in Figure 6.6 on Page 64. The reason for the social network being less dense in this experiment is that when agents do not connect with both their parents, they get a very low fitness. The probability of an agent being chosen as a speaker is related to its fitness compared to the other agents' fitness. When many agents have a very low fitness,




**Figure 6.5:** Social network of experiment 5 at generation 10. Three of the fan-shapes can be seen: one on the left and two on the bottom right.

(a)



(b)



**Figure 6.6:** Comparison of (a) experiment 1 at generation 5 and (b) experiment 5 at generation 100.

the few agents with a high fitness will often be chosen as speakers, which leads to a social network with fewer and stronger edges than in experiment 1.

If the number of new agents being born per generation was decreased, one could expect to see similar results as the previous experiment, because it would lead to fewer newborn agents conducting initial dialogues with their parents. Experiment 6 tested this through performing two experiments with increased the value of both  $k$  and  $n$ . The results show that the population struggled to reach consensus; in fact, within the 150 generations ran in experiment 6a consensus was not reached and experiment 6b reached consensus at generation 178, as seen in Figure 6.2 on Page 60. The result indicates that it is not


only important to have many newborn agents being taught a consistent language in the early years, it is also important to have a certain number of agents being born in each generation in order for the evolution to move forward.

In experiment 6b an even larger portion of the population survived each generation. This means that the selection pressure on the population was decreased and the importance of the agents' ability to learn a language was decreased. When the selection pressure on the agents' languages was decreased, the evolutionary forces at the *phylogenetic* and *glossogenetic* time scale had less of an impact; see Section 2.1 for a description of the evolutionary forces. The low selection pressure lead to the language evolving slower than in experiment 6a, as seen in Figure 6.2 on Page 60. The results from experiment 6 indicate that when the selection pressure on the language is small enough, the traits that concerns language become less important and the evolution of these traits happens slowly. In the experiments where the selection pressure was high, the population reached consensus, but when the selection pressure was decreased the evolution towards consensus slowed down. This indicates that the Baldwin effect affects the *phylogenetic* and *glossogenetic* forces of language evolution.

The average fitness and degree of the agents was higher when a larger part of the population survived to the next generation; see Figure 6.3 on Page 61 for the fitness and Figure 6.4 on Page 61 for the degree. The agents who survive generally are among the fittest of the current population, and so when a larger portion of the agents with a high fitness survive, the average fitness becomes higher. The average degree is directly related to the fitness function, and hence, it is natural that that is higher as well.

In experiment 7a the agents could make up new words with a probability of 25%. One would expect that this made the percentage of successful dialogues to be less and the average vocabulary size to be bigger than in experiment 1. As expected, the average fitness, degree, and percentage of successful dialogues in experiment 7a were lower than in experiment 1, and the average vocabulary size was a lot higher. For the average vocabulary size see Figure 6.1 on Page 59, for the fitness see Figure 6.3 on Page 61, for the degree see Figure 6.4 on Page 61, and for the rate of successful dialogues see Figure 6.7 on Page 66.

One might also expect that the population would have required more generations to reach consensus when more utterances existed. However, this did not significantly affect the evolution towards consensus, as seen in Figure 6.2 on Page 60. The previous simulations



**Figure 6.7:** Comparing the rate of successful dialogues per generation for experiment 1 and 7b.

that had struggled to reach consensus had forced the population to conduct less parent-child dialogues. This simulation actually did not force less parent-child dialogues. In fact, the probability of a parent choosing to utter a word from its vocabulary in at least two out of the three dialogues it has with its newborn agent is 0.84. The probability of at least one parent uttering at least two words from its vocabulary is  $1 - (1 - 0.84)^2 = 0.97$ . So the parent-child dialogues were not particularly affected by this. Also, the added number of words had very little impact, because they were weighted very low compared to the words that had survived for several generations.

In experiment 7b the probability of inventing a word was increased to 45%, and the number of unique words varied a lot, as seen in Figure 6.2 on Page 60. The probability of a parent choosing to utter a word from its vocabulary in at least two out of the three dialogues it has with its newborn agent is 0.57. The probability of at least one parent uttering at least two words from its vocabulary is 0.82. It seems that with a probability for inventing a word that high, several of the newborn agents were taught several words from their parents, leaving the newborns without a consistent language to teach the next generation. When that happened to a large enough portion of the population, new words actually appeared in smaller groups for a few generations until the branch keeping the word alive died. The appearing and disappearing of these smaller languages happened regularly. After 73 generations this simulation reached consensus. However, new words

regularly appeared and disappeared, allowing the number of highest ranked words to vary between one and four from generation 73 to generation 100.



# Chapter 7

## Conclusions and Future Work

This thesis has investigated the effect that cultural and biological evolution has on language evolution, as well as how the relation between language evolution and social networks function. The first research question was concerned with the design and strengths and weaknesses of the state-of-the-art computational models in the field of language evolution. This was answered thoroughly through the discussion of four such models and comparing them to each other: Lipowska (2011); Lekvam (2014); Gong (2011); and Munroe and Cangelosi (2002).

