Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

School of Computer Science and Statistics

Evolutionary Naming Game: an Agent-based Model for the Emergence of Language

Paolo Moloney
Supervisor: Martin Emms

September 14, 2021

A Masters Dissertation submitted in partial fulfilment
of the requirements for the degree of
M.Sc. Computer Science - Augmented and Virtual Reality

Declaration

I hereby declare that this project is entirely my own work and that it has not been submitted as an exercise for a degree at this or any other university.

Signed: _____   Date: _____

Abstract

Computational models have been the standard avenue of research into the emergence and evolution of language in recent decades. One such model, the Evolutionary Naming Game, combines genetic algorithms with a traditional Naming Game, and is reported to demonstrate the Baldwin Effect in its results. This dissertation provides a thorough implementation and evaluation of this model, along with pseudocode descriptions of the algorithms for use by future researchers. The reported results of the original model were verified, but the same results could not be observed in modified versions, putting in doubt the original author's claims of robustness. Some suggestions for continued use of this model in investigating the Critical Period Hypothesis and the evolution of social behaviour are made.

Acknowledgements

Many thanks to Martin Emms for his supervision and guidance over the course of this project, and to Dorota Lipowska for her correspondence and help in the implementation of her model.

Contents

List of Figures

List of Algorithms

Nomenclature

kya	kilo (thousand) years ago
NG	Naming Game
ENG	Evolutionary Naming Game
UG	Universal Grammar
LAD	Language Acquisition Device

1   Introduction

Language is what distinguishes us from every other animal we share the planet with. The ability to transfer a thought or an idea from one mind to another through symbolic gestures and sounds has undoubtedly been one of the reasons behind the success of our species. The question of how language emerged in humans is deeply fascinating and spans disciplines from evolutionary biology, to linguistics, to anthropology.

Working backwards in time from the present moment, linguists have reconstructed the proto-languages that were spoken by our ancestors. The common ancestor to most modern European languages, among others, is the so called Proto-Indo-European language, with the foremost hypotheses as to its origin being the "Kurgan hypothesis" and the "Anatolian hypothesis". The Kurgan hypothesis, currently the most widely accepted, suggests that Proto-Indo-European speaking peoples originated in the area between the Black Sea and the Caspian Sea about 6 kya and cites the presence of haplogroups R1a and R1b in modern Europeans as well as South Asians as evidence [8, 9]. The Anatolian hypothesis places the origin in Anatolia some 8-9.5 kya and is consistent with the known spread of agriculture in the Neolithic period. This hypothesis is backed by a Bayesian analysis carried out in [2], however it has received criticism for its methodology and thus is not as widely accepted.

In any case these methods can only take us back a few millennia, whereas evidence for the origin of language in Africa spans the range from 50-200 kya [10, 11] roughly coinciding with the emergence of anatomically modern Homo sapiens [12, 13]. The theories for the evolution of a language ability can be categorised primarily into two camps for which many different terminologies exist. For the purpose of this dissertation, the theories will be referred to using the terminology described in [14]; namely "Discontinuity" theories, arguing that language is saltational, i.e. it appeared suddenly from one generation to the next in one or very few individuals, and "Continuity" theories, arguing that language is an adaptive behaviour that evolved gradually. These theories are further elaborated on in Section 2.1.

