NTNU – Trondheim
Norwegian University of
Science and Technology

Co-evolving Language and Social Structure Using a Genetic Algorithm

Torvald Lekvam

Master of Science in Informatics
Submission date: June 2014
Supervisor: Bjørn Gamback, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Co-evolving Language and Social Structure Using a Genetic Algorithm

Torvald Lekvam

June 2014

MASTER THESIS

Department of Computer and Information Science
Norwegian University of Science and Technology

Supervisor 1: Professor Björn Gambäck
Supervisor 2: PhD Candidate Lars Bungum

Preface

A Master's thesis in Computer and Information Science at NTNU as part of the study program Artificial Intelligence. The thesis has been carried out in the time frame 2013-08-19 and 2014-06-01.

Trondheim, 2014-06-01

Torvald Lekvam

Problem Description

The original problem as given by given by Lars Bungum and Björn Gambäck:

> Natural language processing grapples with an ever-changing and moving target. The focus of study, natural language, is natural because it changes, interacts and evolves in various directions. The bio-inspired computational methods described as evolutionary computation and/or genetic algorithms create computational models that evolve a population of individuals to find a solution to a given problem. This project will investigate how evolutionary computation can been employed in some natural language processing task, ranging from efforts to induce grammars to models of language development through parameter optimization and search.

Acknowledgment

I would like to express my gratitude to Professor Björn Gambäck for the opportunity to work on such an interesting topic. Moreover, I am grateful for his support and patience with me throughout this work.

Also, a special thanks to PhD Candidate Lars Bungum, for valuable suggestions and encouragement.

An addition thanks to both for helping out getting parts of this work published.

T.L.

Summary

It is interesting how we can take a train of thought and transfer this into an other person's mind by pushing the air around us. Human language, this complex medium that distinctly separates humans from animals, has baffled scientists for centuries. But as it lacks of historical data, researchers wish to benefit from computer science and the field of artificial life to understand the origin of language. This thesis illuminates the potential for using agent-based models to investigate the relationship between biology, culture and behavior on an individual level.

This is done in two parts. First, different theories and computational models experimenting with language evolution are presented. This includes a thorough implementation of and elaborations on one recent paper, where language acquisition is illustrated favorable over multiple evolutionary time scales in an agent-based model. In the second part, a more bio-inspired methodology is proposed to make the former model more robust and better suited for extensions. This is demonstrated by letting the agents evolving some social biases, while they are conducting a naming game in a social structure. A naming game is an abstraction, often used in the research field, to model the spreading and diversity of language. Through pair-wise dialogs, the goal of the game is to reach self-organized agreement on naming an arbitrary object in their environment. Given the assumption that communication is beneficial for social structure and that social structure is beneficial for reproduction, the experimental work demonstrates that agents are able to build social structures that resembles real life social topologies, although the naming game might happen too rapid in respect to the evolving social structure.

Hopefully, with support from other disciplines, the presented model is suited for further investigation of social, or other functional, traits that can influence language evolution.

Sammendrag

Det er interessant hvordan vi kan ta et tanketog og overføre dette til en annen persons sinn ved å skyve luften rundt oss. Menneskelig språk, dette komplekse mediumet som tydelig skiller mennesker fra dyr, har forbløffet forskere i århundrer. Og siden språk har lagt fra seg særlig lite historiske data, ønsker forskere å dra nytte av inforamtikk og metoder i kunstig liv å forstå opprinnelsen til språk. Denne avhandlingen belyser potensialet for bruk av agentbaserte modeller for å undersøke forholdet mellom biologi, kultur og opptreden på individnivå.

Dette gjøres i to deler. Først er ulike teorier og beregningsmodeller som eksperimentere med språkevolusjon presentert. Dette omfatter også en grundig implementering og elaborering av en nyere artikkel, der språketanskaffelse er illustrert gunstige over flere evolusjonære tidsskalaer i en agentbasert modell. I den andre delen, er en mer bioinspirert metodikk foreslått å gjøre den tidligere modellen mer robust og bedre egnet for utvidelser. Dette er demonstrert ved å la agentene få muligheten til å utvikle noen sosiale bias, mens de prøver å gjennomføre et navnespill i en sosial struktur. Et navnespill er en abstraksjon, ofte brukt i forskningsfeltet, for å modellere spredning og mangfold av språk. Gjennom parvis dialoger, er målet med spillet å nå selvorganisert enighet om å navngi en vilkårlig objekt i deres miljø. Gitt antagelsen om at kommunikasjon er gunstig for bygge sosial struktur og at sosial struktur er gunstig for reproduksjon, så viser de eksperimentelle resultatene her at agenter er i stand til å bygge sosiale strukturer som ligner sosiale topologier i virkelig liv, selv om navnespillet kanskje går for hurtig i forhold til utviklingen i den sosial strukturen.

Forhåpentligvis, med støtte fra andre disipliner, kan denne presenterte modellen være egnet for videre undersøkelser av sosiale, eller andre funksjonelle, egenskaper som kan innflytelse språkets evolusjon.

Contents

Chapter 1

Introduction

It is baffling how children can effortlessly acquire the structure and vocabulary of a whole language despite receiving partial and noisy input. Researchers have for several decades tried to understand the complex sociocognitive mechanisms underlying language and its evolution — but there is still no comprehensive solution on the horizon. However, in the last two decades, computational techniques have given the field of language evolution a lot of support, mainly by offering models to validate, explore, and explain different theories.

In this thesis, different theories on language evolution are presented and different computational models are discussed. In more detail, one particular model by Lipowska (2011) is examined. A different methodology is proposed for the same model, more suitable for extensions. The proposal is demonstrated by co-evolving a communication system and a social structure, inspired from recent literature.

This work started out with the topic "genetic algorithms for language processing", two terms equally fascinating. A *genetic algorithm* is in a bio-inspired heuristic search algorithm based on the Darwinian idea of natural selection. It is used in a variety of scientific fields where the solution to a problem lay in a complex and rough search space (Ross & Corne, 1994).

Here, *language processing* is the nail where a genetic algorithm would be the hammer. How language is processed, both by man and machine, is an enormous research field — mainly because it intersects with so many scientific disciplines (Bickerton, 2007; Gong *et al.*, 2014). The research field — including linguistics, biology, neural science, anthropology, archeology, computer science, philosophy, and others — researchers are mostly motivated by two facts. Firstly, understanding and applying language have obviously has great commercial potential, both as a helping aid for those in need, but also as a tool in, for instance, statistical analyzes and information extraction. And secondly, we do not fully comprehend the dynamics involved in language and therefore, scientists wish to identify and understand the origin and constraints of this complex medium to truly come with an explanatory linguistic theory (Kirby, 2002b).

In this intersection, between the evolutionary computation paradigm (which also includes *evolutionary programming* and *evolution strategies*) and language processing, lies a handful of interesting topics. Typical applications in language processing are spelling and grammars checking, text summarization, machine translation systems and induction — all which have been tackled by evolutionary algorithms with various results (Bungum & Gambäck, 2010). These endeavors can all be characterized as typical top-down approaches, whereas modern language is split up into units, groups and relations, and then categorized and structured into statistical and rule-based models. This surely gives great insights to how language behaves and can be used on a daily basis, but it does not say much on its genesis. To study this, a bottom-up approach could be used (Bedau, 2003). A *language game*, which this thesis is focusing on, is such an approach. A language game is originally a philosophical concept, first discussed by Wittgenstein (1953), where a functional language is woven and established trough interactions between speakers and hearers, and particular actions in an environment. From a computer science perspective, this concept was picked up in the 1990s (Steels, 1995). Models on artificial life were adapted to study this and other linguistic phenomena. This later picked up, and have been a corner stone in many bottom-up approaches to language evolution. Computational simulations, such as these, allow researchers to test theories in complex landscapes where intuition often is limited (Christiansen & Kirby, 2003a).

Throughout this thesis the term **language** will be used extensively. When not specified, the term is used in a communication system sense, meaning its not necessarily directly associated with human language unless context implies otherwise.

1.1 Objectives

The main objectives of this Master's project are to address these following points.

1. Investigate the main dynamics in language evolution, and the computational models trying to capture them. Advantages and limitations of these computational models will be addressed.

2. One model from the recent literature will in examined in detail.

3. As language evolution is such a multidisciplinary research fields, a proposals for modeling the different evolutionary forces simultaneously will be offered, by explanation and demonstration. As for the demonstration, this thesis will pick up where others

left off, by simulating relations between the emergence of language evolution and social networks.

Problem Formulation

To understand how language works, how we apply it, its origin and fundamentals; we commonly tend to look at human language, probably our best source for analysis. It is a matter of definition, but there are around 6000-8000 different human languages worldwide today, all holding an extensive diversity of syntax and semantics, phonetics and morphology (Evans & Levinson, 2009). That being said, these non-random languages represent perhaps only 2% of all those thought to have ever existed (Pagel, 2000). And as not enough entropy is lost, 96% of the world's population uses only 4% of the different languages that are alive today (Crystal, 2002). From only studying these languages, can we draw hard constraints for what human language really is, what biological/cultural assumptions it operates on, and how and why it holds so much variation?

The questions raised in the last paragraph cannot simply be answered by analyzing today's situation, we have to look back at our ancestors. But there is a problem in linguistic history — our reconstruction techniques can only take us back six or perhaps seven thousand years. Beyond this point, researchers can only speculate on when and how human language evolved: either as a slowly proceeding process starting millions of years (mega-annum, Ma) ago, e.g., 7 Ma ago with the first appearance of cognitive capacity or 2.5 Ma ago with the first manufacture of stone implements; or through some radical change taking place about 100 ka ago with the appearance of the modern humans or 50–60 ka ago when they started leaving (Tattersall, 2010; Klein, 2008). Before language, an early hominid wanting to trade a banana for some nuts would not have the opportunity to do so, in the act he could just as easily get mistaken for stealing or donating items. But somewhere around here, language emerged and we were able to deal with such ideas. This trait, the ability to learn expressibility, would soon be the foundation for so-called cumulative cultural evolution; how we transmit ideas and wisdom (and language itself) from generation to generation through language.

Researchers want to shed light on the evolution of language to get better perspective and deeper understanding of the dynamics of language, as well as on the brain and the human linguistic phenomena we all are a result of. Moreover, they want to understand how it affects our social affairs, such as how we connect networks (Lupyan & Dale, 2010).

Language evolution has said to be "the hardest problem in science" (Christiansen & Kirby, 2003b), mainly because investigating language change, language acquisition, and language origin is such a highly interdisciplinary research field. Hence, this thesis focuses on how and why computer simulation can help answering parts of these questions.

1.2 Approach

This thesis is mainly investigating agent-based models of language evolution. These models are often characterized by following decentralized and self-organization principles, which are not only fascinating, but also very life like. From local behavior, constrained by certain rules, an overall complex population (or system) can arise with properties greater then the sum of it parts. A complex system is characterized as interacting entities and non-linear interactions among them (Helbing & Balietti, 2011).

Human language can be seen as a complex (adaptive) system, which makes it theoretically suitable for a bottom-up investigation (Loreto & Steels, 2007). Bottom-up models are not necessarily motivated by predicting when human language first emerged — which is often quite hard with stochastic phenomena in nature — but as explanatory toolkits for examining how the dynamics of diversity can unfold (Epstein, 1999).

The approach in this thesis is two-folded. First, there would be some background theory on different computational models experimenting with language evolution — which includes a thorough implementation of and elaborations on one recent paper (Lipowska, 2011), where language acquisition is illustrated favorable over multiple evolutionary time scales. Secondly, another methodology is proposed, and extended with other innate traits, presumably more suited for further research.

1.3 Limitations

Computer simulations of full-blown social phenomena are very rare — simply because they are truly difficult to get right (Gilbert & Terna, 2000). Bottom-up or complex systems are common in physics, with, for instance, observable thermodynamics and statistical models underneath – but in social science, there are no implied laws for how the statistical models should look like. Thus, researchers often practice with mean field (game) theory, lettering a small number of individuals be the object of study when the numbers in reality are extensive. As for language evolution in particular, as it exists in (historical) traceless utterances and neurons firing in our yet unexplored brain — many simplifications must be done when modeling such an adaptive system. Artificial neural networks are extremely simple, if sentences are used in an agent-based model they tend to only be a couple of words long. For instance, languages are often so abstract that simple holistic words are the only linguistic dynamics involved. This intersects with specificity, as single topics can be investigated and others are ignored — for example, no semantics or pragmatics are involved when studying language change, or biology tends to get ignored when simulating cultural aspects (Gong & Shuai, 2013).

1.4 Contribution

This introduction raises some far-reaching questions about language evolution. Admittedly, there will not be drawn any over-ambitious conclusions from social science or linguistic points of view. Instead, this thesis will look at different computations models used to investigate parts of these questions, discuss different *naming game models* within these frameworks, and demonstrate how this can be done slightly differently by combining a few recent models from the literature using a *genetic algorithm*.

Language evolution is a highly interdisciplinary topic, hence it must be tackled from more than one angle simultaneously. As this thesis has a solely computer science background, this work can be taken as insight for other fields — such as mathematical, empirical or other experimental work on language evolution — when they are ready for further examination.

Parts of the two first objectives in this theses have already been published (Lekvam *et al.*, 2014), mainly the reimplementation of a model in the literature. This will be elaborated much more thoroughly through the thesis, before presenting a model and demonstrating this model as stated in the last objective. Hopefully, this can give inspiration and ideas for other researchers in the field of language evolution.

1.5 Structure of the Report

The rest of the report is structured as follows. Chapter 2 will give an overview of related work and introduce some terminology, In Chapter 3, a model from the literature will be presented, reimplemented, and discussed. Chapter 4 will suggest improvements and extensions. The results will be viewed in Chapter 5 and a final discussion with conclusions and directions for further work is given in Chapter 6.

Chapter 2

Background

This chapter will draw a line from different theories and aspects on language evolution and through computational models. Lastly, some words on network theory and genetic algorithms are presented.

2.1 Language Evolution

There are two main ideas in biological evolution as to why humans developed communication through speech. According to the first, language, or more precisely the ability to bear the full structure of language, came as an epiphenomenon, a byproduct of an unrelated mutation This theory, outlined mainly by Chomsky (1986), states that a mental language faculty could not by itself evolve by natural selection; there would simply be too many costly adaptations for it to be possible. He argues that there exists a figuratively *language acquisition device* in the human brain that can understand a *universal grammar*. This universal grammar then holds a finite number of rules that can carry any language. This view is often in the beholders of *nativists*. According to the second, and more mainstream idea, that language emerged in a strictly Darwinian process as Pinker & Bloom (1990) proposed, all explained by natural selection as with other complex traits like echolocation in bats or stereopsis in monkeys. Their much cited paper caused a new spark in the area of language evolution, and most work in recent years is based Darwinian ideas because of this. A very recent debate is how much individual and social learning is influencing this Darwinian approach. Christiansen & Chater (2008), for example, argue that language is entirely build by cultural conventions and thus shaped to fit the existing cognitive abilities in the brain and not the opposite.

As mentioned, the evolution of language is a widespread research field — and to tackle this, one has to approach the domain from several disciplines simultaneously. Foremost, the conventional theories on language evolution must be understood from three different perspectives (Steels, 2012); biological evolution, cultural transmission and learning on the individual and social levels.

This section will elaborate on these three different adaptive systems and look at theories and computational models of language evolution with this in mind.

2.1.1 The Evolutionary Forces

- Biological evolution is the slowest force of the three, and is often said to be working on a *phylogenetic* time scale. Language initially requires some cognitive capacities before emerging within a species. These capacities face a pressure as both increased ecological and linguistic complexity co-emerges (Steels, 2012).

- Culture evolution works on a *glossogenetic* time scale, a much faster time scale. Here, change can be viewed in terms of whole languages. Within a social group or community there are shared linguistic paradigms, rules, norms and phonology that exclusively get transfered culturally, from individual to individual, from generation to generation. In this way, language can co-evolve with biology, as it changes and grows cumulative on top of biological cognitive abilities (Kirby *et al.*, 2008). While cultural change could be influenced by selective pressure in biological evolution, cultural changes can also guide biological evolution to some extent (see Section 2.1.3).

- Individual and social learning must be looked at within an individual's lifetime, on a *ontogentic* time scale. This is clearly an important aspect of how language is built, as newborn humans go from expressing no language to several possible languages in a fairly short time. Through each individual's lifetime, small variations in speech patterns, morphology, and so on works as fuel of the cultural evolutionary engine.

2.1.2 Origins

Before taking a computer science perspective — let us look at the actual language. Prior to the emergence of human language as we know it today, there had to be some intermediate, a simpler language. It is a debated topic, but there is to some point agreement upon the use of a protolanguage with holistic symbols, a simple language toolbox where discrete sounds or gestures would map directly to meanings or concepts (Christiansen & Kirby, 2003a). This naturally assumes that there was something beneficial to communicate about, as for instance the need of food or to warn of danger. Thus, this argues that there should be a connection between the emergence of language and the emergence of cognitive abilities and social intelligence — even though few can agree on what the connections might be (Bickerton, 2007).

Since no other species in nature is known to have the same extensive language capabilities as humans, comparative methods are difficult. Primates can show several of the same communicative signs as humans, for example, giving warning sounds in the case of danger, but as discussed by, for example, Cheney & Seyfarth (2005) they seem not to be able to do this with communicative intent: listeners acquire information from signalers who do not really intend to provide it, at least not in a human sense. Thus, songbirds have shown similarities in their abilities to learn sophisticated auditory perception and vocal production — however, they seem to lack the abilities to map "words" to meanings other essential cognitive abilities (Bolhuis *et al.*, 2010). This is often used as an argument for speech and other cognitive abilities being two different evolved traits (Berwick *et al.*, 2013).

Having syntax, the ability to express a vast of meanings efficiently with the help of structure, is perhaps what distinguishes humans from other animals the most. The question arises, if this syntactic ability evolved before a protolanguage (as a similar neural mechanism for something else) or afterwards — but there does not seem to be any consensus (Bickerton, 2007).

Although, genetic researchers did in 1995 find a language related gene in the human DNA that was functionally different from chimpanzees. They called it Forkhead-box P2 (FOXP2) and mutation of this gene was observed to give both heavy speech and language disorder. This was quickly interpreted as the 'grammar gene', but researchers soon falsified this to be the only genetic factor making humans capable of language. While studying the FOXP2 gene further it was discovered that the gene also is associated with gut, lung, and heart function (Chater *et al.*, 2009) and that we share this gene with a number of animals lacking communication abilities (Fitch, 2005).

2.1.3 The Baldwin effect

Even more interesting — there is the Darwinian idea that cultural learning can guide biological evolution, a process known as *the Baldwin effect* (Baldwin, 1896; Simpson, 1953). The theory argues that cultural learned traits (as for instance an universal understanding of grammar, or a defense mechanism against a predator) could assimilate into a specie's genetic makeup. Teaching each member in a population the same thing over and over again comes with great cost (time, faulty learning, genetic complexity), and the overall population saves a lot of energy if this learned trait would become innate. On the other side, there is a cost connected to genetic assimilation as it can prohibit plasticity in future generations and make individuals less adaptive to unstable environments. So in other words, the Baldwin effect is a two step process. First, a population of individuals can increase their reproduction abilities if they are able learn to deal with a curtain threat or benefit, through their lifetime, in their environment. In the next step, if the environment is sufficiently stable, there will be a selective pressure towards individuals who get these traits innate and thus do not need more costly learning. There have been much debate in recent time whether language is a result of the Baldwin effect or not (Zollman & Smead, 2010; Evans & Levinson, 2009; Chater *et al.*, 2009; Lipowska, 2011; Kuechle & Rios, 2012; Suzuki & Arita, 2013), but questions, hypotheses, and computer simulations flies in multiple directions. For instance, Chater *et al.* (2009), based on their simulation argue that only the stable parts of language may be able to assimilate into the genetic makeup, as the variation (in word-order, morphology, cases, and so on) within the linguistic environment is to unstable to be a target of natural selection. Watanabe *et al.* (2008), on the other hand, use a similar model and argue that genetic assimilation not necessarily needs an unstable linguistic environment to unfold (see Section 2.6 for an elaboration on the methods in these models). The Baldwin effect should however not be confused with the Lamarckian mechanism, a very similar effect. The theory of Lamarckian inheritance states that acquired characteristics (through one's lifetime) can be genetically inherited directly (Sasaki & Tokoro, 1999), in contrast to the Darwinian approach where only the learning apparatus can be genetically inherited.

2.1.4 Computational Models

We can divide computational models on language evolution into three groups; evaluation, exploration and exemplification (Christiansen & Kirby, 2003a).

Models that evaluate mathematical methods can help researchers determine and understand, or simply underline, the soundness of different theories. Computational models by exploration have the advantages, unlike models in for instance psychology, to audit every agent's mental state at any time, as well as add and remove cognitive or external influences and study the effect. Although it must be done in a careful manner, this could give rise to new theories. The most common might be the models of exemplification where scientists from different disciplines wish to illustrate a point by constructing artificial life in a suitable environment. However, simulations can be combinations of two, or all three types,

The field is highly influenced by the prominent work of Steels (see his review (Steels, 2011)). A recent review also worth mentioning is Gong & Shuai (2013).