Research question two was answered by presenting a computational model that was an extension of the work done by Lekvam. The naming game, social network, and evolutionary algorithm that Lekvam used, were also used by this computational model. However, the fitness function, the method for calculating the weight of a word, and the genome were changed in an attempt to make the model more realistic and to study the impact cultural and biological evolution has on language evolution.

Research question three was concerned with how social networks have an impact on language evolution. This thesis' contribution to the field of simulating language evolution was the new fitness function which was designed to allow both extrovert and introvert strategies to be viable. This did provide different results than Lekvam had, but all experiments performed in this study evolved similar social networks. In order to conclude on how social networks and language impact each other's evolution, more specific research on this topic needs to be performed.

The fitness function implemented seemed to fulfil its purpose, both extrovert and introvert

strategies had advantages in this fitness function. The simulations preferred an extrovert strategy during the initial generations, but in the end, an introvert strategy was favoured. In real life, the vast majority of conversations a person conducts are with other persons that the person already knows, meaning that an introvert strategy is favoured.

Parent-child dialogues were the most important factor in order for the simulations to reach a consensus. This indicates that the initial dialogues a newborn conducts are of crucial importance for its language, and it was virtually impossible to learn a new language properly after having lived for two generations in the simulations. Two generations in the simulation is not equal to two generations in real life, it merely represents that it is almost impossible to become fluent at a new language after having lived for a while. This result indicates that being taught the same languages during the early years is an important factor on the *ontogenetic* time scale. Both in this model and in reality, language is easier to learn while one is young and the older one becomes, the harder it is to learn new languages.

The more the selection pressure was reduced, the more the simulations struggled to reach consensus. When the selection pressure was reduced, the effects of the *phylogenetic* (biological evolution) and *glossogenetic* (cultural evolution) time scales were reduced. This indicates that these two evolutionary forces have an impact on language evolution. It also indicates that the Baldwin effect (Baldwin, 1896) affects the cultural and biological forces of language evolution. Other simulations have also concluded that the Baldwin effect has an impact on language evolution (Lipowska, 2011; Zollman and Smead, 2010; Chater and Christiansen, 2010; Lekvam et al., 2014).

The results from this model indicate that it partly simulates language evolution successfully. The model indicates that the first years are of great importance for an individual's language, that mostly acting introvertly is beneficial, and that the Baldwin effect has an impact on language evolution. However, some parts of the model were made too simplistic to simulate language evolution in a sufficiently realistic manner. These weaknesses will have to be improved upon in order to learn more about how human language has evolved into what it is today.

This computational model also attempted to investigate the relation between social networks and language evolution. The social network evolved into a network where all agents were transitively connected in all simulations. The simulations represented a small population, theoretically living with no distance between each other. A small town would

likely have a network where almost everyone was transitively connected to each other, but small subgroups might exist within the population. Extending the computational model with geographical distances within which groups would form and simulating movement of the groups would be interesting. Implementing geographical distance could provide a better understanding of the relation between social networks and language evolution.

The genome evolved towards one local maximum, and all agents practically ended with the same personality. That personality favoured an introvert strategy, speaking to its newborns, and a relatively high learning ability. The personality based genome worked well in terms of depicting the important values to have in this model, but it did not depict the large variance between personalities that exists in real life. In reality, a population has some traits that always are favourable, and some traits that have positive and negative aspects to them. Considering that evolutionary algorithms are intended to search for optimal solutions in a complex search space, it is natural that the variance in personality was not represented properly in this model. The model would have to be altered in a manner which allows more variance and multiple personality strategies to be viable in order to incorporate personalities into the model.

The naming game used in this model was a substantial simplification of human language. There are several possible ways to improve the language used in this computational model. The simplest method would be to have more objects to name. In reality, new objects appear, forcing the population to agree upon words for the new objects. To simulate this, the naming game could be extended by adding objects to name at certain points during a simulation. Replacing the naming game entirely could also be considered. Two other methods for simulating language was presented in the literature study: the word order regularity model made by Gong et al. (2004) and the model using artificial neural networks as the agents' language interpreters made by Munroe and Cangelosi (2002).

The most important findings of this thesis indicate that mostly acting introvertly is beneficial compared to an extrovert strategy and that language evolution to a certain extent is influenced by the Baldwin effect. Further investigations on how social networks affect language evolution could be the key to understanding how the human language evolved to what it is today.




# Appendix A

## Results Not Included in the Main Report


This appendix displays all results that were not relevant enough to be a part of the results chapter. This mostly consists of the social networks from the later experiments. The results not included from experiments 6b and 7b will not be shown because they were conducted with one specific purpose, so the other graphs are not relevant.

## A.1 Experiment 4


(a)




(b)



(c)




(d)




**Figure A.1:** Snapshots of the social network of experiment 4 at the following generations: (a) 5, (b) 20, (c) 40, and (d) 100.