The lack of direct evidence makes any theorising beyond these dates speculative in nature, and as such has led to the prevalence of computational models and simulations being the main drivers of research in the field since the 1990's.
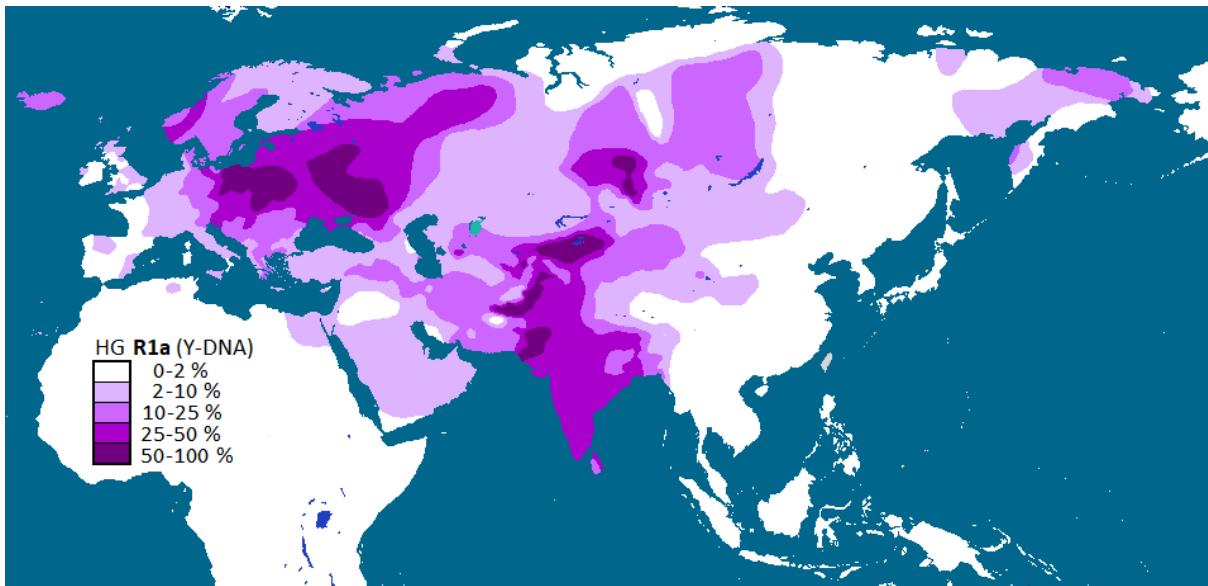
Figure 1.1: Distribution of Y-DNA haplogroup R1a. Areas of high incidence are consistent with the hypothesised spread of Indo-European people groups and languages. Shared with permission under Creative Commons Attribution-Share Alike 4.0 International license from [1].
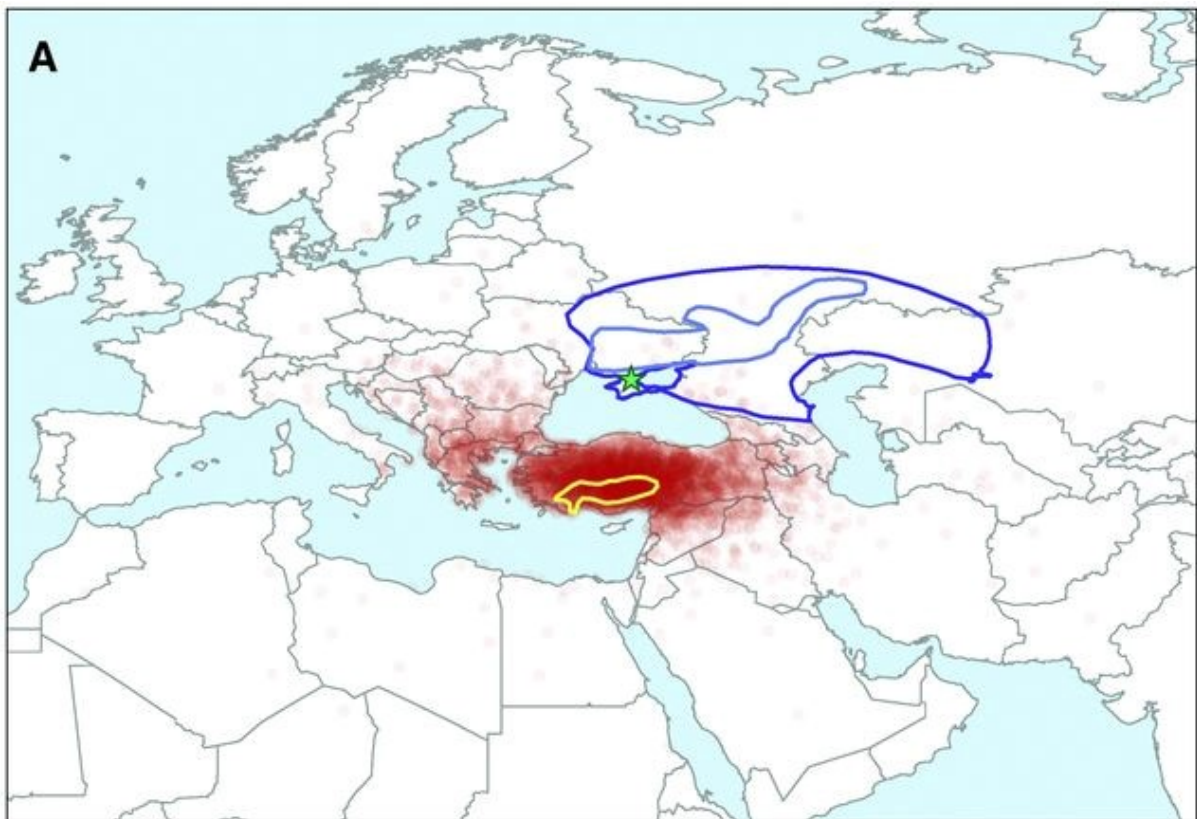
Figure 1.2: Inferred geographic origin of the Indo-European language family according to the Bayesian analysis in [2]. Areas of high probability of origin are plotted in red, such that darker red corresponds to higher probability. Blue and yellow polygons correspond to the hypothesised origins according to the Kurgan and Anatolian hypotheses respectively. Adapted from [3].