Simulations of this kind are of course not full-blown simulations of how the human language works, the idea of such as simulation is too far fetched (Helbing & Balietti, 2011; Gong & Shuai, 2013). Computational research in this field are limited to modeling very simplified features of human language in isolation, such as strategies for naming colors (Bleys & Steels, 2011; Puglisi *et al.*, 2008), different aspects of morphology (Dale & Lupyan, 2012), grammatical agreement (Beuls & Steels, 2013), simulations of dialect diversity (Livingstone, 2002), word-order (Gong, 2011), how our phonetics might have evolved (de Boer, 1997), emergences of different language families (De Oliveira *et al.*, 2008), and so on. There have also been a lot of studies trying to establish the relationships between linguistic features and social structures, both on syntactical and phonological structure of language (Lupyan & Dale, 2010). This simplicity is important to keep in mind, since it is conceivable that certain features of language can be highly influenced by other features in real life.

There have also been studies trying to cover the emergence of compositional structures, although under the assumption that there is a specialized acquisition mechanism prior to the emergence of syntax, which is not necessarily true (Vogt, 2005; Kirby, 2002a). However, there have been studies trying to capture the role of such a learning apparatus, mainly through studying the Baldwin effect (Lipowska, 2011; Munroe & Cangelosi, 2002). Hinton & Nowlan (1987) were the first to computational demonstrate this in their much cited paper, with the use of a simple neural network and a genetic algorithm. They argued that learning can facilitate evolution of in finding "a needle in the haystack" (see Section 2.6).

2.2 Language Games

A computational language game (hereafter called a language game), first introduced by Steels (1995), is a model where artificial agents in turn interact with each other in order to reach a cooperative goal; to make up a shared language of some sort, all while minimizing their cognitive effort. All agents are to some degree given the cognitive ability to bear language, but they are not given any prior knowledge of how language should look like or how the consensus should unfold. There are no centralized anchors involved, a simulation is all self-organized. In that paper, Steels argued that this kind of mechanism, aligning communication systems among artificial agents through conversations, would open up for new exciting lines of linguistic research — and he was right.

In a basic language game, two agents are chosen (mostly at random) as hearer and speaker, and they are made to exchange an utterance about a certain arbitrary concept or meaning in their environment. If the agents use the same language — meaning the utterance is understood by both parties — the conversation is a success. If the speaker utters something unfamiliar to the hearer, the conversation is termed as a failure. As mentioned in the introduction, language does not need to be verbal, language games consider both verbal and gesture interaction, together or apart.

If an agent wants to express a certain meaning or concept without having any term for it, the agent is assumed to have the ability to make one up and add this to its memory. While interpretation in real life is a complex affair, it is mostly assumed that there is a more direct connection between utterance and actual meaning in language game models (emotions and social situations do not bias how language is interpreted). An utterance can be represented as a simple binary signal (Cangelosi & Parisi, 2002), holistic words (Lipowska, 2011), or structured grammar with multiple words (Gong *et al.*, 2004). How these are implemented in the different models varies from neural networks (Cangelosi & Parisi, 2002; Batali, 1998), Fluid Construction Grammar (Steels, 2011), rule-based systems (Kirby, 2002a), and others.

Further, most language games deal with *alignment strategies* (Lei *et al.*, 2010; van Trijp & Steels, 2012), a term used for scoring or weighting each agent's utterance-concept mapping. Alignment strategies help the agents reach agreement faster and more efficiently by giving feedback and adjust weights after each dialog. Strategies tend to vary, but lateral inhibition might be the most common one, where successful usage of a particular mapping also dampers other mappings with the same utterance. The most extreme use of lateral inhibition is when only the transmitted utterance is kept after a successful dialog (see Section 2.2.7).

A language game is typically split up into four categories (Wagner *et al.*, 2003), as in a two by two matrix, determined by the two factors *structured* and *situated*. A language game is termed as structured if it aims to cover language with grammars and unstructured if the language is holistic. Further, a language game is said to be situated if the simulation takes place in a virtual world where agents perceive their environment and act upon it in addition to verbal interactions. In contrast, a non-situated language game's only focus is the language interactions, therefore agents can be abstract models without any meaningful environment.

Almost all language games share a common goal of reaching a single coherent and consistent language, but there are also studies where the aim is to simulate the emergence of bilingualism or multilingualism with different models (Lipowska, 2011; Roberts, 2013; Castelló *et al.*, 2008).

As there have been a lot of computational work on this, language games have also been studied on humans, often based on earlier computational models (Scott-Phillips & Kirby, 2010).

If used, cultural evolution is captured by

1. horizontal communication between adults or individuals in the same generation, often with bi-directional dialogs, or

2. vertical communication from adults to children, where dialogs most often are uni-

directional. The latter typically lets the agents breed, age and die (this endeavor is more recently known as the *iterated learning model*, see Section 2.3).

In the next subsections a presentation of different classical language games will be given, and some examples will be brought up.

2.2.1 Guessing game

A guessing game is a setup — most often used with embodied robots — where there are only two agents involved, one acts as a speaker and one as a hearer. The speaker's role is to utter a term about an object or concept in an environment and the hearer tries to point out what object or concept the speaker have in mind. The best, and maybe the first, example of this simple language game is the Talking Head Experiment by Steel's team (Steels, 1999). Here, the environment is perceived and then modeled with discrimination trees or similar. How the world gets modeled is the most important aspect of a guessing game as it forms the conceptualizations and thereafter the language itself.

2.2.2 Grounding game

The grounding game is a lot like a guessing game, but there are often more agents, and they are endorsed with improved cognitive functions for conceptualizations and semantics. In every simulation, agents are — embodied or not (much of L. Steels' leading work is about grounded embodied agents) — situated in an environment with objects and with perceptual input from a camera or similar. Every agent uses a *semiotic network* (Steels & Loetzsch, 2012) to conceptualize the arbitrary objects (often called prototypes) and their features. Semiotics is the study of meaning-making, and how signs and symbols can refer to different concepts. Agents try to agree upon a shared vocabulary with ordinary language game conventions by relating utterances between the closest prototypes (by some measure in a suitable search space). By doing this relation, the agents are said to meet on "common ground", hence the name.

2.2.3 Classification Game

A classification game shares similarities with grounding games, but focuses more on a particular problem, which autonomous agents are set out to solve or classify, all through evolving a shared language (see Swarup & Gasser (2009, 2010)).

2.2.4 Signaling Games

A signaling game puts the natural language aside and focuses on more direct (and often innate (Noble, 1999)) communication systems. While natural language (words, letters, and sentences) can be difficult to handle in a neural network, a signaling game can facilitate this by using only signals (as discrete or continuously numbers) as word representation. This makes it easier to model the emergence of compositionality and recursive structures, studied in the work by Suzuki & Arita (2008) and Zollman & Smead (2010).

2.2.5 Spatial Naming Games

There have also been studies where agents in an environment have to communicate and differentiate between the location of objects in the environment. This is a language game where choice between different strategies is essential for reaching consensus. For instance, which spatial references you are using when naming different objects in a room, does your utterance rely on your view or your hearer's view, or does one simply use a global reference system such as a cardinal system (Spranger, 2013).

2.2.6 Action Game

Another language game is one that tries to model the cognitive dynamics of naming different body actions. The research done on action games is mainly by Steels *et al.* (2012), naturally with the use of embodied agents. The agents build up semiotic networks to bind relations between invented words, visual prototypes and motoric behaviors.

2.2.7 Naming Game

A naming game (Steels, 1995) takes the semiotic dynamics of a language game for granted and looks more closely at the spreading of simple words in a population. Such a game can be seen as a grounding game where there is only one object in the environment to perceive. This is an often used abstraction for investigating cultural transmission horizontally, between agents of the same generation. It also works over generations (Lipowska, 2011), for capturing biological aspects such as the Baldwin effect (as discussed in Section 2.1.3).

A simple naming game is therefore typically non-structured (no evolving grammars) and non-situated (no concrete environment). Agents, all equipped with internal inventories containing known utterances, attempt to agree on one particular utterance through pairwise dialogs. The inventories before and after a failed and successful dialog, with a simple alignment strategy without weights, can be viewed in Tables 2.1 and 2.2.

In an ordinary, non-structured naming game, a simulation is characterized by two phases (Baronchelli *et al.*, 2006a). In the first phase, words spread around in the population, and every agent's lexicon grows even more rapidly (often up to $populationSize/2$, since inventories are empty at the first time step and the first dialogs share one word with two agents). As agents occasionally get successful dialogs, words start to disappear. This naturally increases the possibilities of new successful dialogs, and the simulation enters its second phase; a steady progress towards consensus, the state when every agent share the same word and can understand each other. The simulation slows down at the end of the last phase, as getting rid of the last words requires some chance. Thus, the function of successful dialogs through the naming game simulation can resemble a cumulative distribution function (S-shaped).

2.3 Iterated Learning Model

While animals often use innate communication systems, humans' mappings between signals and meaning are not bound by genetic constrains. Oliphant (1999) presented a computational model for investigating this shift — going from a simple innate communication system to a system based on learning — in a population of agents where no reinforcement signal is used (in contrast to most language games). Kirby (2002b) follows this paradigm in a cultural setting, by addressing the concern of human language to solely being a result of natural selection. He presented a tool called the *iterated learning model* and applied this to different linguistic characteristics on a cultural timescale. The most interesting findings must be the emergence of compositionality (Smith *et al.*, 2003). Under the assumption that there exists some innate ability to express words, the model shows that the structure of language can

Speaker	Hearer
Gral	Qjar
Kron	Tram
Fjap	

Speaker	Hearer
Gral	Qjar
Kron	Tram
Fjap	**Kron**

Table 2.1: Schematic rules for a dialog failure. The speaker utters the word “kron” which the hearer is not familiar with. The hearer adds the transmitted word to his inventory.

Speaker	Hearer
Gral	Qjar
Kron	Tram
Fjap	**Kron**

Speaker	Hearer
Kron	**Kron**

Table 2.2: Schematic rules for a dialog success. The speaker utters the word “kron” which the hearer is familiar with. Both agents clear their inventory for other words.

be built on top of a strictly cultural level — even when the poverty of stimulus is present (Chomsky's (1980) idea that a language can be fully learned with minimal input, and that there exists innate language "organ" that supports this). In their model, adults expose their children to spoken language (so-called e-language, or external language) and the children acquire this (or a similar) language (i-language, internal language). The learners rarely hear all possible utterances and are therefore victims of a *transmission bottleneck*. As the language travels through this bottleneck multiple times and more meanings are being transmitted, the language adapts to be generalizable and to contain more structure — thus, the transmission error rate goes down and the language becomes more learnable. However, most of their simulations where conducted with relatively few individuals.

These findings have been supported from an empirical point of view. It has been argued that languages tend to be morphologically simpler, less redundant and more regular as the population increases and the community undergoes iterations of learning (Lupyan & Dale, 2010).

2.4 Network Theory

Before reviewing language games in more realistic graphs (next section), let first elaborate on theories behind networks and why they matter. In recent years, an increasing amount of research has been devoted to the ecology of multi-agent language games. Instead of distributing the agents in a regular lattice (Lipowska, 2011; Baronchelli *et al.*, 2006c; Lu *et al.*, 2009) or not structured at all (Noble, 1999; Baronchelli *et al.*, 2006b), the focus has shifted towards social networks (Gong *et al.*, 2004; Lipowska & Lipowski, 2012).

In modern network theory (very recently also termed "network science") the most prominent discovery must be the fact that universality of real network topologies converges to similar architectures across a lot of disciplines (Barabási, 2009; Baronchelli *et al.*, 2013). Many regularities in social science have been uncovered by studying the complex system through network theory (Baronchelli *et al.*, 2013). Other applications which have analytically benefited from network theory is for instance Internet traffic, aviation, epidemic spreading, and free word associations (Hernando *et al.*, 2010) (see Costa *et al.* (2011) for an extensive survey). All of these share the characteristics of non-trivial topologies, as they follow either random behavior, nor regular patterns.

2.4.1 Definitions

Thus, networks that hold such non-trivial topologies are termed *complex networks*. A network is in physics an applied *graph*, which is the mathematical term. A graph is a set of *nodes* (often referred to as vertices) that are connected to each other by *edges*. A particular node can have zero, one or multiple edges to other nodes, and the number of edges is called the node's *degree*. The *neighbors* of a given node, n, are the nodes connected to n by edges. Further, edges are either *directed* (as arcs) or *undirected*. If not specified, edges can be assumed to be undirected when spoken about through the rest of the thesis. An edge's intensity can be represented by adding *weights*.

Given two random nodes, their *shortest path* is intuitively the shortest path between the two nodes through the graph. If the graph is weighted, the lowest weighted path can also be used.

2.4.2 Measurements

A graph's *degree distribution* is the probability distribution over the nodes' degree through the entire graph. $P(k)$ is then the fraction of the nodes who degree is k. Each node has an associated *clustering coefficient* which represents how concentrated the neighborhood is (which corresponds to the term *transitivity* in social sciences (Baronchelli *et al.*, 2013)). The clustering coefficient for the node n is defined as

$$C_n = \frac{\text{number of edges between neighbors}}{\text{number of potential edges between neighbors}}$$

If then $C = 1$ for a particular node, every node in its neighborhood would be connected together. In a *complete graph*, where every node is connected to every other node, the average clustering coefficient would naturally be 1. Degree is a measure of *centrality* in graph theory. There do exist other measures under the same category, such as betweenness (how often a node, or an edge, is included in the total set of shortest paths) and closeness (to calculate a nodes independence) (Freeman, 1979) — but those are perhaps more important for statistical approaches.

Different types of graphs can be categorized by their characteristics. The *random graph* Erdös & Rèyi (1959) model can be viewed as the simplest complex network — as edges are added between each possible pair of nodes with a probability p. This way, both the degree distribution and the average clustering coefficient can be easily predicted. Unfortunately, random networks are homogeneous and not very realistic in terms of modeling real world systems. In contrast, the Watts & Strogatz (1998) model, often referred to as the *small-*

world model has shown more resemblance. This model is characterized by having a low average shortest path, as illustrated with the famous "six degrees of separation" concept — while still having a high clustering coefficient. A small-world network can be constructed by initially adding all nodes to a lattice with k nearest neighbors and then rewire every edge with a probability p. There is also the *scale-free model* presented by Albert & Barabási (2002). This model is driven by *preferential attachment*, meaning that new edges are attached to nodes proportional to their degree. These networks naturally get a degree distribution which to some extent follows a power law

$$P(k) \sim k^{-\gamma}$$

where γ typically is $2 < \gamma < 3$ (Newman, 2003).

Although a network can be classified as either a random graph or a scale-free graph, both can also resemble and be termed as small-world graphs at the same time.

2.5 Language Games on Social Networks

The relationships among humans are of course so complex and intricate that no mathematical models are anywhere near verifying every choice of trust, communication, partnership, or religious beliefs. Ambitiously, the first systematic studies of social relations began around a century ago, and since then both books and whole journals have been devoted to social networks (Costa *et al.*, 2011). Some may argue that there has been a whole industry committed to the field in the last, or the last couple of decades — following the exponential growth of technology, Internet and social media.

The last decade, simulations on language and language change on social structures have picked up a lot of attention. There is of course a vast number of social features that could possibly influence linguistic (Gong *et al.*, 2008), such as age, gender or religion — but with both careful and bold abstractions, researchers keep exploring.

The paper by Baronchelli *et al.* (2006a) is a good example of the dynamics involved in a naming game conducted on a social network. In contrast to ordinary non-structured language games, with high amounts of total words and relatively fast consensus time — on scale-free networks, the total amount of words in the simulation is rather low, but reaching consensus takes much longer time. This is mainly because, when choosing speakers at random in scale-free networks, there is a high probability that agents with a low degree are picked. When the majority of the hearers are agents with high degrees, words do not spread that easily — and when agents with high degrees occasionally are chosen as speakers, they tend to have such large inventories that their contribution is of little worth. Thus, when choosing speakers in a naming game on a social network, the random agents should always be chosen proportional to their degree.

An other approach to minimizing the consensus time in such simulations is by manually making some long distance edges between random nodes in the network (Dall'Asta *et al.*, 2006). If these connections are not too long, this favors agreement in small-world networks. This has also been supported when conducting similar experiments on empirical gathered data (Lu *et al.*, 2009). In this study, the social networks (high school friendships) where manually altered by adding a small number of long distance edges between communities to mimic a mass media effect in a naming game. They also experimented with simulations where some agents in the population insisted on uttering only one certain word throughout the naming game. In both cases, these alterations induced a much faster consensus among the agents — as without, the population consensus was rarely reached. Whether or not the analogy towards a mass media effect is coherent or not, the study highlights at least some of the dynamics and tensions in social networks — as *tipping point*, where small influences can make sudden shifts in the system (Helbing & Balietti, 2011).

As for what could influence language on a social network, take for example learning (Maity *et al.*, 2013b). In this paper, a naming game was attempted on both homogeneous fully-connected networks as well as heterogeneous scale-free networks. In both cases, results show that learning hinders the consensus and instead facilitates the emergence of small and strong communities.

Recently, (Lipowska & Lipowski, 2012) examined a naming game on an fully connected network. Here, every relation, has an associated weight. Through dialogs, the weights get increased and decreased on dialog successes and failures respectively, with a small user-given ϵ. A speaker is chosen at random. The speaker chooses among all the agents in the population for a possible hearer, but proportional to the weights in the network. In most cases, the simulations converges against a handful of words, spread out in contracted clusters of different sizes.

It is said that the dynamics of a complex network comes from the interplay between two factors, the state of the network and the topology of the network (Gross & Blasius, 2008). A network's state can be thought of as the nodes' values or features, and for the latter, the pattern of connections between the nodes. When the nodes' states change, this can spark a topological change — and a change in the topology can again change the states of different nodes. This can happen in both local parts of the network, or in the network as whole. Although there have been several endeavors on co-evolving topologies with other ecological phenomena — for instance in epidemic spreading (Marceau *et al.*, 2010), or opinion dynamics (Nardini *et al.*, 2008) — few simulations on language evolution have done the same (Quillinan, 2006); experiments are instead mostly carried out on static topologies.

2.5.1 Social strategies

The animal kingdom is full of social learners. Social learning, in contrast to individual (or asocial) learning, is often assumed to be preferred as it is a cheaper and safer way to new information. Rendell *et al.* (2011) presented several different learning mechanisms; random copying, strategies that are dependent on others choosing the same strategies, strategies where individuals are only pointed in the right direction, strategies that depended on age, gender, or dominance ranks and so on. The list is long, and it emphasizes how intricate modeling culture can be in computer science and game theory. Even so, under certain circumstances, social learning might not be more favorable then individual learning, as the information gained through social learning might be wrong or outdated — a phenomenon referred to as "Rogers' paradox" (Rendell *et al.*, 2010; Hashimoto *et al.*, 2010).

The literature on self-emergent language games is full of examples where agents uses social learning. For instance, Maity & Mukherjee (2012) reflected dominance by choosing speakers with a probability proportional to their communication success rate. In a more thorough fashion, Gong *et al.* (2008) have experimented both having more popular agents among "ordinary agents", and the balance between short and long distance communication. As expected, having popular (and thus dominating) agents helps towards faster agreement. Further, having multiple hearers has been tested (Li *et al.*, 2013), as well as overhearing (Maity *et al.*, 2013a), both facilitating faster agreement in the population and significantly decreasing cognitive efforts in terms of memory capacities.

2.6 Genetic Algorithms

A genetic algorithm (GA) (Holland, 1992) is a heuristic optimization algorithm which mimics the Darwinian theory of natural selection. The idea works as follows; a population of possible solutions to a particular problem is randomly generated. Each individual in the population is evaluated by an objective function, hereafter suitably called a fitness function. Over generations, poorly fitted individuals die and the surviving individuals are allowed to reproduce. Through genetic operations and the variation that they bring, new solution candidates are born and new fitness values can be calculated.