## A.2 Experiment 5


(a)




(b)



(c)




(d)



**Figure A.2:** Snapshots of the social network of experiment 5 at the following generations: (a) 5, (b) 20, (c) 100, and (d) 150.


### A.3 Experiment 6




**Figure A.3:** Snapshots of the social network of experiment 6 at the following generations: (a) 5, (b) 20, and (c) 100. (d) Three graphs that show the evolution of the traits in the genome. The green function represents the average probability of conducting parent-child dialogues, the blue function represents the average learning rate, and the red function represents the average probability of acting extrovertly.

## A.4 Experiment 7


(a)



(b)



(c)



**Figure A.4:** Snapshots of the social network of experiment 7 at the following generations: (a) 5, (b) 20, and (c) 100.



# References

- Baldwin, J. M. (1896), ‘A new factor in evolution’, *The American Naturalist* **30**(354), 441–451.
- Bickerton, D. (2007), ‘Language evolution: A brief guide for linguists’, *Lingua* **117**(3), 510–526.
- Chater, N. and Christiansen, M. H. (2010), ‘Language acquisition meets language evolution’, *Cognitive Science* **34**(7), 1131–1157.
- Chomsky, N. (1986), *Knowledge of language: Its nature, origin, and use*, Greenwood Publishing Group.
- Chomsky, N. and Halle, M. (1965), ‘Some controversial questions in phonological theory’, *Journal of Linguistics* **1**(02), 97–138.
- Christiansen, M. H. and Kirby, S. (2003), ‘Language evolution: Consensus and controversies’, *Trends in cognitive sciences* **7**(7), 300–307.
- Gong, T. (2011), ‘Simulating the coevolution of compositionality and word order regularity’, *Interaction Studies* **12**(1), 63–106.
- Gong, T., Ke, J., Minett, J. W. and Wang, W. S. (2004), A computational framework to simulate the coevolution of language and social structure, *in* ‘Artificial Life IX: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems’, Tremont Hotel, Boston, MA, USA, pp. 158–64.
- Hauser, M. D., Chomsky, N. and Fitch, W. T. (2002), ‘The faculty of language: what is it, who has it, and how did it evolve?’, *Science* **298**(5598), 1569–1579.
- Jain, A. K., Mao, J. and Mohiuddin, K. M. (1996), ‘Artificial neural networks: A tutorial’,

- Computer* **29**(3), 31–44.
- Lekvam, T. (2014), Co-evolving language and social structure using a genetic algorithm, Master’s thesis, Norwegian University of Science and Technology.
- Lekvam, T., Gambäck, B. and Bungum, L. (2014), Agent-based modeling of language evolution, *in* ‘Proc. of 5th Workshop on Cognitive Aspects of Computational Language Learning (CogACLL)@ EACL’, pp. 49–54.
- Lipowska, D. (2011), ‘Naming game and computational modelling of language evolution’, *Computational Methods in Science and Technology* **17**(1-2), 41–51.
- Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (2013), *Machine learning: An artificial intelligence approach*, Springer Science & Business Media, Berlin, Germany.
- Müller, F. M. (1861), Lectures on the science of language, *in* ‘The theoretical stage, and the origin of language’, Longman, Green, Longman, and Roberts, London, UK, pp. 7–8.
- Munroe, S. and Cangelosi, A. (2002), ‘Learning and the evolution of language: The role of cultural variation and learning costs in the Baldwin effect’, *Artificial Life* **8**(4), 311–339.
- Pinker, S. and Bloom, P. (1990), ‘Natural language and natural selection’, *Behavioral and brain sciences* **13**(04), 707–727.
- Pinker, S. and Prince, A. (1994), ‘Regular and irregular morphology and the psychological status of rules of grammar’, *The reality of linguistic rules* **321**, 51.
- Quillinan, J. (2006), Social networks and cultural transmission, Master’s thesis, University of Edinburgh.
- Schwartz, S. H. (1992), ‘Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries’, *Advances in Experimental Social Psychology* **25**, 1–69.
- Schwartz, S. H. (2012), ‘An overview of the schwartz theory of basic values’, *Online readings in Psychology and Culture* **2**(1).
- Steels, L. (1997), ‘The synthetic modeling of language origins’, *Evolution of communication* **1**(1), 1–34.
- Steels, L. (2011), ‘Modeling the cultural evolution of language’, *Physics of Life Reviews*

8(4), 339–356.

Steels, L. (2012), *Experiments in cultural language evolution*, John Benjamins Publishing, Institute for Evolutionary Biology, Barcelona, Spain.

Suzuki, R. and Arita, T. (2008), How learning can guide evolution of communication, *in* ‘ALIFE’, Nagoya University, Nagoya, Japan, pp. 608–615.

Tomasello, M. and Rakoczy, H. (2003), ‘What makes human cognition unique? From individual to shared to collective intentionality’, *Mind & Language* **18**(2), 121–147.

Zollman, K. J. and Smead, R. (2010), ‘Plasticity and language: an example of the Baldwin effect?’, *Philosophical studies* **147**(1), 7–21.