2.6.1 Definitions

In more technical terms, a GA can be viewed as a set of *genotypes* analogous to a part of a species' DNA. These genotypes are mostly represented as strings of bits where each group of bits can be interpret as a gene. *Grey-encoding* is one of the most used mechanisms for doing this. It is a straight forward casting from bits to real-values, but with some alterations so that a single bit flip mutation always can both make slight and large changes to the decoded value. As human DNA works as a blueprint for humans, a genotype works as a blueprint for its *phenotype*. Therefore, a phenotype in a GA can be one of many candidates to a solution to for instance, an electronic circuit design, the traveling salesman problem, or other problems where the search space is huge and a suboptimal solution is good enough. Initially, a random set of genotypes is generated and *developed* into a phenotype. Each gene in the genotype is typically interpreted as a real-value or discrete building block for the phenotype. A much used application is the construction of artificial neural networks, where genes are interpreted as weights to the neural network (Belew *et al.*, 1990). Further, the phenotype can undergo some interaction with its environment before a *fitness* is calculated, rating the candidate solution in terms of objectives. How the phenotype adapts to an environment is termed *phenotypic plasticity*, thus high plasticity makes a phenotypic less constrained by its genotype. When the first random generated candidates are tested and rated, a *generation* is over. Here, a handful of the best fitted solutions are chosen by some *selection strategy*, typically by a *roulette wheel selection* or *tournament selection* to get a good and biased balance of both fitted solutions and solutions that are heading towards other (and perhaps more profitable) parts of the search space. A roulette wheel selection is when an agent is chosen at random, but proportional to its fitness. A tournament selection is when k agents are selected for a "tournament" and n agents "win" by some metric. As these solutions are chosen, they reproduce by combining their genotype in ways known to work in mother nature. The combination is generally done by *crossover*, an operation which splits a pair of genotypes in two or several parts, often between genes, and weave them together to new genotypes. Occasionally, the genotypes are victims of mutations where one or more bits get flipped. These new genotypes then get developed to new phenotypes and the algorithm runs through generations until the solutions start to converge.

While this sounds like following a recipe, the challenge is choosing suitable development and fitness functions. Development can be done in several ways, but the key is to make the mapping between genotypes and phenotypes so that small mutations in a genotype also makes small changes in a phenotype. Another important matter is to make sure all phenotypes in the search space can be mapped from the genotypes. As for the fitness function, it is highly dependent on the problem — but the essence is to make a function which can output an arbitrary number for how good a solution is (objective), or how much better it is then something else (subjective).

Overall, a GA provides a set of operators that can be used to achieve exploration (crossover), exploitation (mutations), and both at the same time (with different selection strategies). The term *fitness landscape* is extensively used with in the research field of evolutionary computation. There is not any distinct definition of what a fitness landscape is, it is mainly used figuratively to describe how “peaks”, “valleys”, and “ridges” can look like in vast search space (Jones, 1995). The landscape intuitively enough resembles where the global and local optimum or optima are located in terms of fitness. Particular problems for GAs are often described as how their landscapes look like — for example, from unimodal to rugged, with or without deceptions, or isolated “needle-in-a-haystack”, regions with misleading gradient information. They are all typical problems in that must be accounted for by tuning the GA in different directions (Weise *et al.*, 2012).

2.6.2 Applications

Genetic algorithms have many applications other than optimization. They can for example also be used for programming, where genotypes work as programing operators, or in economics for development different bidding strategies, or in simulating ecology in terms of host-parasite co-evolution and disease patterns (Mitchell, 1998). However, they can also translate into a framework for social behavior, agent-based systems, and game theory. While it might not translate directly, the methodology is still valid. For instance in the language evolution territory, an objective function is not necessarily the best choice as a fitness function — since the best communication system might not be known. Thus, a *subjective function* can be applied to value a solution candidate relatively to other solutions by some measure. An example of this is reviewed later on (Suzuki & Arita, 2013).

Surprisingly few have used GAs to investigate self-organizing language systems in agent-based simulations. Noble (1999) used a GA in a “Prisoner’s dilemma” inspired signaling game where cooperative and non-cooperative strategies evolve in an non-situated environment. A recent paper from Gong & Shuai (2012) used a typical GA to model the co-evolution of joint attention (the non-verbal activates that often aid communication, for instance pointing) and language over both cultural and biological timescales. They demonstrate that assimilating not-language specific abilities, such as different strategies for joint attention, could support the emergence of language. Another interesting paper is the one from Munroe & Cangelosi (2002). They simulated agents from an ecological perspective where agents were using neural networks to evolve a simple grammatical language. The agents applied this language while walking around in a situated world filled with edible and poisonous mushrooms. The agents needed to agree upon what to eat and what not to eat, by perceiving information in the environment and sharing this information with their offspring.

The simulation was split up into two phases. First, a number of agents, each provided with a random weighted feed-forward neural network, were placed in an environment for a finite number of generations. Here they formed into a set of agents capable of finding, handling and eating edible mushrooms. After this phase, the best genotypes (essentially the neural network's weights) were selected on the basis of the phenotype's fitness and stepped in as parents for a new group of visually impaired agents. The new group of agents could only perceive the mushrooms and their features 10% of the time, and depended on the caretakers output language to survive. From this point, the newborn had to learn what features made mushrooms edible or not throughout listening to their parents instead of just visually perceiving. They adjust their neural network weights with backpropagation at each time step (analogous to how humans do individual learning, by trying and failing in general tasks). Before the children become parents, they are allowed to visually perceive their whole surrounding for some iterations, just enough to be get a better grip of what visual features map to what linguistic signals. Over generations visual and linguistic features got more and more assimilated into the genome and offspring got less and less dependent on learning. By tweaking the simulation parameters, as (noise in) cultural variation, the Baldwin effect and learning costs can be encouraged and damped. The compositional language in this model is viewed as clusters of signals, where the different clusters server as different word types and signals as different words (edible, poisonous, wash, cut, ...). It is an interesting model because it is one of the first simulations carried out with a GA and where both the biological and cultural aspect are represented in a situated, artificial life setting. As written theories tend to be spoken about in very general terms, this model is good because it gives a somewhat clear picture of how a concrete language can evolve.

From a more abstract perspective, Suzuki & Arita (2013) wanted to investigate how negative effect of frequency-dependent selection could be tackled by evolving phenotypic plasticity. Frequency-dependent selection in biology terms is when the fitness for a phenotypic trait gets greater when the trait becomes more common (positive) or less common (negative). An example used in the paper is the hypothesis that it can be hard for mutations in a hard-wired language system to be beneficial if fellow participants do not understand the change. In their GA, they let the individual's genotype contain two real-values, one for an arbitrary trait and one to express the phenotype's plasticity. Communication success between agents is measured by the distance between the agents' trait values, and fitness is calculated from this. Plasticity in this model allows the agents to adaptively skew their trait value (learning) towards better fitness through their lifetime. Overall, Suzuki & Arita (2013) show that learning allows the population to avoid getting stuck in local optima.

Chater *et al.* (2009) uses a similar model when they argue that only stable parts of the linguistic environment encounter the Baldwin effect. In their simulations, they demonstrate that arbitrary linguistic properties in human language (such as word order, grammatical agreement, and case) do not assimilate into the genetic makeup purely through adaptation if cultural variance is too high. Their genotypes are made up by genes, where each gene can have one of three discrete values; positive, negative, or a neutral. The values represent biases towards a certain language principle, one bias for each principle. If an agent (has a polarized bias towards the same principle used by other agents, there would be no need for learning (high fitness) — but if the innate bias and the language principle are conflicting, the communication will not succeed. Their experiment goes over three simulations, first, language is assumed static and unchanged through generations. All the agents quickly adapt to having genes favoring this language and thus the Baldwin effect is operating. In the second simulation, the target language is changing between generations by flipping random principles to represent cultural change. This dampers or eliminates the Baldwin effect completely, depending on the change rate — and all agents adapt to having only neutral biases towards the target language, relaying on a learning orientated communication system. In their last simulation, they let the target language be influenced by the genetic biases among the agents to if the Baldwin effect reemerges. Hence, it is only by "implausible" high values of genetic influence the Baldwin effect emerges. These are the same people who proposed that language is shaped by the brain instead of vice versa. They have later used the same methodology where populations split up and evolve side by side and later merge without any problem, similar to how Australasian populations have been largely isolated for 50,000 years but still have no problem learning European languages (Baronchelli *et al.*, 2012).

Chapter 3

The Lipowska Model

As this thesis wants to investigate the different evolutionary forces shaping human language from a computer science perspective, it is appropriate to first look at previous studies where similar questions have been raised. There have been other endeavors (Nowak *et al.*, 1999; Hashimoto *et al.*, 2010; Munroe & Cangelosi, 2002), but the paper by Lipowska (2011) sticks out. It is an interesting model because it was motivated by the combination of a model from the interactive language game paradigm (which is focused on static one-generation interactions) and the iterated learning model (Kirby, 2002a) (which concentrates on the interactions between parents and children). On top of this, the model looks at how learning can affect the emergence of a self-organizing language system over the different timescales involved; the biological, the cultural and at an individual level.

Following are three sections; method, results, and discussion. The method includes the basic setup Lipowska used in her experiments. This method will be implemented and results will be presented. The discussion will explain some problems with the methodology and present arguments for extensions.

3.1 Method

As this model adopts a language game, namely the naming game, it is said to be non-structured (no grammar). However, it is to some degree situated as all of the agents are structured in a regular lattice with size L.

The simulation unravel as follows; at each iteration i

1. a random agent is chosen, and
2. the agent is allowed, by a user-given probability p, to

   initiate a dialog with a random neighbor and act as speaker, or by the probability $1-p$, face a “population update”,

Thus, this p is of essence, being the ratio between communication attempts and the population turnover. If p is small, there would be little interaction in the simulation — and if it is high (or simply higher), agents have statistically more dialogs throughout their lifetime.

As with naming games, every agent has an internal lexicon consisting of words. Here, every agent has weights, $w(w > 0)$, associated to their words. Whenever a chosen speaker is to utter a word w, the agent selects a word w_i from its lexicon with the probability of $w_i/\sum_j w_j$. If the lexicon is empty, an agent is assumed to have the ability to make up a new word. One of the nearest nine neighbors in the lattice is chosen as the hearer. If both agents know the transmitted word, the game is deemed a success and if not; a failure. Upon success, both selected agents increase the transmitted word’s weight in their lexica by a learning ability variable (see Tables 3.1 and 3.2). This is an alignment strategy that permits this learning ability to be favorable (and even possible), hence it also preserves the culture variance as word do not get deleted upon dialog success (like in Section 2.2.7). Each agent is equipped with a learnability variable l $(0 < l < 1)$. An increasing learning ability is meant to, in its simplicity, reflect the genetic assimilation of a language learning apparatus, as discussed in Subsection 2.1.3.

Apart from the communication, the chosen agent is occasionally updated, with the probability $1-p$. This is to establish the generation turnover emphasized in the iterated learning

Speaker ($l = 0.6$)	Hearer ($l = 0.8$)
Gral ($w = x$)	Qjar ($w = x$)
Kron ($w = 1.8$)	Tram ($w = x$)
Fjap ($w = x$)	

Speaker ($l = 0.6$)	Hearer ($l = 0.8$)
Gral ($w = x$)	Qjar ($w = x$)
Kron ($w = 1.2$)	Tram ($w = x$)
Fjap ($w = x$)	**Kron** ($w = 0.8$)

Table 3.1: Schematic rules of a dialog failure. The speaker utters the word “kron” which the hearer is not familiar with. The hearer adds the transmitted word to its inventory with its learnability, $l = 0.8$, as weight — and similar, the speaker subtracts $l = 0.6$.

Speaker ($l = 0.6$)	Hearer ($l = 0.8$)
Gral ($w = x$)	Qjar ($w = x$)
Kron ($w = 1.2$)	Tram ($w = x$)
Fjap ($w = x$)	**Kron** ($w = 0.8$)

Speaker ($l = 0.6$)	Hearer ($l = 0.8$)
Gral ($w = x$)	Qjar ($w = x$)
Kron ($w = 1.8$)	Tram ($w = x$)
Fjap ($w = x$)	**Kron** ($w = 1.6$)

Table 3.2: Schematic rules of a dialog success. The speaker now utters the word “kron” which the hearer is familiar with after a dialog in the past. Both agents add additional weights to the utterances with their learnability learnability l.

model. Agents die or survive with the probability

$$p_{surv} = exp(-at)[1 - exp(-b\sum_j w_j/\langle w\rangle)] \tag{3.1}$$

by that takes age (t), knowledge (lexicon weights in respect to the population's average weights $\langle w\rangle$), and user-given coefficients into account (typically $a = 0.05$ and $b = 5$). Intuitively, this directly favors the agents' learnability variable if there are any successful dialogs. How age is calculated is not formally presented in Lipowska's (2011) paper, therefore the assumption is made here that every time an agent is allowed to communicate, its age gets incremented. If the agent has a high-weighted lexicon and is young of age, the agent will most likely survive. If the agent does survive, the agent is allowed to breed if there are empty spaces among its neighbors. If an agent reproduces (asexually), a new agent spawns in the empty cell and inherits the parent's learnability with the probability $1 - m$. m is here the mutation rate, so if a mutation comes about, the new agent acquires a new random learnability value. Lastly, the offspring inherits the parent's highest weighted word with the same probability $1 - m$, and can otherwise make one up.

Lipowska (2011) goes a bit further and experiments with a gradually increasing p in some of her simulations, as well as seeding the population with pre-assimilated agents. These experiments will not be implemented here, but discussed afterwards.

3.2 Results

Lipowska (2011) uses different parameters in her model. 0.2–0.3 is used for p , $a = 0.05$, $b = 5$, $m = 0.001$, $L = 60$. Following are some figures (3.1 to 3.4) from a simulation run with similar given parameters as in (Lipowska, 2011); However, with $p = 0.4$. The high p is due to the fact that no significant progress is taking place with values lower then $p = 0.3$ (similar to Lipowska's value $p <= 0.15$, many words and small clusters). Moreover, 200,000 timesteps are performed, the chances of learnability mutation is $m = 0.01$, and lattice size $L = 25$ (625 agents). Another factor that are not clearly stated are the initial values. Lipowska has used different initial values for learning ability in her analysis — here, $l = 0.5$ is used, which makes a decrease in learnability a part of the search space as well. For initial weights on invented words, $weights = 0.5$ is used. All parameters which is not exactly the same are skewed in favor of a sharper coherence.

In Figure 3.1a to 3.1e, each agent is a dot in each "heat map" and each dot is colored to represent an agent's value. Black dots indicate a dead agent, or in other words, an empty space. The left column is the agent's learnability, the center column is number of words in

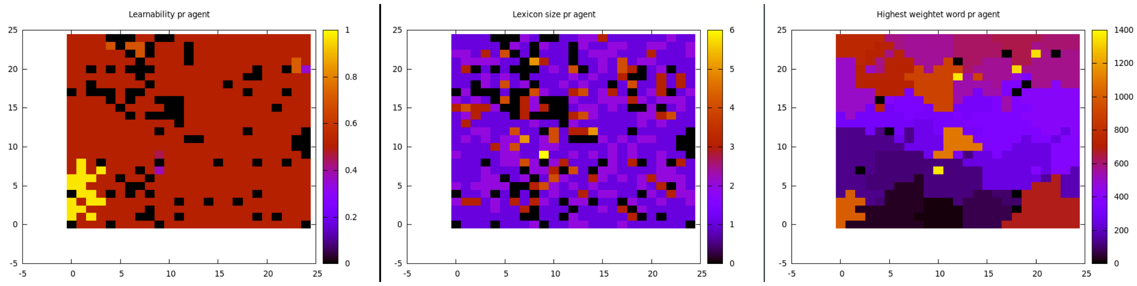

(a) After 10k timesteps

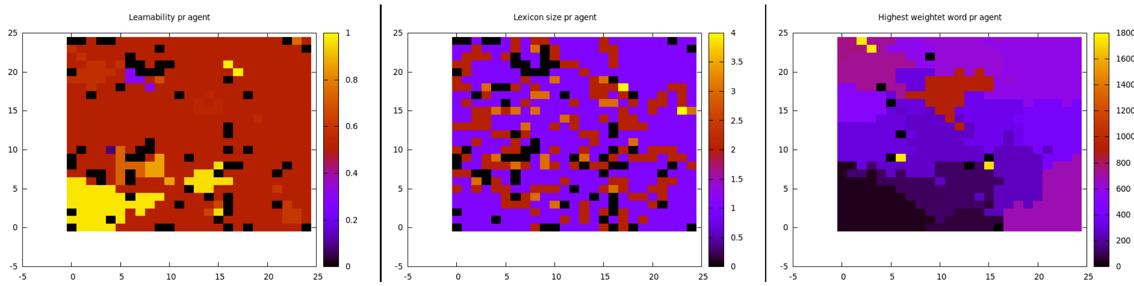

(b) After 50k timesteps

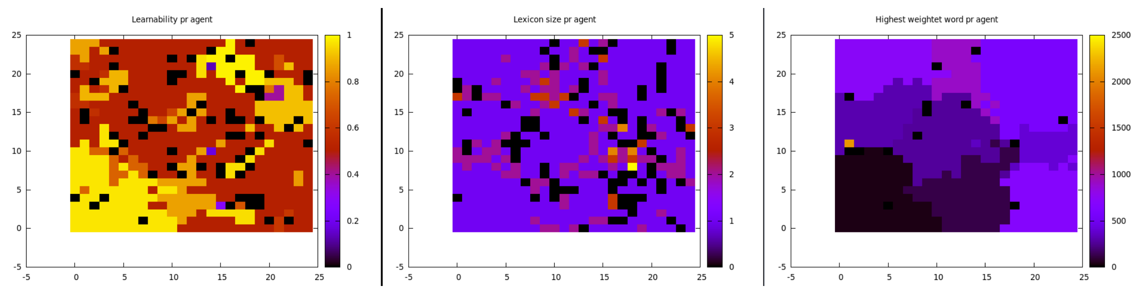

(c) After 100k timesteps

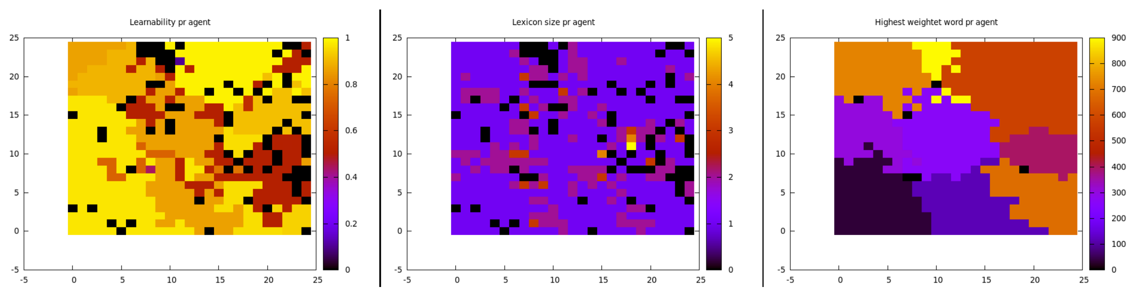

(d) After 150k timesteps

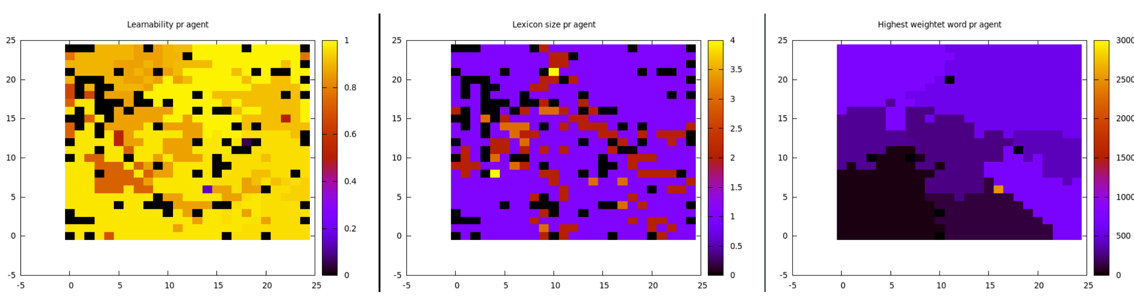

(e) After 200k timesteps

Figure 3.1: Naming game simulation. Every diagram is a heat map, representing some value from every agent in the population. The left column shows learnability, the middle column shows lexicon size, and the right column shows the highest weighted word in the agent's inventory (the colors are arbitrary).

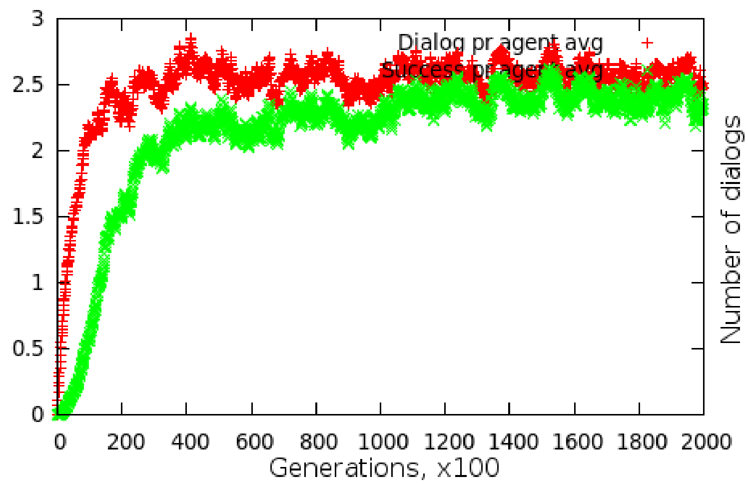

Figure 3.2: Average number of communication dialogs per agent. The x-axis is one point per 100 timesteps.

an agent's lexicon, and the right column represents the highest weighted word per agent (the colors and numbers are somewhat arbitrary — the interesting facts here are where the borders go, and how the clusters characteristics over time). Already from timestep 10k, small groups of agents agree upon different words (as seen in the right column), this abrupt transition from disorder to order is mainly because of all the deaths. In only 10,000 iterations, a statistical 60% of random agents have faced death through Equation 3.1 (since $p = 0.4$), whereas many have not survived, judging by there age and knowledge. With these, a lot of entropy vanish. As mutations start to spread (left column), the groups turn fewer and their sizes grow. Mutations that do not favor learnability quickly lead to death. In the center column, it is easy to see where the different groups meet. The agents in these areas naturally know both words on each side (or more) and can abstractly be seen as multilingual. By timestep 200k, the population do not reach a single word agreement (as in most of Lipowska's simulations), but consensus get reached if the simulation is allowed to run further.

In Figure 3.2, the upper line represents the average number of communication dialogs per agent alive. The other line represents the average number of successful dialogs. The number of successful dialogs quickly follows the number of total dialogs after small linguistics groups settle. The average age of each agent correlates with the upper line as it is measured from the number of dialogs it participates in.

In Figure 3.3, the upper line represents the fraction of agents alive in the lattice. It is initially fully populated, but the population decreases with time and balances at a point where both death and birth are equally tensioned. The other graph is here the average learnability in the population. This is s-shaped as a result of slow mutation rate in the first

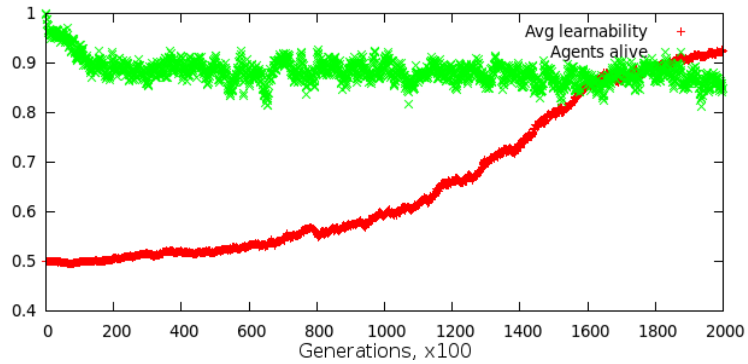

Figure 3.3: Fraction of agents alive and average learnability. The x-axis is one point per 100 timesteps.

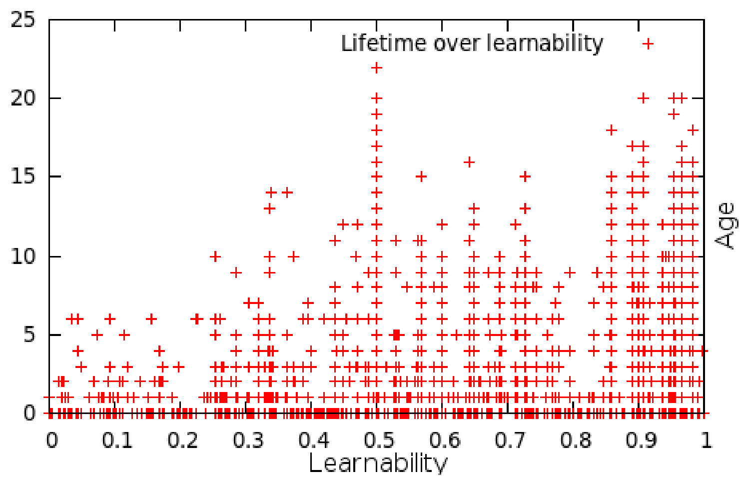

Figure 3.4: Agent's lifetime as a function of learnability. The x-axis is one point per 100 timesteps.

phase (few die of old age and no space to breed), accelerated around 150k as the favorable learnability also gets inherited, and decreasing towards the end because a mutation is more likely to ruin an agent learnability then (when l is at its upper limit).

In Figure 3.4, the agent's lifetime as a function of learnability is plotted. Since the surviving probability, p_s, is influenced by the total number of words and weights in the lexica, agents with higher learnability tend to live longer.

3.3 Discussion

Firstly, in order to replicate Lipowska's (2011) results quite a few parameters where altered (for instance, higher mutation rate, higher possibility of survival with richer lexicon/higher age). Presumptively, this is due to how age is calculated and how this affects the chance of survival. In the replica, age is assumed to correspond to the number of dialogs an agent participates in. Another possibility is to increment every agents' age at every iteration, or at every time there is a dialog, but then the results divert even more — and it makes little sense in terms of small values of p (very little chance of communication). With that said, having such a model where the calculation of age is so crucial seems curious. For instance, if age is not carefully weighted in the survival equation (Equation 3.1), more agents might die than agents that are allowed to breed. This effect is further enhanced if p is changed, as more population updates make agents confront death more often. Furthermore, for different values of p, the number of iterations needs to be adjusted to reach the same amount of dialogs. This is also the case when changing the population size — if the population for instance is increasing, one particular agent is less often chosen. Lipowska herself has compensated with higher mutation rate on simulations with different population sizes, which could be two variables somewhat more independent from each other in an ideal simulation.

At every iterations, or timestep, the simulation follows a sequential technique — where each iteration can end in multiple closings. For instance, if an agent is chosen for population update and it survives — it is allowed to breed, but if there are not any empty cells around, the iteration just ends — and the simulation meets this endpoint almost 25% of the time. From earlier literature on naming games, where only cultural transmission is simulated on static populations, this timestep design is quite reasonable as it only has two outcomes per dialog, success or failure. Reusing this design when expanding a from cultural inspection only to both culture *and* biological analysis can thus be arguably messy.

If this naming game paradigm is the right tool for abstractly studying semiotic dynamics on a phylogenetic time scale, generation turnover needs to be more vigorous than happening in between dialogs. While this certainly is more lifelike in a contemporaneous universe, it makes it difficult to make good computational models for further analysis. A typical genetic algorithm could be a reasonable substitution, and could make the model more open for studying other innate traits. This is the opening argument for the model presented in Chapter 4.

As for the overall intention, illustrating the Baldwin effect, the model captures the essence. Given metastable clusters of words (and thus a high average communication success rate), agents with higher learnability are not killed as often as those who are without — meaning they have greater possibilities for acquiring an empty cell for breeding. Although survival is calculated upon their performance (or acquired characteristics), its is rather obvious that learning is favorable. At least it is, when the linguistic performance is normalized on the whole population. It is arguably less self-emergent if the fitness of an agent is valued in relation to other distant and non-threatening agents.

However, Lipowska's main arguments are sound — it is clear that the ratio between communication and population update is significant. There exists a balance, in which she demonstrates in her paper, between cultural and biological transmission. Or said in other words, how much one speaks before dying. This directly influences the relative variation on both levels, which is the allover topic in the Baldwin effect debate in language evolution (Chater *et al.*, 2009).

Chapter 4

Methodology

Given the discussion in Chapter 3 (see Section 3.3) this thesis will examine a similar model in more robust framework; a traditional genetic algorithm (in the sense of Section 2.6, hereby just called a genetic algorithm). While Lipowska's (2011) model can be seen as a genetic algorithm, or at least is inspired by genetic algorithms, was constructed to show a specific point and not very suitable for extensions. When exploring with learning, and other innate traits, and how it influences agents' lifespans in various environments, it is important — in an computational model — to keep track of causality and make the right abstractions based on this.

This chapter has three parts. Firstly, a brief clarification and justification is given for the choices made in this chapter. Secondly, the experiment from Chapter 3 will be implemented using a genetic algorithm. And finally, a similar naming game will be demonstrated using a genetic algorithm where other innate traits are taken into account, namely social strategies and some cognitive constraints. This is done to highlight the possibilities of using a genetic algorithm while simulating aspects of language evolution.

4.1 Choices

Using a genetic algorithm simulating language evolution in artificial life will be justified by the following two observations. Firstly, it is clear that language evolution must be looked at, and dealt with, from three different viewpoints, namely the evolutionary forces surrounding biology, culture, and choices done on an individual level (through learning) (Steels, 2011). A good model for this must be capable of distinguishing between these co-existing adaptive systems and somehow be able to give insights into the interactions between them. A genetic algorithm is able to take this into account (Bedau, 2003; Hashimoto *et al.*, 2010). Assuming that learning, or a language learning apparatus, is assimilated and can be used — there may exists many other traits facilitating communication indirectly, such as social strategies **??**. A species' assimilated (carefully chosen, abstract) biases and strategies for learning, social interactions and other cognitive abilities, can be captured with a vector of bits — a genotype. From this genotypes, development takes place at each individual's birth. Then, depending on the cultural environment, the individuals acts upon these strategies throughout their life — as phenotypes. In other words, biology can be mapped to genotypes, culture to the environmental interactions, and the individuals to phenotypes. Although life is continuous and concurrent, a genetic algorithm applies temporal turnover by iterating generations. As a genetic algorithm originally is a search optimization algorithm, it is common to only take a few selected individual from generation i_{t-1} and let them breed and mutate extensively to a fully populated generation i_t. By skewing these parameters towards fewer deaths per generation and fewer interactions between each generation, a more realistic continuing community of artificial life can be achieved. Secondly, in the reviewed literature (Chapter 2) it has been mentioned that a genetic algorithm would be a step in the right direction for simulating different aspects of language evolution on social graphs, and no one seems to have picked it up. A genetic algorithm was proposed as future work in the paper by Gong *et al.* (2004), in their experiment where simple rule based syntax was evolving on different social networks. Another example is the thesis by Quillinan (2006), where one-generation social networks co-evolve with an abstract model of language and learning, the same proposal is pointed out as an extension.

4.1.1 The Lipowska Model using a Genetic Algorithm

In Section 4.2, the experiment from Chapter 3 will be outlined using a genetic algorithm. This is done for the purpose of demonstrating that the problems easily translated from a sequential timestep model to a model where genotype and phenotypes are interacting between generations, although it might over-engineer particular problem. This is justified as it makes the model more suitable for extensions.

4.1.2 Co-evolving Communication Systems and Social Structure

As there is good indications for communication systems to affect social structure (Lipowska & Lipowski, 2012) and vice versa (Quillinan, 2006) this will outline a model for investigating such experiments in greater detail.

For the type of computer models (Subsection 2.1.4) this work is aimed at exemplifying and highlighting the methodology towards this field of study. However, both motivation and inspiration is gathered from the recent literature, so the overall contribution has an explorational touch to it.

To test and demonstrate the model, a total of seven different single-run simulations will be conducted, where four innate traits are investigated.

- **Extraversion**. How far out in the social graph do an agent search for new friend. This is inspired by Gong *et al.* (2008), where agents have a user-given probability for speaking with other agents in either their own community or cross-communities. In the discussion, they argue that these user-given probabilities should co-evolve with the language evolution and social structure. This extraversion gene is a response to that discussion. The gene's value will be treated as probabilities to reflect innate biases and not deterministic innate behavior.

- **Teach child**. This value reflects the probability that a parent is the first one to talk to the child. If both parents have high values, a child gets its first dialogs from them — and if the values are low, the child would have to make its first dialog with a random agent. This was tested in a paper by Nowak *et al.* (1999), but not on a social graph. They showed that learning from your own parents is highly favorable, but it could be interesting to see how a social structure would respond to such behavior (or lack of behavior). The gene's value will also be treated as probabilities to reflect innate biases and not deterministic innate behavior.

- **Lexicon limit**, a value which constrains the number of words an agent can remember at once, as it is a rational for an agent to not have unlimited (computational) powers (Epstein, 1999).

- **Speech ability**, a value reflecting the agent's reliability of uttering the actually chosen word, and not a random one. This is because it is a bit odd that every agent spontaneously is able to make distinct difference between a conceptual thought and a holistic sound. This particular trait will only be added in one of the simulation, to see how it affects the evolutionary process.

4.2 The Lipowska Model using a Genetic Algorithm Defined

Going from the timestep model from from Chapter 3 to a traditional genetic algorithm is relatively straight forward. The learnability translate to a single gene, encoded as a real values from 0 to 1 using Grey encoding. However, the genetic operator crossover usually splits between genes, and is therefore transparent in this case. Mutation m at birth works as expected in a genetic algorithm (one bit gets flipped). Regarding population turnover, a simple tournament selection is applied, k random agents are chosen and n of the best fitted agents "win" the tournament. Between generations, $d \cdot populationSize$ dialogs are randomly distributed. Every speaker chooses a random neighbor and act accordingly to the experiment in Chapter 3. For the fitness function, p_{surv} is used (Equation 3.1), where age is corresponding to the number of generations an agent has survived.

4.3 Co-evolving Communication Systems and Social Structure Defined

This model can be split up in into two co-occurring processes, namely evolution of language and evolution of social structure. The language evolution is adapted from the naming game, as in the model in Chapter 3, except for learnability. In human language, it is not necessary true that language arose simultaneously as with social structure. There is no clear answer to what came first, and if, to what degree language worked as a catalyst to social structure, or inversely. However, Dunbar (1997) argues that the two strongly correlates, and if any connection, communication is perhaps the most natural cause of humans social structure. Thus, in this model it is assumed that both social structure and a language arose coinciding. So to set the scene, consider for example a group of early hominids with no strict relations other then having part of the same environment.

The essence of the model is the following. Every agent seeks relationships, and only though successful dialogs can these be established. This follows the work of Quillinan (2006), as agents are only allowed to make ties with those they are similar with (often termed homophilia bias or preferential attachment), although Quillinan experiments with other strategies for attachment. Hence, fitness is evaluated in terms of number of friends, or in other words, the node's degree in the social network. The assumption that naturally follows is that having more friends is beneficial for reproduction, and communication success (or gossip) is beneficial for making such relationships. This might seem simplistic, but it is at least a more pragmatic than letting a higher weighted lexicon favor the odds of survival directly. However, every agent can not acquire unlimited relationships. As for establishing new friends, each agent can only reach out to z new agents before concentration on maintaining these friendships. This does not mean that an agent can only acquire a maximum of z friendships — if a lonesome agent is to make contact with a popular agent with z relations, a connection can be made if their dialog is a success. This hopefully forces agents to make introvert decisions when having no relations and extrovert when social groups are established. This is also meant to reflect an agent's capacity to reach out for new friends — meanwhile, it makes the visualization easier. In the network, every friendship is weighted with the number of successful dialogs. If there happens to be more failed dialogs then successful ones, the friendship falls apart and the edge is removed. This is regulated by adding 1.0 to weights on successful dialogs and subtract 0.5 on failed one. It seems reasonable to give an established relationship a "second change" if there suddenly is a miscommunication. The only other way a friendship can decay, is if agents die.

Regarding the chosen traits mentioned in Section 4.1, the genetic makeup consists of four genes (a Grey-encoded bit vector) decoded into values between 0 and 1.

- **Extraversion**. If a hearer has between 1 and z relations, the extraversion variable would be the probability of searching only ply further (among friend' friends) in the graph before initiating a dialog. The maximum ply an agent can search is 3. A low value will then increase the probability of the agent only choosing among its closest friends.

- **Teach child**. This value reflects the probability that a parent is the first one to talk to the child. If both parents have high values, a child gets its first dialogs from them — and if the values are low, the child would have to make its first dialog with a random agent.

- **Lexicon limit**. The limit is calculated from $limit = geneValue \cdot 20$. If an agent gets more than $limit$ words in its lexicon, the lowest weighted word gets pruned.

- **Speech ability**. The value is used as a probability. For low values, an agent will just invent a new word despite of having chosen one in its lexicon — and for high values the intended word is mostly uttered.

In generation 1, a population with a fixed number p of agents is developed from random generated genotypes (random bit vectors). They interact, for every agent in the population, d dialogs are distributed. Thus, for every generation there would be a total of $p * d$ dialogs taking place. For every dialog, the speaker is first chosen with a probability proportional to the agent's number of relations (degree). If a speaker has no relations, a hearer is picked out at random. If the agent has less then z relations, it searches for hearer accordingly to its extraversion strategy, and if it has z or more relations, he chooses among one of its nearest to reflect a sort of maintenance of neighborhood. When all dialogs are completed, a fitness score is calculated for every agent by

$$fitness = [exp(0.02 * numberRelations) - 1] * exp(-0.05 * age) \tag{4.1}$$

This gives a stable climb from 0 to 1 for $0 < numberRelations < 30$ and by the time an agent is around 15 generations old, it gets a 0.5 fitness, given it has a very high degree (see Figure 4.1).

When fitness is calculated, the poorest fitted are chosen to die by a simple tournament selection. k random agents are selected where the n ($n <= k$) best fitted are transfered into the next generation. When $1 - n$ of the population is gone, the $1 - n$ best fitted agents are selected for breeding. Each adult chooses a partner from its closest neighborhood, or a random one if $z = 0$, to account for associative mating (so mating do not follow a random pattern). Crossover is done between the adults' genotypes, and the random selected genotype (as crossover gives two) can mutate with the probability $m = 0.01$ (one bit gets flipped). The newborns are allowed to join the next generation, and the simulation goes on like this for i generations.

In summary, simple biological traits are captured in the genotypes and variation here is offered by the different genetic operators. Culture can be seen as the diversity of words in the population, and the evolving topology. On an individual level, agents wish to establish relations with other agents by following their phenotypic constraints.

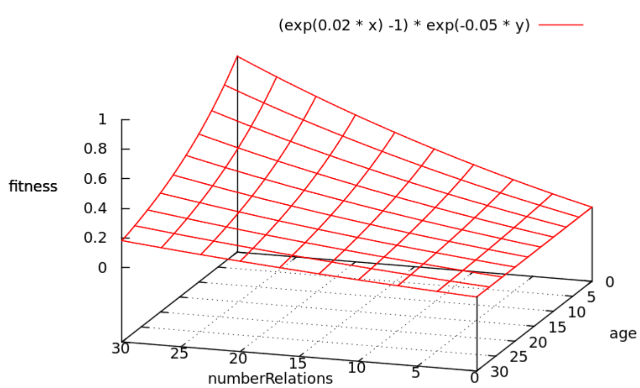

Figure 4.1: A representation of how the fitness function (Equation 4.1) looks like. An ideal agent has a low age and has many relations.

Chapter 5

Results

In this chapter, the results from working using a genetic algorithm is presented. First, a brief illustration and elaboration of how Lipowska's model could be performed on a genetic algorithm — and then some experiment work will be presented, using a genetic algorithm to co-evolve language and social structure.

5.1 The Lipowska Model using a Genetic Algorithm

Here are the results from the model descried in Section 4.2. Simulations are done with learnability being a solely gene in a simple genotype. The simulation goes over $i = 200$ generations, and with $d = 5$ dialogs per agent per generations. Population size is $p = 1225$ $(35 \cdot 35)$, $m = 0.1$, both the initial weights on learnability and invented words are 0.5 as in the experiment from Chapter 3. Tournament selection uses $k = 0.8$ and $n = 0.5$, meaning 40% of the poorest fitted agents dies after each generation. That implies that 4500 dialogs are held and 360 agents die thereafter, in every generation.

Figure 5.2 illustrates a very similar progress as earlier,

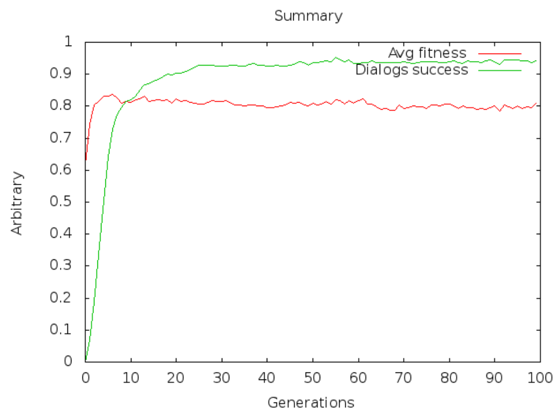

Figure 5.1: Average fitness and the number of successful dialogs over a total of dialogs.

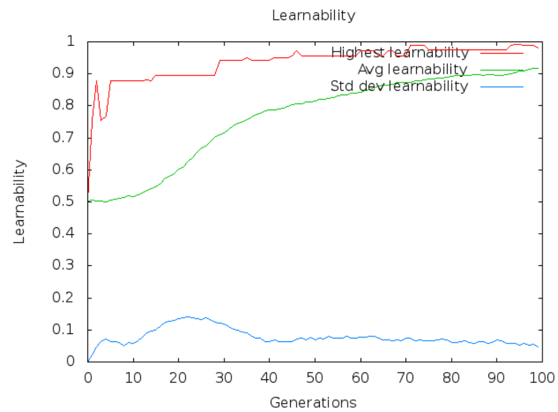

Figure 5.2: Learnability over all agents through the simulation.

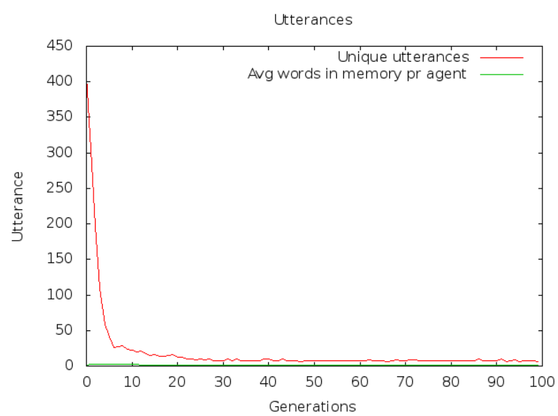

Figure 5.3: Average fitness and the number of successful dialogs over a total of dialogs.

By naively using Lipowska's p_{surv} (Equation 3.1), which can be characterized as a subjective function (and not an objective function), the average fitness score holds little information about what goes on in regard to the agents' weights (See Figure 5.1), only information about the agents' age is extractable. There are of course a small deviation in the fitness among the agents, and better fitted agents are more often chosen as survivors.

There is another issue that can be accounted for in this simulation. Going from above 450 utterances (initially, one word is invented for every unique speaker-hearer pair, and some words vanish as poor fitted agents die — so at the end of the first generation there is not a total of 1225 words) to about 25 happens extremely fast. In only 5 generations, the dead agents drag along with most of the words, leaving “lucky ones” to keep spreading their highest weighted words. Some of these lucky ones are probability newborn with a favorable mutation, eager to promote the parent’s word.

As for learnability (Figure 5.2), this follows a steady progress as in the replica of Lipowska’s model (Figure 3.3). The reduction from 25 to only a handful of dominant words happens as the derivation of learnability is at its highest (around generation 15 to 30), and as positive learnability spread out in the population the forces towards a single word are reduced. This is a result of a decreased *evolvability*, meaning fewer and fewer offsprings turn up better than their parents. In this particular simulation, the population ends up with two equally sized groups of words (and some tiny groups with no future) — fewer words as in results presented in Chapter 3; presumptively because of an increase in the number of total dialogs.

The little dip in fitness (and the learnability’s standard deviation) around generation 10 is a result of starting all agents with age 0. Almost everyone grow old simultaneously, creating some noise for the fitness function.

5.2 Co-evolving Communication System and Social Structure

Following are the results from seven single-run simulations from the outlined methodology in the previous chapter (Section 4.3). Just to emphasize, every simulation has genotypes evolving the three traits; extraversion, teach child and lexicon limit. Only one of the simulations, Simulation 2, has the additional speech ability gene. The first simulations can be seen as the main results, the following six, Simulation 2–7 illustrates how the model responds to other parameters. Every simulation uses the parameters in Table 5.1 if not stated otherwise.

Parameter	Description
$a = 1000$	Number of agents.
$i = 100$	Number of generations.
$d = 1$	Number of dialog per agent per generation, making a total of 100,000 dialogs.
$z = 5$	As maximum new relations an agent can make.
$k = 0.9$	90% of the whole population are allowed in tournament of generation turnover.
$n = 0.8$	80% of the k best fitted agents win, leaving 20% of the agents for death.
$genesSize = 8$	The number of bits a gene is made up from.
$m = 0.01$	1% chance for a mutation in a newborn.

Table 5.1: The parameters used in simulations.

Simulation 1, Main Results

This is the main results from the model, with parameters all equal to those in Table 5.1. The first, and most important (See Simulation 3), trait facing selective pressure is the one making parents teach their own children a language after birth. This is not very surprising as this is the probability of the cheapest way to make a relationship in the population. Only afterwards do extroversion/introversion behavior emerge, and agents quickly find it favorable to reach far out in the social network for new friends. Now, remember that an agent can permit incoming dialogs even though the agent has reach z relations. As it might be, by (selfish) intuition, preferable to bind with close one first (as these are more likely to share your language/opinion) — biology might find it more favorable for a descendant if an agent passes along a individual unprofitable mutation, as this might increase the descendant's degree (and therefor fitness) having such agents around.
The last phenotypic trait, memory limit, appear very sporadic in most simulations. Seemingly, with these parameters, their is no need for either high, or low, memory ability. In contrast to naming games on static topologies (Wang *et al.*, 2007), a very low memory limit is sufficient when agents are continuously dieing and breeding. As for coherence, it is only once in a while a simulation reaches consensus. Two or three words are often left after 100,000 dialogs. This is very similar to work found in the literature on static scale-free networks. In particular, Baronchelli *et al.* (2006a) constructed a scale-free network with a power law degree distribution and an average degree of 4, populated with 2000 agents and reached consensus right after 10,000 dialogs (although they where not choosing speakers proportional to their degree, but rather letting a random hearer choose its speaker) — this network is also showing a power law degree distribution (See Figure 5.12a) and a average degree of 4.072. (See Figure 5.4)

Count of most used words
536
461
3

Avg degree	Avg shortest path
4.072	7.094

Table 5.2: Simulation 1, the count of most the used words, average degree, and average shortest path in the network at generation 100.

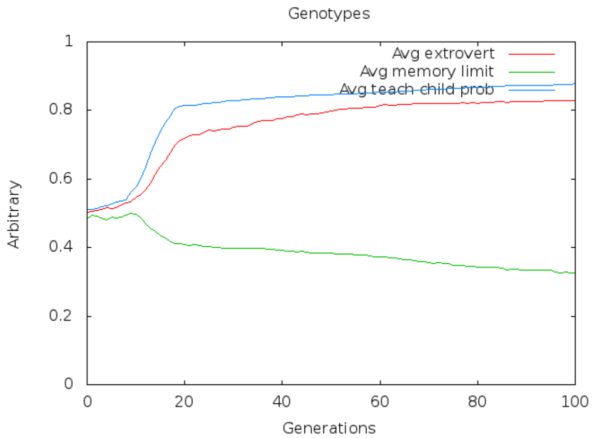

(a) The average of each gene over all the agents.

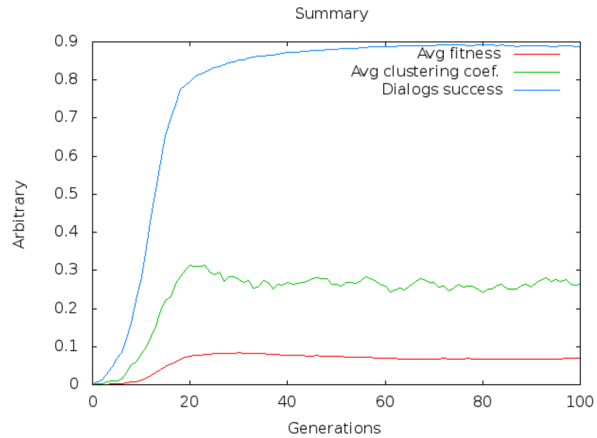

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

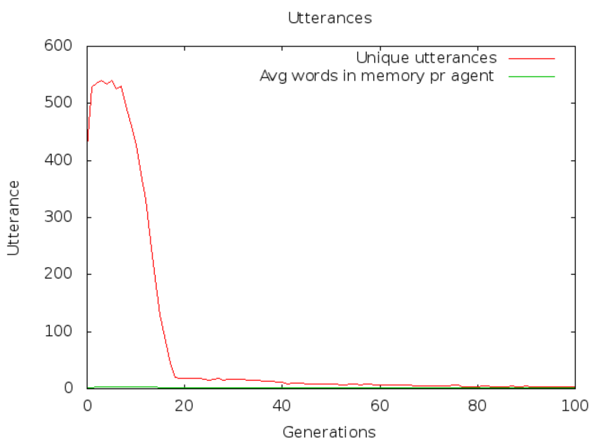

(c) The number of unique highest-weighted words and average inventory size.

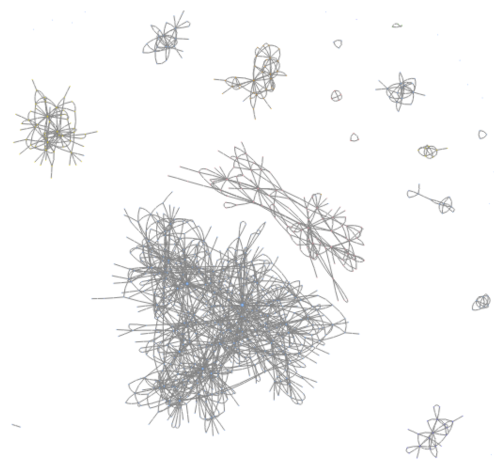

(d) A snapshot of the social network at generation 20.

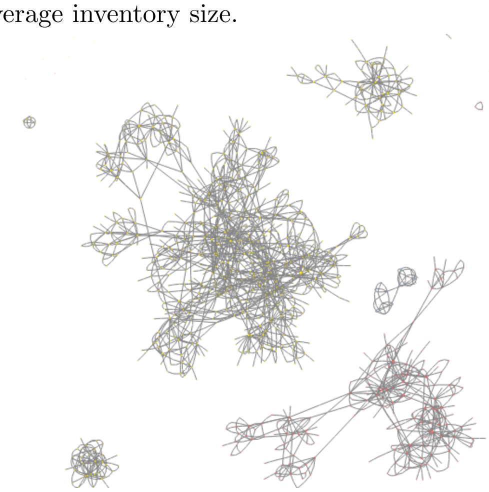

(e) A snapshot of the social network at generation 60.

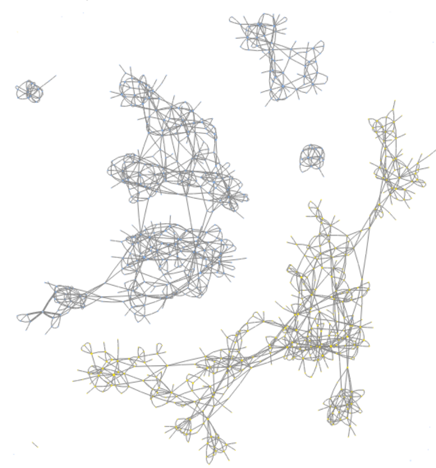

(f) A snapshot of the social network at generation 100.

Figure 5.4: Overview of simulation 1 in terms of communication success, genotypic values, fitness, utterances, and social network.

Simulation 2, Speech Evolving

An additional gene is added to capture the emergence of a dependable speech apparatus.

In this simulation, when no agent have the ability to make themselves understood, the evolutionary progress is just drifting (no selection pressure) for the first 20 generations (See Figure 5.5a). Here, three observations can be commented.

Firstly, the three different traits, speech, teach child probability, and extroversion are all happening right after each other, respectively. This is often known as *domino convergence*, and happens when features are encoded separately and have different contribution to the fitness (Weise *et al.*, 2012). Secondly, there is actually — in contrast to Simulation 1 — a slight pressure towards introversion when the agents start to build a social structure. This can be due to drifting, but the extroversion behavior is at least emerging much later then in Simulation 1. When a few agents get there speech apparatus going and they happens get a couple successful dialogs, these are very likely to stick around, because of higher fitness. They breed, and those how are learning their children their highest-weighted words are even more favorable. By generation 20 (see Figure 5.5d), a little community where everyone share their highest weighted word is settled — and this group grows and dominates the simulation, both biological and cultural. Thus, this leads to relatively safe consensus in the population. Still, there are many disconnecting communities, however, this is not very surprisingly when a small community agree on a language system before all parents adapt to letting their newborn be a part of the community. This converges to one single cluster if the simulation is allowed to run longer, as the extroversion behavior goes up.

Count of most used words	Avg degree	Avg shortest path
975 10 1	3.762	4.470

Table 5.3: Simulation 2, the count of most the used words, average degree, and average shortest path in the network at generation 100.

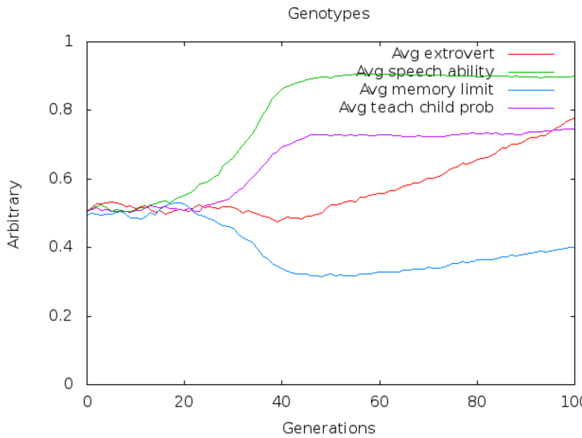

(a) The average of each gene over all the agents.

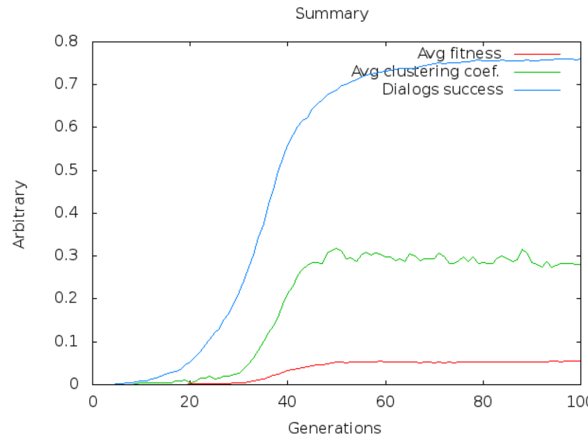

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

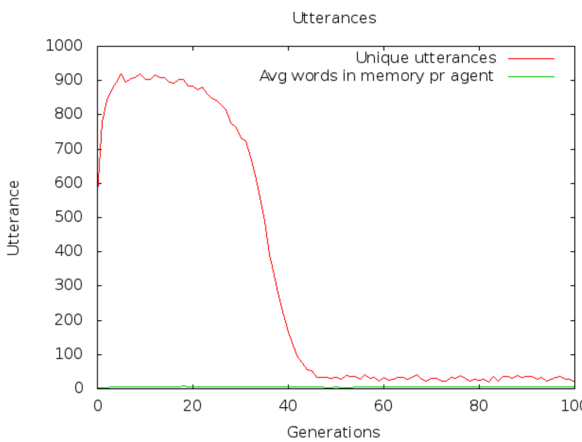

(c) The number of unique highest-weighted words and average inventory size.

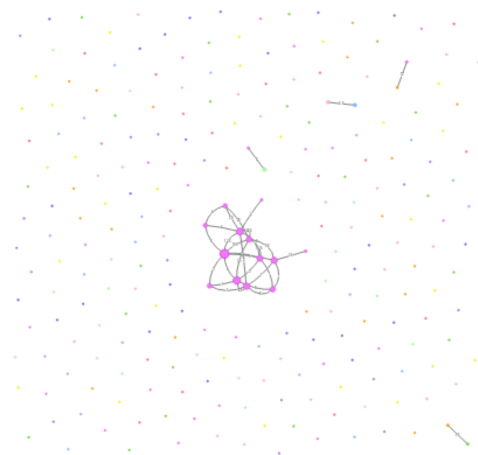

(d) A snapshot of the social network at generation 20.

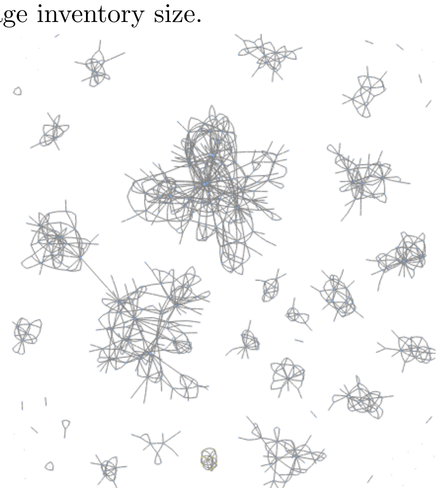

(e) A snapshot of the social network at generation 60.

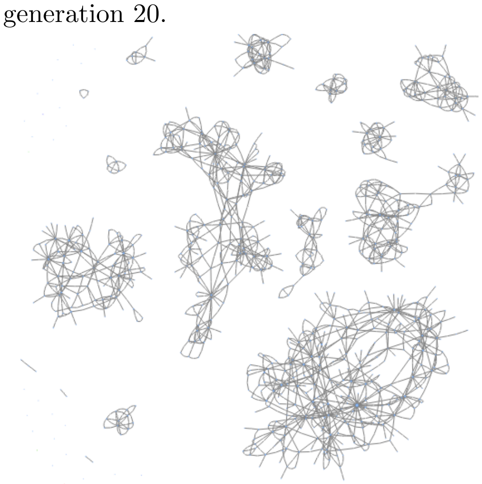

(f) A snapshot of the social network at generation 100.

Figure 5.5: Overview of Simulation 2 in terms of communication success, genotypic values, fitness, utterances, and social network.

Simulation 3, Hard-wiring Low Child Learning Probability

From Simulation 1 and 2, a high probability for a parent to teach its own child tend to evolutionary favorable. In this simulation, this trait is hard-wired to 0.15, meaning that each parent initiates a dialog with its newborn only once in a while after breeding.

This one is rather intuitive, when only 15% of all parents get to have their their first dialogs with their newborns, there are very few connections being made. Although, when some agents do occasionally make a successful dialog, they are more likely than others to breed (as they gain higher fitness), and if newborns occasionally learn from only one of their parents, the word entropy starts to decrease (See Figure 5.6).

Count of most used words
135
122
107
36

Avg degree	Avg shortest path
0.440	1.172

Table 5.4: Simulation 3, the count of most the used words, average degree, and average shortest path in the network at generation 100.

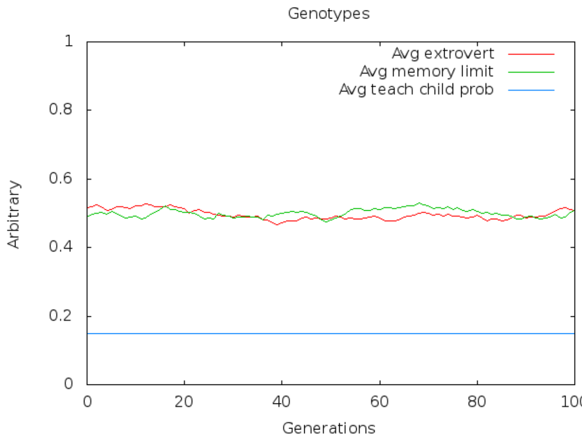

(a) The average of each gene over all the agents.

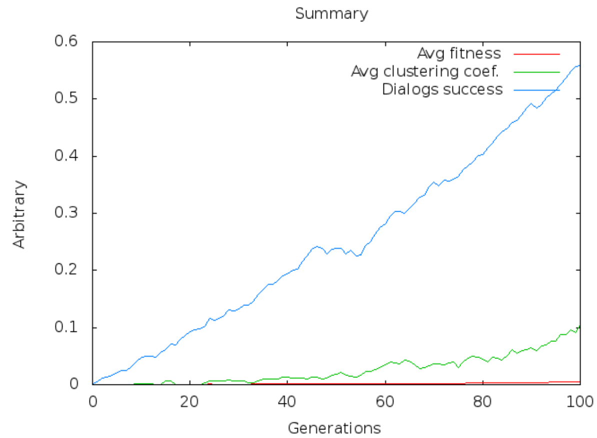

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

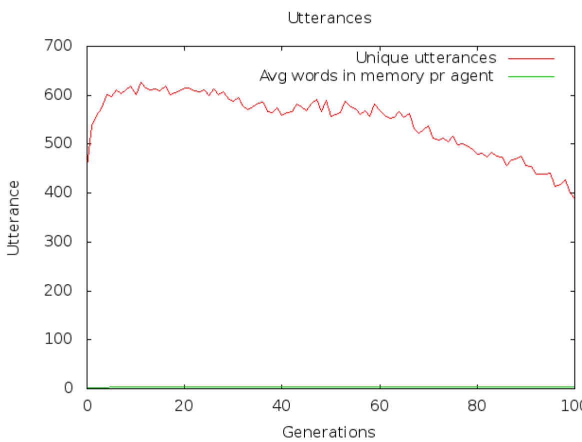

(c) The number of unique highest-weighted words and average inventory size.

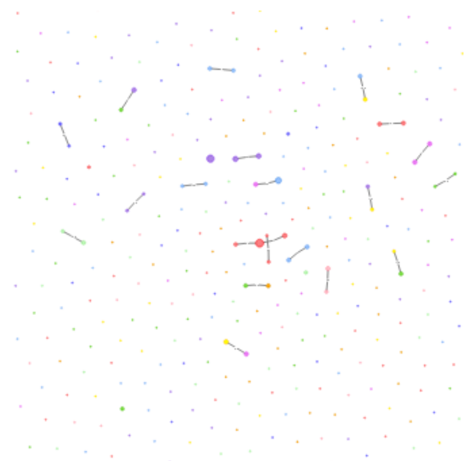

(d) A snapshot of the social network at generation 20.

(e) A snapshot of the social network at generation 60.

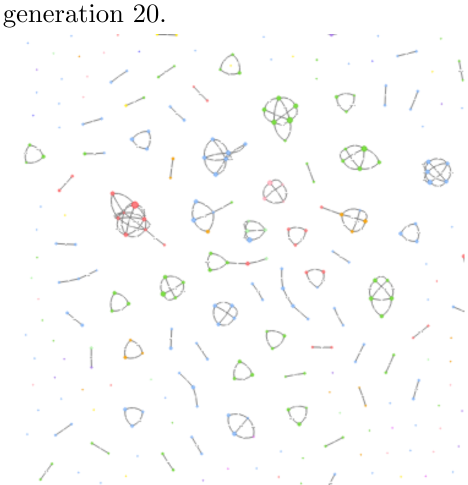

(f) A snapshot of the social network at generation 100.

Figure 5.6: Overview of Simulation 3 in terms of communication success, genotypic values, fitness, utterances, and social network.

Simulation 4, Hard-wiring Low Extroversion Probability

As seen in Simulation 1–2, the extroversion probability tend to be evolutionary favorable. In this simulation where the value is hard-wired to 0.15, meaning that the agent most often only initial new friendships with friend's friends and not friends beyond that.

This share resemblance with Simulation 1 in terms of consensus (see Figure 5.7c) and degree distribution (see Figure 5.12d) — but the social network is very disconnected as there are no longer far reaching connections holding the graph together.

Count of most used words
552
413
21

Avg degree	Avg shortest path
2.236	2.653

Table 5.5: Simulation 4, the count of most the used words, average degree, and average shortest path in the network at generation 100.

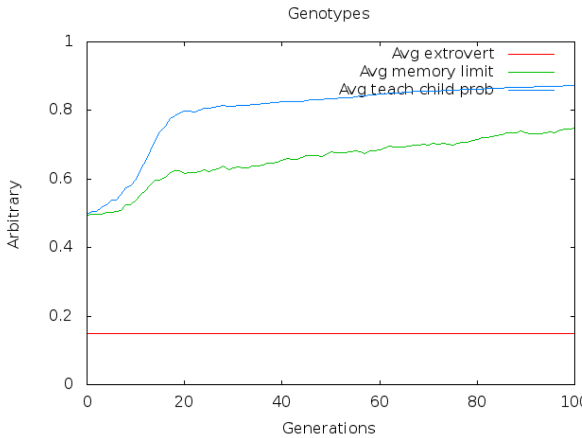

(a) The average of each gene over all the agents.

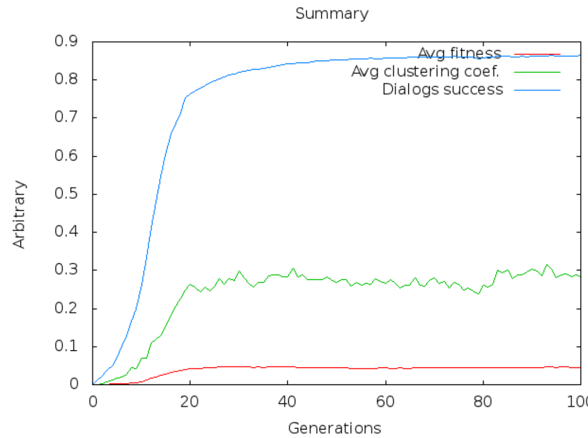

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

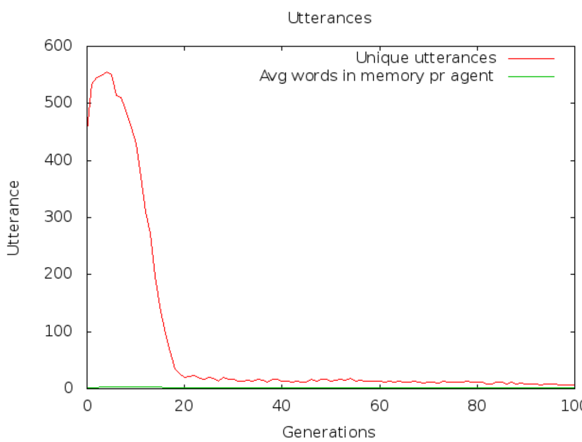

(c) The number of unique highest-weighted words and average inventory size.

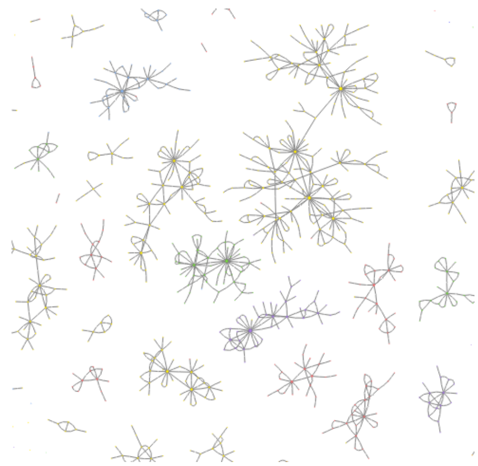

(d) A snapshot of the social network at generation 20.

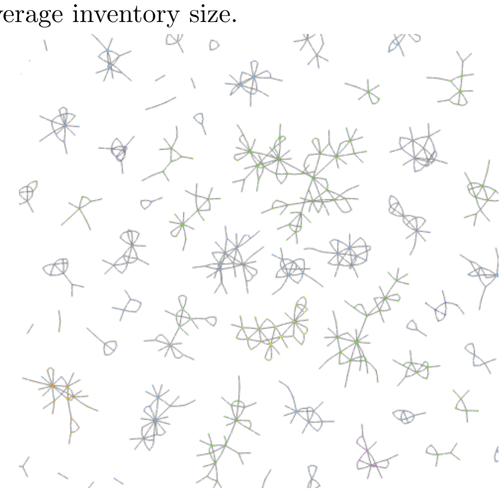

(e) A snapshot of the social network at generation 60.

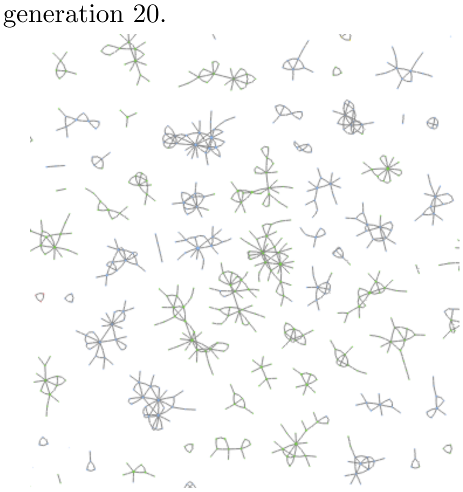

(f) A snapshot of the social network at generation 100.

Figure 5.7: Overview of Simulation 4 in terms of communication success, genotypic values, fitness, utterances, and social network.

Simulation 5, Increasing Number of Dialogs

In this simulation the ratio between culture interaction and biological turnover is adjusted by going from $d = 1$ to $d = 5$. To account for the total number of dialogs, there are only 25 generations in this simulation.

This is somewhat equal to adjusting Lipowska's p in earlier experiments (see Section 3.1) - as this adjusts the ratio between communication and "population update". Changing the number of dialogs in each generations allows for more connection to be made. This highly condenses the communities; the average degree goes up to 5.570 and average shortest path goes down to 3.461. It puts less pressure on extroversion behavior as a tight community do not fall apart that easy — which leads to many invulnerable communities, each holding its most dominant word. This is also the simulation which diverts the most from the power law degree distribution (see Figure 5.12e).

Count of most used words
326
196
155

Avg degree	Avg shortest path
5.570	3.461

Table 5.6: Simulation 5, the count of most the used words, average degree, and average shortest path in the network at generation 25.

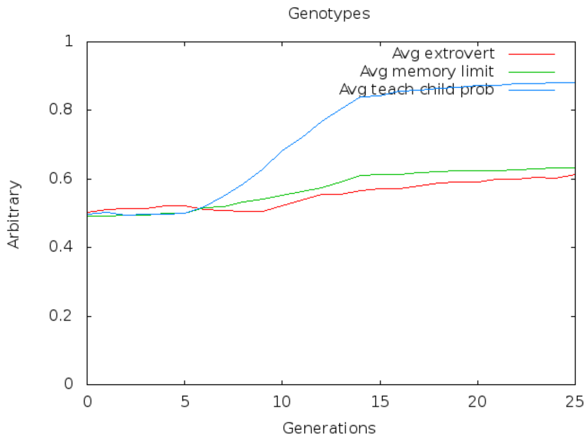

(a) The average of each gene over all the agents.

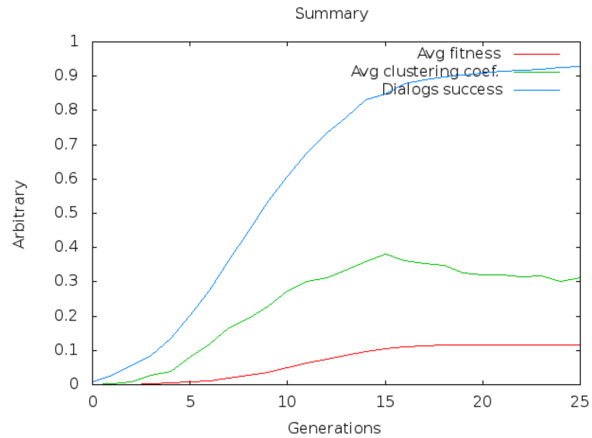

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

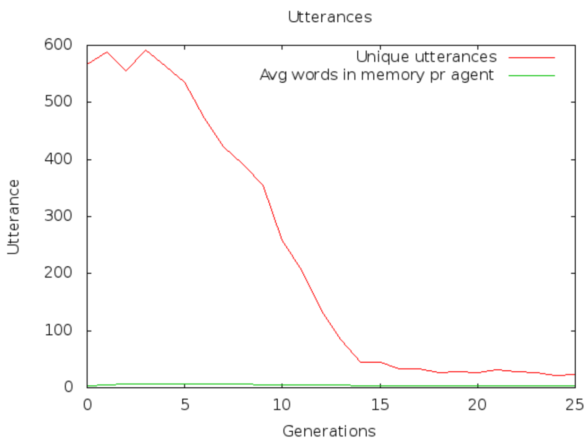

(c) The number of unique highest-weighted words and average inventory size.

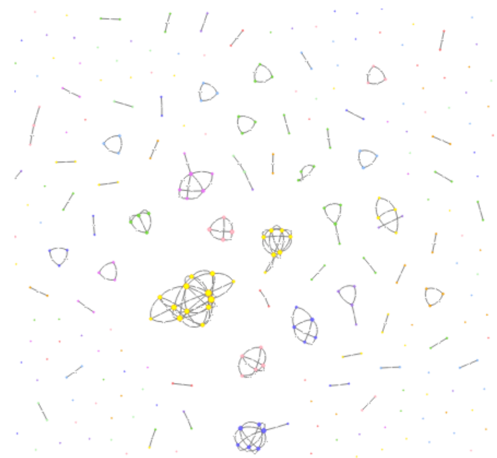

(d) A snapshot of the social network at generation 5.

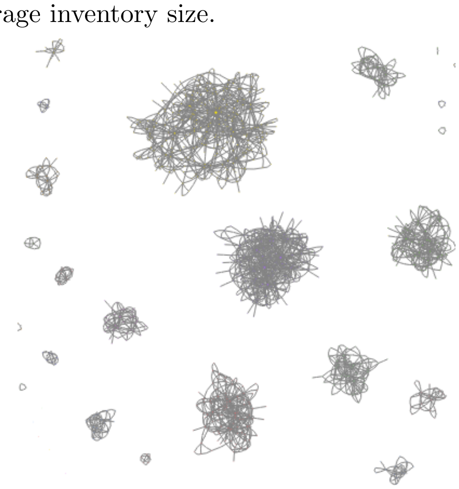

(e) A snapshot of the social network at generation 15.

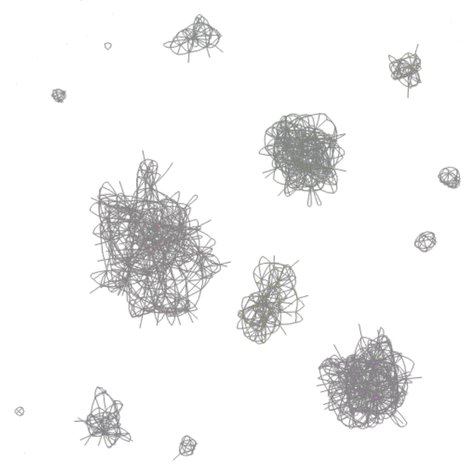

(f) A snapshot of the social network at generation 25.

Figure 5.8: Overview of Simulation 5 in terms of communication success, genotypic values, fitness, utterances, and social network.

Simulation 6. Scaling, 100 Agents

In this simulation the model is tested on a much smaller set of agents, only 100 agents.

Despite for the fact that the average shortest path goes down, this looks very much as Simulation 1. The Figures 5.9d to 5.9f shows (with arbitrary colors matching words) how very clearly that different agents are following different strategies, and hence making it both non-regular and yet not random.

Count of most used words
76
24

Avg degree	Avg shortest path
4.220	3.099

Table 5.7: Simulation 6, the count of most the used words, average degree, and average shortest path in the network at generation 100.

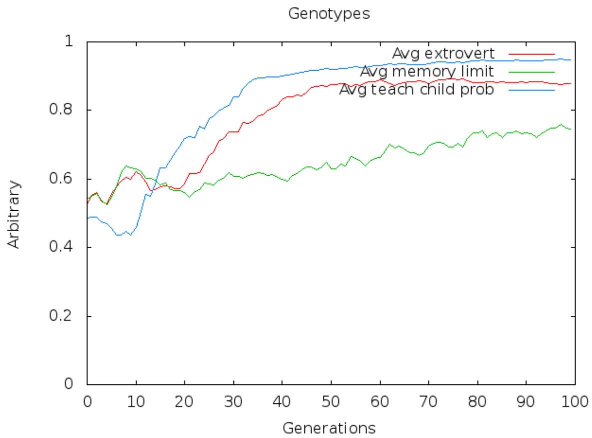

(a) The average of each gene over all the agents.

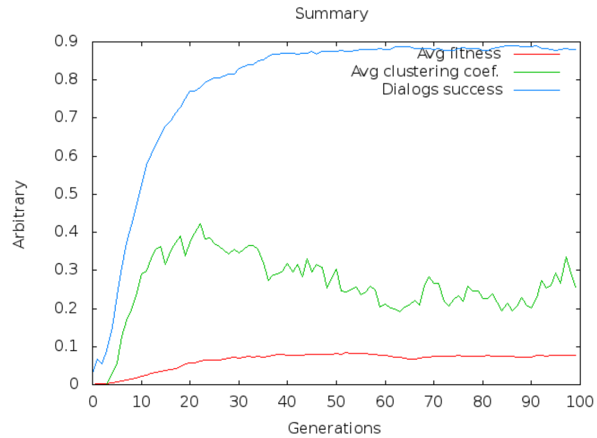

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

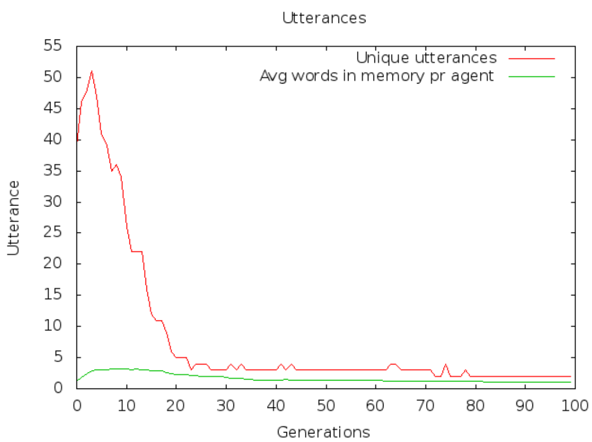

(c) The number of unique highest-weighted words and average inventory size.

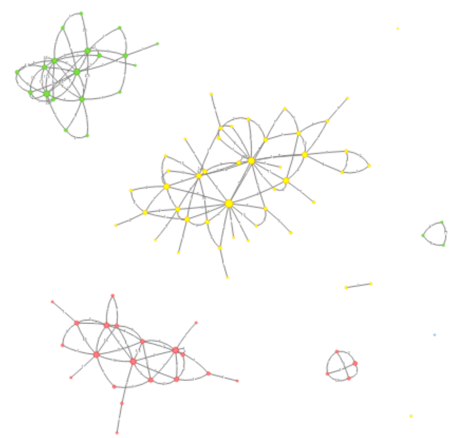

(d) A snapshot of the social network at generation 20.

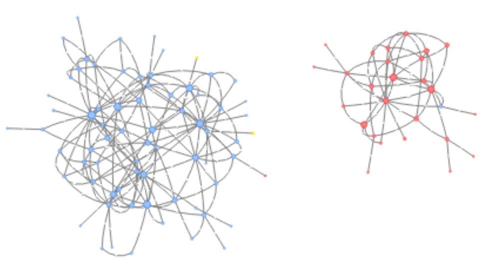

(e) A snapshot of the social network at generation 60.

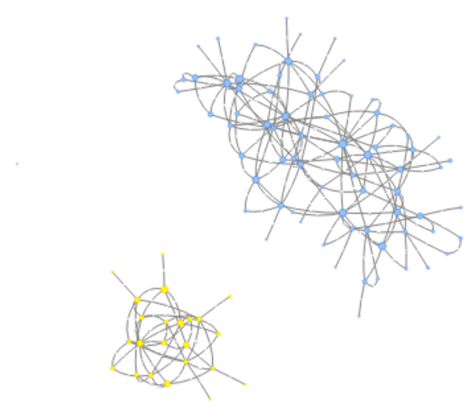

(f) A snapshot of the social network at generation 100.

Figure 5.9: Overview of Simulation 6 in terms of communication success, genotypic values, fitness, utterances, and social network.

Simulation 7. Scaling, 10,000 Agents

In this simulation, the number of agents are increased to 10,000 agents. The simulation is limited to 50 generations for computational reasons.

The model seems to scale up quite nicely and shows resemblance with Simulation 1 in almost every aspect. It takes about 5 generations longer to reach the stable “few words” phase, but thats not surprisingly as there are 10 times more words in the population (see Figure 5.10c).

Count of most used words
8365
1201
141

Avg degree	Avg shortest path
4.167	7.480

Table 5.8: Simulation 7, the count of most the used words, average degree, and average shortest path in the network at generation 50.

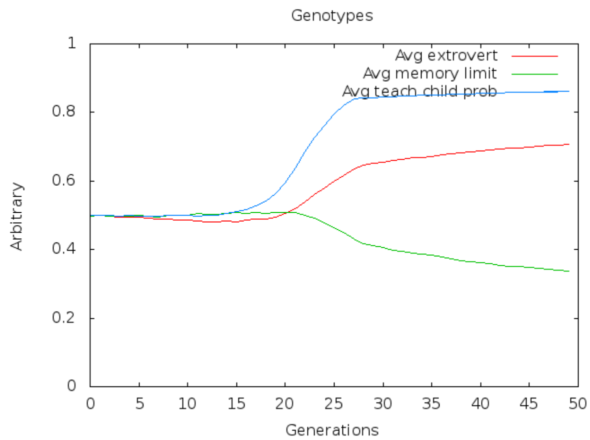

(a) The average of each gene over all the agents.

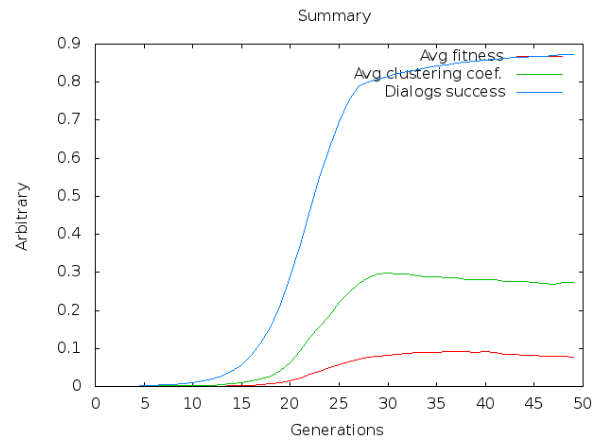

(b) Average of fitness, the ratio between successful and failed dialogs, and the clustering coefficient in the population.

(c) The number of unique highest-weighted words and average inventory size.

Figure 5.10: Overview of Simulation 7 in terms of communication success, genotypic values, fitness, utterances, and social network.

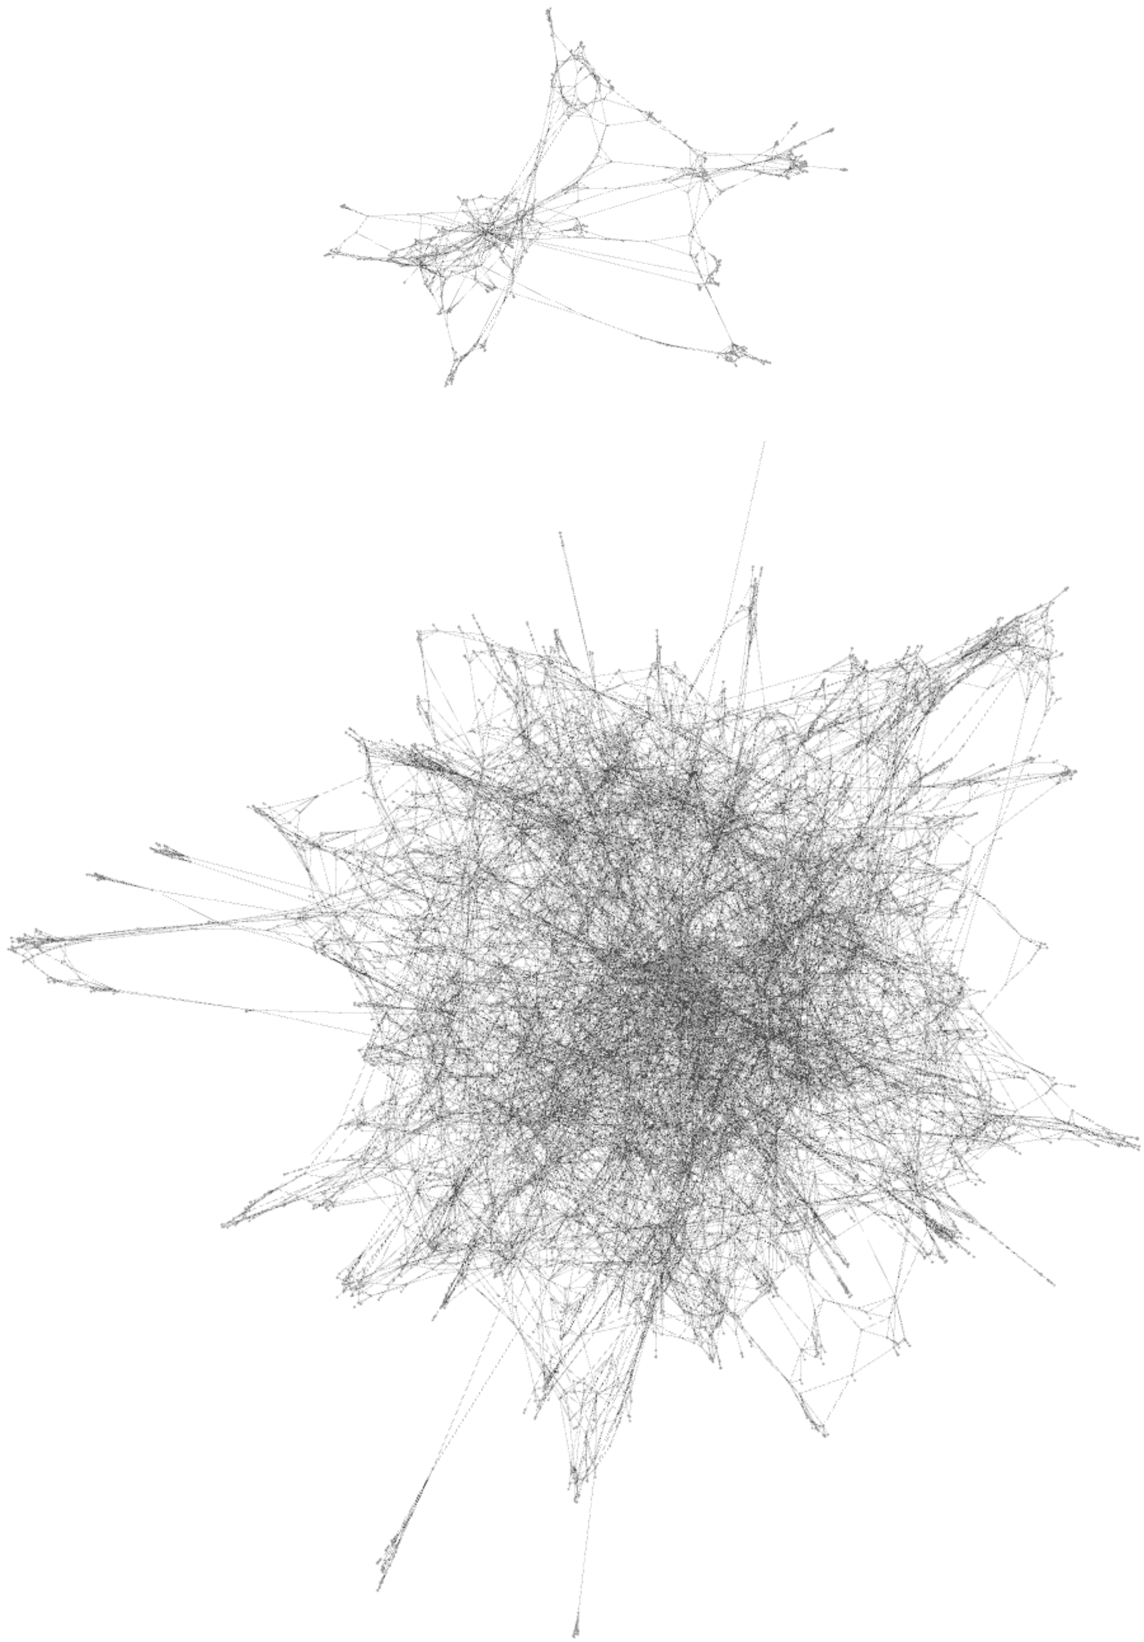

Figure 5.11: The social network at generation 50 with 10,000 agents, one dominant word and two smaller groups. The last one, with 141 agents, is out of scope.

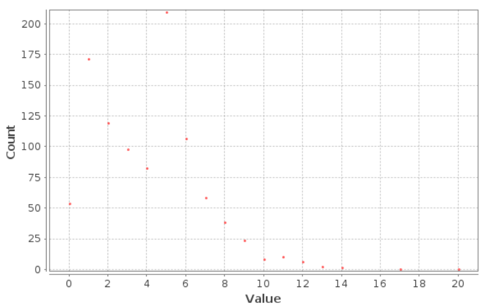

(a) Simulation 1 at generation 100.

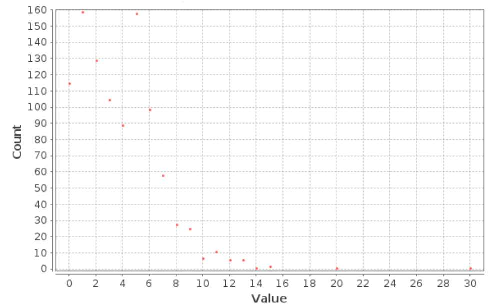

(b) Simulation 2 at generation 100.

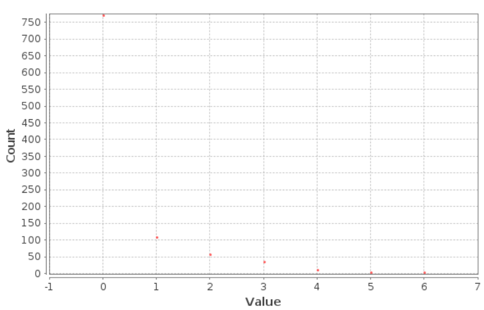

(c) Simulation 3 at generation 100.

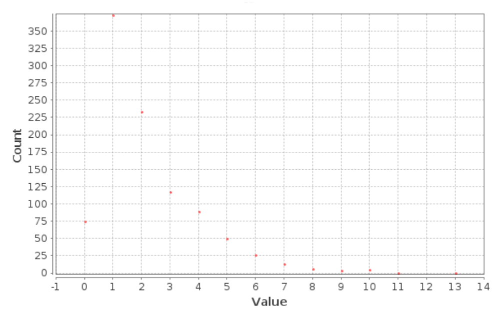

(d) Simulation 4 at generation 100.

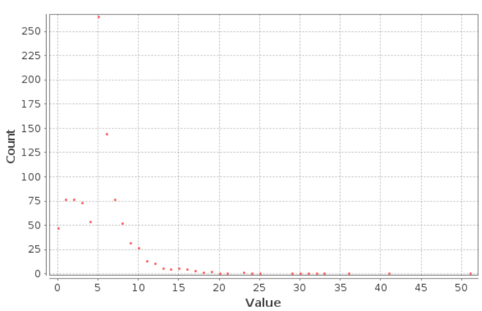

(e) Simulation 5 at generation 25.

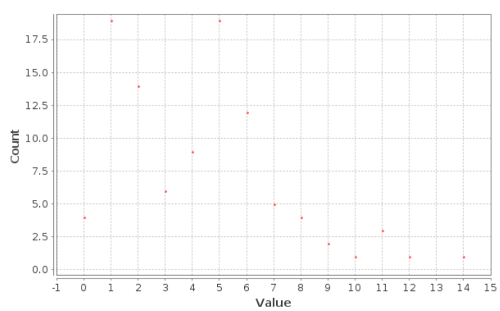

(f) Simulation 6 at generation 100.

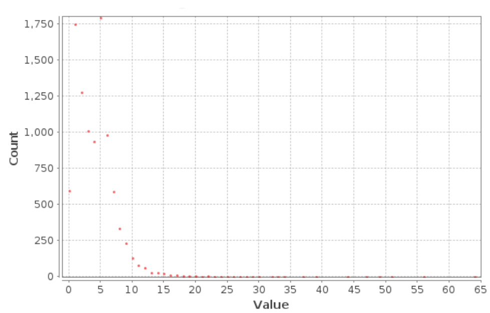

(g) Simulation 7 at generation 50.

Figure 5.12: Degree distributions from Simulation 1-7.

Not mentioned are the clustering coefficients. They all tend to stabilize to around 0.3, even when the extrovert behavior is hard-wired to a low value. Precisely the cause is difficult to say without having more data, but it is seems safe to assume that there are design flaws involved, as there should have been more fluctuation while skewing different parameters.

A few words on the degree distributions (see Figure 5.12a to 5.12g). They all (with the exception of Simulation 5) tend to display similarities with scale-free networks, having few nodes with very high degree and a very dense amount of few-degree nodes. Although, there are two irregularities with these distributions, most interestedly in Simulation 1 and Simulation 7, as they represent the main results. Firstly, the number of nodes having zero connections, hovering about 5%. This is due to remaining chance of child abandonment (low teach child probability), and as friendships can decay through death and failed dialogs. Another deflecting number is the degree of 5. This is naturally because of the constrain z in the model. Agents (living in the last generation) quickly fill up their quota after birth and are then left for distant speakers to initiate dialogs. Having a fixed constrain z over all agents might be a bit rough — handing out z variables following a normal distribution is perhaps a better way to go. In other words, this 5-degree irregularity would probability ease out if every agent had a more distribution mental capacity for establishing new friends.

In summary

- Conducting a naming game on a co-evolving social network (from nothing but nodes and made-up words), shows multiple similarities to naming games in the literature where the topology is pre-constructed.

- The population goes from high to low word entropy very quickly, even in Simulation 7 with 10,000 agents, as soon as small clusters are forming.

- There is no clear pressure towards having neither a high nor low memory constraint.

- The population rarely reaches consensus, except when there is a need for pre-adaptations, as in Simulation 2.

- The count of different words, and the clusters' density, at the end of the simulation is highly influenced by the number of dialogs per generation as demonstrated in Simulation 5.

Chapter 6

Discussion and Conclusions

Considering the literature and the results from the two experiments from Chapter 5 — a discussion is reasonable before drawing conclusions. The discussion has three subsections, one for each model in Chapter 5, and one where they are seen together. Conclusions are highlighted before giving suggestions for further work.

6.1 Discussion

The Lipowska Model using a Genetic Algorithm

First, some thoughts on using a genetic algorithm on the experiment from Chapter 3, demonstrating the Baldwin effect. Using a genetic algorithm when the genotype only consists of one gene might be an over-engineering — but it provides a modular and flexible implementation environment (with genetic operators, different strategies for survival and adult selection, easy to add new genes). While this is good — it also introduces another problem. Earlier, the parameter p was the only factor determining the ratio between biological turnover and culture interaction — now, this is a dispute between number of dialogs per generation d and the number of survivors w. The search for good simulation parameters was not very problematic in this case, but it was at least more work then skewing a single p parameter. Although there is a lot of conditional choices happening behind Lipowska's (2011) population update, this emphasizes the importance of her work in studying the relationship between cultural and biological transmission.

Further, the p_{surv} (Equation 3.1) was reused as a fitness function. This makes the fitness function a subjective function, in contrast to an objective function which is ordinarily used in evolutionary algorithms. In this simulation, it means that every agent is valued after how every one else are doing. While this is not necessarily a wrong approach, it is a bit diverting in terms of self-emergence. For instance, if this model is to be tested on a vast number of agents, it would not be fair to let agents be judged after how someone else is doing in a distant community. At the very least, if a subjective function is to be used in a model like this, a agent should be compared against other agents in its immediate surroundings. Of course, this might introduce noise in a model which only serves as an example, but it is surely more realistic. To clarify, an objective function would in comparison be something like "how well is an agent doing in terms of optimal performance" — but then you would need to know what the best agent, or solution, looks like. An objective function in this presented model would not be a problem, but it could clearly make a difference in an emergent human culture — with accumulating wisdom and health care — where there are no comprehensive solution for individual performance.

Co-evolving Communication Systems and Social Structure

Following a sociobiological track, the presented genetic algorithm demonstrate how researchers in the field of language evolution can investigate genetic constrained mechanisms with highly cultural influence — both co-evolving with social structure, or other phenomena assumed to be influenced by, or from, language change. However, this model lacks several important factors before conclusions on language evolving could be made. For instance, it is not necessarily true that a naming game is a suitable abstraction for the semiotic dynamics, the process of mapping signals and signs to concepts or meanings. At least not when it is co-evolving with a social structure (discussed later). When making this abstraction, there is a loss in adaption on the individual level, as agents only participates in social learning and not individual learning. The learning mechanisms used in this experiment is also arguably naive, as agents only choose words randomly from their inventory proportional to success weights. Furthermore, there are no distinct difference between who is the speaker and who is the hearer, or in other words, who learns from whom — except for a child's first dialog. As in the work from both Hashimoto *et al.* (2010) and Munroe & Cangelosi (2002), agents are assigned roles as teachers and students depending on their age which allows for a deeper investigation of between social and individual learning. However, there is the obvious problem in computer simulations when dealing with animal and human learning; what should the associated costs be? This is clearly a challenge in the presented model as well. For instance, weights on the social network are added after successful dialogs with the weight 1.0 and subtracted by 0.5. In the present model, skewing this weights, or making them equal, does not change the overall behavior much, but it matters. This is also true for the overall parameters given to the simulations (Table 5.1), regarding the speed of cultural and biological turnover. In both cased, the numbers have been set (although inspired by the literature) by what sounds and gives intuitive results — this emphasizes the importance of involvement from other disciplines. In the same way, the costs of having a extrovert behavior in this simulation is no more than risking a potential successful dialog. While this might be enough to damper extrovert behavior at the start of each simulation, the real risk in making contact with total strangers are naturally much higher.

The social strategies in the presented model follow a probabilistic nature, meaning no agents have deterministic behavior solely based on their genetic makeup. While this might make a more lifelike scene, it makes it difficult to determine what diverting behavior looks like and how it influences the allover system, since agents are per definition diverting from their strategies every now and then by chance. It is also debatable if the initial setup of homogeneous agents is realistic — as social structure might preexist in the animal kingdom prior to language. There are, for instance, evidence of social communities among animals, depending on sex, age, and territories (Wolf *et al.*, 2007). How strong these relations are in contrast to relations establish on the basis of language and culture is still an opened question. Hence, the model presented should perhaps start of with pre-seeded social networks, reflecting a typical distribution of age and demography, instead of going from a homogeneous group of "free-riders". Although the topology in the final results from Simulation 1 (Section 5.2) display high resemblance with social networks in modern human history (power law degree distribution, low average shortest path) (Baronchelli *et al.*, 2013), it is questionable if these are the main characteristics in a social network right after the emergence of a language within a species.
This could point to biases in the implementation and in the assumptions made in the model — a known problem when modeling social phenomena (Helbing & Balietti, 2011).

As for some other observations. Simulation 2 (Section 5.2), which also includes the emergence of a coherent speech apparatus, highlights an important point in language evolving, namely pre-adaptations. Language, as we see it in humans, would indeed require some adaptations that might not necessarily be directly favorable for reproduction itself, such as certain cognitive abilities for mapping holistic sounds or gestures to concepts in nature, or for instance strategies for joint attention (pointing, eye-gazing) (Gong & Shuai, 2012; Christiansen & Kirby, 2003a). In this simulation, there is only one simple pre-adaptation, working as a "needle-in-a-haystack" in the fitness landscape, but in reality there could be a lot more. This is one of the most used arguments from nativists against Darwinian explanations to language.

Further, as viewed through all the simulations, there does not seem to be any pressure facing the agents lexicon size. On one-generation static topologies, it has been shown that the time to reach consensus on naming games is highly influenced by the lexicon size, and intuitively, bigger lexica is better (Wang *et al.*, 2007). But in this model it makes no difference, agents are simply dying before reaching their cognitive limits, or are born, raised and retired in the same one-word stable community.

Overall

Conducting a naming game on a co-evolving social structure where fitness is based on the number of friendships has some implications. When only a few agents get successful dialogs and small cluster are emerging, these agents are naturally highly favored. It only takes about 20 generations before all agents without relationships are replaced by newborns from the small emergent clusters. This is also the only period in the simulation where there are any clearly selection pressures. After this rapid transition, the majority of agents are mainly satisfied with their adaptations. Only extrovert behavior can be biological preferable - as this increases the low-degree agents' chances to be selected as hearer by others. Even in Simulation 2 (Section 5.2), not having a coherent speech apparatus makes lot of noise between the agents, the transition happens relatively fast. Thus, the naming game might be "too easy" and happening too fast compared to the evolution of social structure — which is unfortunate for the overall behavior, as it is assumed that language is the functional cause for social structure in this experiment. Although in general, complex linguistic systems needs many generations to forge — assuming there exists a language capability, a single generation could be able to construct simple human languages through few, or even a single, generation (Kirby, 2002b).

This leads us to an important point, namely the fitness function and if it includes the right ingredients. In this simulation, language is naively assumed to be favorable for making social ties. Perhaps this causality is too weak, and there should be even more in between, for instance by using language as a coordination tool in resource gathering (Munroe & Cangelosi, 2002; Vogt, 2009). Language is by not doubt a complex affair, and when choosing good parents for future generations, other aspects should also be considered — for instance, cognitive effort, learning costs, teaching costs. While it is temping, we should at least not calculate fitness solely based on communication success, as it would relate to "the classic case of looking for your car-keys where the street-lamps are" (Bickerton, 2007, p. 522). Also important is, what — or whom —- shall agents be compared to? In this thesis, both a subjective function normalized over the population, and a objective function (number of relationships) have been used. A better approach might lay somewhere in between — for instance, the number of social ties, weighted after the friends' knowledge or popularity? This could provoke conflicting behavior within the population, making the agents oscillate between different optimal strategies.

A strength with a genetic algorithm is that we can change the whole experiment's motivation and outcome with only a few lines of code, but that also emphasizes how careful the function must be chosen initially. Allover, with support from other disciplines, this work outline great potential in combining different theories and models with the use of an open minded and explorational framework.

6.2 Summary and Conclusions

When exploring, validating or and exemplifying different theories in language evolution with computation models it is necessary to make assumptions where knowledge is limited and where computers lack representation. Although there is a lot of these assumptions, researchers tend to lay little weight on the importance of one particular consensus; that language evolution is influenced by both biology, culture, and different types of learning — all simultaneously. While there are many simulation which obvious do not depend on biological variation, a genetic algorithm as presented here, can still give a reliable set of tools for finding good cultural parents in overlapping generations. In this thesis, several computational models have been brought up, hopefully to give some insight to the potential in self-emergent agent-based models in language evolution. One model from recent literature is examined in detail, a model demonstrating the advantages assimilating an abstract learning apparatus and applying it in a naming game, over both biological, cultural and individual timescales. A slightly different and more robust approach is proposed, making a foundation for applying other theories regarding language evolution over multiple timescales. Using the same naming game, a co-evolution between language and social relations is demonstrated, where language works as a dependency for assembling strong social networks. A naming game, which work as an abstraction to semiotic dynamics, is perhaps a far too unrealistic model in terms of evolving a early protolanguage when co-evolving social ties. However, both the methodology and the experimental work reveal a great potential for future work.

6.3 Recommendations for Further Work

Referring to a part of this thesis' objectives, this model is highly open for extensions and explorations in different directions. Earlier work which has not considered either dynamic social networks or biological turnover, could be combined. The first and most interesting aspect would be the emergence of compositionality and word order regularities with rule-based systems (Gong, 2011; Vogt, 2005), or even neural networks (Munroe & Cangelosi, 2002). Adapting neural networks could not only incorporate language, but also the social dynamics and the perceptions in the environment. An additional emergence of compositionality would probably increase the complexity so that the turnover of language and social structure would become more equal, in terms of evolution.

A less dramatic change could be to add more arbitrary objects to the naming game for agents to agree on. Or, the guessing game could be adapted, where all agents get to perceive some object or objects, having, for instance, different color and size features. This is one of many other language games often used when embodied agents are socially cooperating (Steels, 2009).

This model may also be a foundation for studying the diversification of language and language families with variation on both biological and cultural level (Levinson & Gray, 2012). In particular, it might facilitate in the questions whether phylogenetic trees are adequately to capture the diversification of language, or if a network is more suitable, as different language families might interact over communities. Hence, demography and geography could be introduced to make the model more realistic (Patriarca & Heinsalu, 2009).

Of course, experimenting with other types of social strategies could be done. For example, adapting a more "Prisoner's dilemma" type of interaction would be interesting. If initiating a dialog with a total stranger has additional cost, there might be agents deflecting from this strategy in order to only exploit other's extrovert behavior. If everyone does this, no social network would assemble, creating a social dilemma (Fehl *et al.*, 2011).

Another interesting effort could be to add more distinct social roles, assigning the agents as teachers and learners, depending on their age, to further investigate the learning dynamics in an evolving graph.

Bibliography

Albert, Réka, & Barabási, Albert-László. 2002. Statistical mechanics of complex networks. *Reviews of modern physics*, **74**(1), 47.

Baldwin, James Mark. 1896. A new factor in evolution. *The American Naturalist*, **30**(354), 441–451.

Barabási, Albert-László. 2009. Scale-free networks: a decade and beyond. *Science*, **325**(5939), 412–413.

Baronchelli, Andrea, Loreto, Vittorio, Dall'Asta, Luca, & Barrat, Alain. 2006a. Bootstrapping communication in language games: Strategy, topology and all that. *Pages 11–18 of: Proceedings of the 6th International Conference on the Evolution of Language*, vol. 2006. World Scientific Press.

Baronchelli, Andrea, Felici, Maddalena, Loreto, Vittorio, Caglioti, Emanuele, & Steels, Luc. 2006b. Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics: Theory and Experiment*, **2006**(06), P06014.

Baronchelli, Andrea, Dall'Asta, Luca, Barrat, Alain, & Loreto, Vittorio. 2006c. Topology-induced coarsening in language games. *Phys. Rev. E*, **73**(Jan), 015102.

Baronchelli, Andrea, Chater, Nick, Pastor-Satorras, Romualdo, & Christiansen, Morten H. 2012. The biological origin of linguistic diversity. *PloS one*, **7**(10), e48029.

Baronchelli, Andrea, Ferrer-i Cancho, Ramon, Pastor-Satorras, Romualdo, Chater, Nick, & Christiansen, Morten H. 2013. Networks in cognitive science. *Trends in cognitive sciences*, **17**(7), 348–360.

Batali, J. 1998. Computational simulations of the emergence of grammar. *Pages 405–426 of:* Hurford, J. R., Studdert-Kennedy, M., & C., Knight (eds), *Approaches to the Evolution of Language: Social and Cognitive Bases*. Cambridge: Cambridge University Press.

Bedau, Mark A. 2003. Artificial life: organization, adaptation and complexity from the bottom up. *Trends in cognitive sciences*, **7**(11), 505–512.

Belew, Richard K, McInerney, John, & Schraudolph, Nicol N. 1990. Evolving networks: Using the genetic algorithm with connectionist learning. *In: In.* Citeseer.

Berwick, Robert C, Friederici, Angela D, Chomsky, Noam, & Bolhuis, Johan J. 2013. Evolution, brain, and the nature of language. *Trends in cognitive sciences*, **17**(2), 89–98.

Beuls, Katrien, & Steels, Luc. 2013. Agent-Based Models of Strategies for the Emergence and Evolution of Grammatical Agreement. *PloS one*, **8**(3), e58960.

Bickerton, Derek. 2007. Language evolution: A brief guide for linguists. *Lingua*, **117**(3), 510–526.

Bleys, Joris, & Steels, Luc. 2011. Linguistic Selection of Language Strategies. *Pages 150–157 of:* Kampis, G., Karsai, I., & Szathmáry, E. (eds), *Advances in Artificial Life. Darwin Meets von Neumann*, vol. 2. Springer.

Bolhuis, Johan J, Okanoya, Kazuo, & Scharff, Constance. 2010. Twitter evolution: converging mechanisms in birdsong and human speech. *Nature Reviews Neuroscience*, **11**(11), 747–759.

Bungum, Lars, & Gambäck, Björn. 2010. Evolutionary Algorithms in Natural Language Processing. *In: Norwegian Artificial Intelligens Symposium (NAIS).* Tapir Akademisk Forlag.

Cangelosi, Angelo, & Parisi, Domenico. 2002. Computer Simulation: A New Scientific Approach to the Study of Language Evolution. *Chap. 1, pages 3–28 of:* Cangelosi, Angelo, & Parisi, Domenico (eds), *Simulating the Evolution of Language.* London: Springer Verlag.

Castelló, Xavier, Toivonen, Riitta, Eguíluz, Víctor M, Loureiro-Porto, Lucía, Saramäki, Jan, Kaski, Kimmo, San Miguel, Maxi, *et al.* . 2008. Modelling language competition: bilingualism and complex social networks.

Chater, Nick, Reali, Florencia, & Christiansen, Morten H. 2009. Restrictions on biological adaptation in language evolution. *Proceedings of the National Academy of Sciences*, **106**(4), 1015–1020.

Cheney, Dorothy L, & Seyfarth, Robert M. 2005. Constraints and preadaptations in the earliest stages of language evolution. *The Linguistic Review*, **22**(2-4), 135–159.

Chomsky, Noam. 1980. Rules and representations. *Behavioral and brain sciences*, **3**(01), 1–15.

Chomsky, Noam. 1986. *Knowledge of language: Its nature, origins, and use.* Greenwood Publishing Group.

Christiansen, Morten H, & Chater, Nick. 2008. Language as shaped by the brain. *Behavioral and Brain Sciences*, **31**(05), 489–509.

Christiansen, Morten H, & Kirby, Simon. 2003a. Language evolution: Consensus and controversies. *Trends in cognitive sciences*, **7**(7), 300–307.

Christiansen, Morten H, & Kirby, Simon. 2003b. Language evolution: The hardest problem in science? *Studies in the evolution of language*, **3**, 1–15.

Costa, Luciano da Fontoura, Oliveira Jr, Osvaldo N, Travieso, Gonzalo, Rodrigues, Francisco Aparecido, Villas Boas, Paulino Ribeiro, Antiqueira, Lucas, Viana, Matheus Palhares, & Correa Rocha, Luis Enrique. 2011. Analyzing and modeling real-world phenomena with complex networks: a survey of applications. *Advances in Physics*, **60**(3), 329–412.

Crystal, David. 2002. *Language death.* Cambridge University Press.

Dale, Rick, & Lupyan, Gary. 2012. Understanding the origins of morphological diversity: the linguistic niche hypothesis. *Advances in Complex Systems*, **15**(03n04), 1150017.

Dall'Asta, Luca, Baronchelli, Andrea, Barrat, Alain, & Loreto, Vittorio. 2006. Agreement dynamics on small-world networks. *EPL (Europhysics Letters)*, **73**(6), 969.

de Boer, B. 1997. Generating vowel systems in a population of agents. *In:* Husbands, P., & Harvey, I. (eds), *ECAL97.* Cambridge, MA: MIT Press.

De Oliveira, Paulo Murilo Castro, Stauffer, Dietrich, Wichmann, Søren, & De Oliveira, Suzana Moss. 2008. A computer simulation of language families. *Journal of Linguistics*, **44**(3), 659–675.

Dunbar, Robin IM. 1997. Groups, gossip, and the evolution of language. *Pages 77–89 of: New aspects of human ethology.* Springer.

Epstein, Joshua M. 1999. Agent-based computational models and generative social science. *Generative Social Science: Studies in Agent-Based Computational Modeling*, **4**(5), 4–46.

Erdös, & Rèyi. 1959. On random graphs I. *Publ. Math. Debrecen*, **6**, 290–297.

Evans, Nicholas, & Levinson, Stephen C. 2009. The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, **32**(05), 429–448.

Fehl, Katrin, van der Post, Daniel J, & Semmann, Dirk. 2011. Co-evolution of behaviour and social network structure promotes human cooperation. *Ecology letters*, **14**(6), 546–551.

Fitch, W.Tecumseh. 2005. The Evolution of Language: A Comparative Review. *Biology and Philosophy*, **20**(2-3), 193–203.

Freeman, Linton C. 1979. Centrality in social networks conceptual clarification. *Social networks*, **1**(3), 215–239.

Gilbert, Nigel, & Terna, Pietro. 2000. How to build and use agent-based models in social science. *Mind & Society*, **1**(1), 57–72.

Gong, Tao. 2011. Simulating the coevolution of compositionality and word order regularity. *Interaction Studies*, **12**(1), 63–106.

Gong, Tao, & Shuai, Lan. 2012. Modelling the coevolution of joint attention and language. *Proceedings of the Royal Society B: Biological Sciences*, **279**(1747), 4643–4651.

Gong, Tao, & Shuai, Lan. 2013. Computer simulation as a scientific approach in evolutionary linguistics. *Language Sciences*, **40**, 12–23.

Gong, Tao, Ke, Jinyun, Minett, James W, & Wang, William SY. 2004. A computational framework to simulate the coevolution of language and social structure. *Pages 158–64 of: Artificial Life IX: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems.*

Gong, Tao, Minett, James W., & Wang, William S-Y. 2008. Exploring social structure effect on language evolution based on a computational model. *Connection Science*, **20**(2-3), 135–153.

Gong, Tao, Shuai, Lan, & Comrie, Bernard. 2014. Evolutionary linguistics: theory of language in an interdisciplinary space. *Language Sciences*, **41**, 243–253.

Gross, Thilo, & Blasius, Bernd. 2008. Adaptive coevolutionary networks: a review. *Journal of the Royal Society Interface*, **5**(20), 259–271.

Hashimoto, Takashi, Warashina, Katsuhide, & Yamauchi, Hajime. 2010. New composite evolutionary computation algorithm using interactions among genetic evolution, individual learning and social learning. *Intelligent Data Analysis*, **14**(4), 497–514.

Helbing, Dirk, & Balietti, Stefano. 2011. *How to do agent-based simulations in the future: From modeling social mechanisms to emergent phenomena and interactive systems design.* Tech. rept. Tech. Rep. 11-06-024, Santa Fe Institute, NM, USA (June 2011), santa Fe Working Paper.

Hernando, A, Villuendas, D, Vesperinas, C, Abad, M, & Plastino, A. 2010. Unravelling the size distribution of social groups with information theory in complex networks. *The European Physical Journal B*, **76**(1), 87–97.

Hinton, Geoffrey E, & Nowlan, Steven J. 1987. How learning can guide evolution. *Complex systems*, **1**(3), 495–502.

Holland, John H. 1992. Genetic algorithms. *Scientific american*, **267**(1), 66–72.

Jones, Terry. 1995. *Evolutionary algorithms, fitness landscapes and search.* Ph.D. thesis, Citeseer.

Kirby, S. 2002a. Learning, Bottlenecks and the Evolution of Recursive Syntax. *Chap. 6 of:* Briscoe, Ted (ed), *Linguistic Evolution through Language Acquisition: Formal and Computational Models.* Cambridge University Press.

Kirby, Simon. 2002b. Natural Language from Artificial Life. *Artificial Life*, **8**(2), 185–215.

Kirby, Simon, Cornish, Hannah, & Smith, Kenny. 2008. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *PNAS*, **105**(31), 10681–10686.

Klein, Richard G. 2008. Out of Africa and the evolution of human behavior. *Evolutionary Anthropology: Issues, News, and Reviews*, **17**(6), 267–281.

Kuechle, Graciela, & Rios, Diego. 2012. A game-theoretic analysis of the Baldwin effect. *Erkenntnis*, **77**(1), 31–49.

Lei, Chuang, Jia, Jianyuan, Wu, Te, & Wang, Long. 2010. Coevolution with weights of names in structured language games. *Physica A: Statistical Mechanics and its Applications*, **389**(24), 5628–5634.

Lekvam, Torvald, Gambäck, Björn, & Bungum, Lars. 2014. Agent-based modeling of language evolution. *Pages 49–54 of: Proc. of 5th Workshop on Cognitive Aspects of Computational Language Learning (CogACLL)@ EACL.*

Levinson, Stephen C, & Gray, Russell D. 2012. Tools from evolutionary biology shed new light on the diversification of languages. *Trends in cognitive sciences*, **16**(3), 167–173.

Li, Bing, Chen, Guanrong, & Chow, Tommy W.S. 2013. Naming Game with Multiple Hearers. *Communications in Nonlinear Science and Numerical Simulation*, **18**(5), 1214–1228.

Lipowska, Dorota. 2011. Naming game and computational modelling of language evolution. *Computational Methods in Science and Technology*, **17**(1–2), 41–51.

Lipowska, Dorota, & Lipowski, Adam. 2012. Naming Game on Adaptive Weighted Networks. *Artificial Life*, **18**(3), 311–323.

Livingstone, Daniel. 2002. The Evolution of Dialect Diversity. *Chap. 5, pages 99–118 of:* Cangelosi, Angelo, & Parisi, Domenico (eds), *Simulating the Evolution of Language*. London: Springer Verlag.

Loreto, Vittorio, & Steels, Luc. 2007. Social dynamics: Emergence of language. *Nature Physics*, **3**(11), 758–760.

Lu, Qiming, Korniss, Gyorgy, & Szymanski, Boleslaw K. 2009. The naming game in social networks: community formation and consensus engineering. *Journal of Economic Interaction and Coordination*, **4**(2), 221–235.

Lupyan, Gary, & Dale, Rick. 2010. Language structure is partly determined by social structure. *PloS one*, **5**(1), e8559.

Maity, Suman Kalyan, & Mukherjee, Animesh. 2012. Understanding how dominance affects the emergence of agreement in a social network: The case of naming game. *Pages 508–513 of: Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*. IEEE.

Maity, Suman Kalyan, Mukherjee, Animesh, Tria, Francesca, & Loreto, Vittorio. 2013a. Emergence of fast agreement in an overhearing population: The case of the naming game. *EPL (Europhysics Letters)*, **101**(6), 68004.

Maity, Suman Kalyan, Porwal, Abhishek, & Mukherjee, Animesh. 2013b. Understanding how learning affects agreement process in social networks. *Pages 228–235 of: Social Computing (SocialCom), 2013 International Conference on*. IEEE.

Marceau, Vincent, Noël, Pierre-André, Hébert-Dufresne, Laurent, Allard, Antoine, & Dubé, Louis J. 2010. Adaptive networks: Coevolution of disease and topology. *Physical Review E*, **82**(3), 036116.

Mitchell, Melanie. 1998. *An introduction to genetic algorithms*. MIT press.

Munroe, Steve, & Cangelosi, Angelo. 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.

Nardini, Cecilia, Kozma, Balázs, & Barrat, Alain. 2008. Who's talking first? consensus or lack thereof in coevolving opinion formation models. *Physical review letters*, **100**(15), 158701.

Newman, Mark EJ. 2003. The structure and function of complex networks. *SIAM review*, **45**(2), 167–256.

Noble, Jason. 1999. Cooperation, conflict and the evolution of communication. *Adaptive Behavior*, **7**(3-4), 349–369.

Nowak, Martin A, Plotkin, Joshua B, & Krakauer, David C. 1999. The evolutionary language game. *Journal of Theoretical Biology*, **200**(2), 147–162.

Oliphant, Michael. 1999. The learning barrier: Moving from innate to learned systems of communication. *Adaptive behavior*, **7**(3-4), 371–383.

Pagel, Mark. 2000. The history, rate and pattern of world linguistic evolution. *In:* Knight, Ch., Hurford, J.R., & Studdert-Kennedy, M. (eds), *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*. Cambridge University Press.

Patriarca, Marco, & Heinsalu, Els. 2009. Influence of geography on language competition. *Physica A: Statistical Mechanics and its Applications*, **388**(2), 174–186.

Pinker, Steven, & Bloom, Paul. 1990. Natural language and natural selection. *Behavioral and Brain Sciences*, **13**, 707–784.

Puglisi, Andrea, Baronchelli, Andrea, & Loreto, Vittorio. 2008. Cultural route to the emergence of linguistic categories. *Proceedings of the National Academy of Sciences*, **105**(23), 7936–7940.

Quillinan, Justin. 2006. *Social Networks and Cultural Transmission*.

Rendell, Luke, Fogarty, Laurel, & Laland, Kevin N. 2010. Rogers' Paradox Recast and Resolved: Population Structure and the Evolution of Social Learning Strategies. *Evolution*, **64**(2), 534–548.

Rendell, Luke, Fogarty, Laurel, Hoppitt, William JE, Morgan, Thomas JH, Webster, Mike M, & Laland, Kevin N. 2011. Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends in cognitive sciences*, **15**(2), 68–76.

Roberts, Sean Geraint. 2013. Evolutionary approach to bilingualism.

Ross, Peter, & Corne, Dave. 1994. Applications of genetic algorithms. *AISB Quaterly on Evolutionary Computation*, **89**, 23–30.

Sasaki, Takahiro, & Tokoro, Mario. 1999. Evolving learnable neural networks under changing environments with various rates of inheritance of acquired characters: comparison of Darwinian and Lamarckian evolution. *Artificial Life*, **5**(3), 203–223.

Scott-Phillips, Thomas C, & Kirby, Simon. 2010. Language evolution in the laboratory. *Trends in cognitive sciences*, **14**(9), 411–417.

Simpson, George Gaylord. 1953. The Baldwin effect. *Evolution*, **7**(2), 110–117.

Smith, Kenny, Kirby, Simon, & Brighton, Henry. 2003. Iterated learning: A framework for the emergence of language. *Artificial Life*, **9**(4), 371–386.

Spranger, Michael. 2013. Evolving Grounded Spatial Language Strategies. *KI-Künstliche Intelligenz*, **27**(2), 1–10.

Steels, Luc. 1995. A self-organizing spatial vocabulary. *Artificial Life*, **2**(3), 319–332.

Steels, Luc. 1999. *The Talking Heads Experiment. Volume 1. Words and Meanings.* Antwerpen: Laboratorium.

Steels, Luc. 2009. Cognition and social dynamics play a major role in the formation of grammar. *In: Biological Foundations and Origin of Syntax*, vol. 3. MIT Press.

Steels, Luc. 2011. Modeling the cultural evolution of language. *Physics of Life Reviews*, **8**(4), 339–356.

Steels, Luc. 2012. Self-organization and selection in cultural language evolution. *Experiments in Cultural Language Evolution. John Benjamins, Amsterdam.*

Steels, Luc, & Loetzsch, Martin. 2012. The Grounded Naming Game. *Pages 41–59 of:* Steels, L. (ed), *Experiments in Cultural Language Evolution.* John Benjamins.

Steels, Luc, Spranger, Michael, Trijp, Remi, Höfer, Sebastian, & Hild, Manfred. 2012. Emergent Action Language on Real Robots. *Pages 255–276 of:* Steels, Luc, & Hild, Manfred (eds), *Language Grounding in Robots*. Springer US.

Suzuki, Reiji, & Arita, Takaya. 2008. *How Learning Can Guide Evolution of Communication.*

Suzuki, Reiji, & Arita, Takaya. 2013. A simple computational model of the evolution of a communicative trait and its phenotypic plasticity. *Journal of theoretical biology*, **330**, 37–44.

Swarup, Samarth, & Gasser, Les. 2009. The Iterated Classification Game: A New Model of the Cultural Transmission of Language. *Adaptive Behavior*, **17**(3), 213–235.

Swarup, Samarth, & Gasser, Les. 2010. The classification game: combining supervised learning and language evolution. *Connection Science*, **22**(1), 1–24.

Tattersall, Ian. 2010. Human evolution and cognition. *Theory in Biosciences*, **129**(2–3), 193–201.

van Trijp, Remi, & Steels, Luc. 2012. Multilevel alignment maintains language systematicity. *Advances in Complex Systems*, **15**(03n04), 1250039.

Vogt, Paul. 2005. The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence*, **167**(1–2), 206 – 242. ¡ce:title¿Connecting Language to the World¡/ce:title¿.

Vogt, Paul. 2009. Modeling interactions between language evolution and demography. *Human biology*, **81**(2), 237–258.

Wagner, Kyle, Reggia, J.A., Uriagereka, J., & Wilkinson, G.S. 2003. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, **11**(1), 37–69.

Wang, WX, Lin, BY, Tang, CL, & Chen, GR. 2007. Agreement dynamics of finite-memory language games on networks. *The European Physical Journal B*, **60**(4), 529–536.

Watanabe, Yusuke, Suzuki, Reiji, & Arita, Takaya. 2008. Language evolution and the Baldwin effect. *Artificial Life and Robotics*, **12**(1-2), 65–69.

Watts, Duncan J, & Strogatz, Steven H. 1998. Collective dynamics of 'small-world'networks. *nature*, **393**(6684), 440–442.

Weise, Thomas, Chiong, Raymond, & Tang, Ke. 2012. Evolutionary optimization: Pitfalls and booby traps. *Journal of Computer Science and Technology*, **27**(5), 907–936.

Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Basil Blackwell.

Wolf, Jochen BW, Mawdsley, David, Trillmich, Fritz, & James, Richard. 2007. Social structure in a colonial mammal: unravelling hidden structural layers and their foundations by network analysis. *Animal Behaviour*, **74**(5), 1293–1302.

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