1.1 Goal

The goal of dissertation is to investigate ways in which agent-based models, namely through the use of Naming Games (NGs), can be used to simulate the emergence of language in humans. This is mainly achieved through looking at and implementing one particular Evolutionary Naming Game (ENG), first published by Dorota Lipowska in [7].

To facilitate the discussion, the project can be split into a number of research questions:

Research Question 1 Can the results of the Lipowska ENG be verified?

> This question firstly involves implementing the ENG model described by Lipowska [6, 7]. The second phase involves recreating the experiments from the original papers and comparing the results in order to verify the model. This section of the project will be discussed in Chapter 3.

Research Question 2 How robust is the Lipowska ENG to modifications?

> The Lipowska ENG is claimed by the author to be robust to modifications of the survival probability function and parameters, and different breeding and mutation rules. The descriptions of some modifications and experiments run will be provided in Chapter 4.

1.2 Contribution

The main contributions of this dissertation are the following:

- The verification of the Lipowska ENG described in [6, 7].

- The definition of clear pseudocode descriptions of the algorithms used in the Lipowska ENG for use by future researchers.

- An analysis of the purported robustness of the Lipowska ENG with respect to some modifications of its rules.

1.3 Overview

Chapter 2 of this dissertation contains a review of the background theory and related work relevant to this project. Chapter 3 describes the implementation of Lipowska ENG in detail, followed by an analysis of the results with respect to the original findings. Chapter 4 details the modifications made to Lipowska's model and the experiments carried out. Chapter 5, contains a discussion and analysis of the results of the models described in Chapters 3 and 4. Finally, Chapter 6 concludes this dissertation with some closing remarks and suggestions for future work in Section 6.1.

2 Background & Related Work

This chapter deals with the background knowledge required to understand the material discussed in subsequent chapters. Firstly, the prevailing theories behind the Emergence of Language are discussed in some detail, followed by an introduction to some techniques used to simulate the emergence of language, with a particular focus on Naming Games.

2.1 Theories on the Emergence of Language

Chomsky's Universal Grammar (UG) hypothesis [15, 16] attributes the faculty of language to a "system of categories, mechanisms, and constraints ... considered to be innate" [17]. According to Chomsky there is an underlying universal structure that allows the construction of all natural languages from a set of rules, i.e. a grammar, and that such a grammar is genetically encoded into the human brain in the form of an organ or module, sometimes referred to as a Language Acquisition Device (LAD) [16]. In Chomsky's own words: "UG may be regarded as a characterization of the genetically determined language faculty. One may think of this faculty as a 'language acquisition device', an innate component of the human mind that yields a particular language through interaction with presented experience, a device that converts experience into a system of knowledge attained: knowledge of one or another language." [18]

This has become the most widely accepted model (since the 1950's) for *how* humans are able to acquire and use language, but the main area of debate centers around the *origin* of such an LAD. As mentioned in Section 1.1, these origin theories are categorised by Parravicini et al. [14] into Discontinuity theories and Continuity theories, and will be referred to as such throughout this dissertation. The following sections describe these theories in more detail.

2.1.1 Discontinuity Theories

In describing the origin of language faculty, Chomsky argues that language, being such a complex and rare behaviour in the natural world, could not have emerged gradually as a Darwinian adaptation. Instead the LAD, which is in its most basic form a set of symbols and

rules, and since an organism cannot possess half a symbol or follow 5% of a rule, must be acquired on an all-or-nothing basis and *not* through a gradual process [14]. What evolutionary advantage does possessing a rudimentary version of a complex organ carry if the organ can only perform its function when it is fully developed? This paradox is why, according to the discontinuist approach, natural selection cannot explain language and there must be an evolution-free explanation.

2.1.2   Continuity Theories

Pinker and Bloom, two former pupils of Chomsky, take a different approach in their paper "Natural language and natural selection" [19]. If language is an innate and universal faculty in humans, then it follows that it *must* be a biological adaptation, and therefore necessarily must be explained by Darwinian evolution. In order for language to be explicable by natural selection, a number of prerequisite conditions must be true:

Genetic variation There must have been genetic variation in the grammatical competence of individuals.

Intermediate steps There must have been a series of steps leading from no language at all to language as it exists now, with each intermediate step being useful to its possessor.

Reproductive advantage Every element of the language faculty must have brought with it a reproductive advantage large enough to have become fixed in the ancestral population.

Phyletic continuity There must have been enough evolutionary time and genomic space separating our species from our nonlinguistic primate predecessors.

There is little conclusive evidence for these prerequisites and the discontinuist approach conflates this absence of evidence with evidence of absence, which Pinker and Bloom regard as a premature conclusion. Instead they argue that there is in fact biological evidence for these postulates.

- Within the "normal" range of grammatical ability, there is a spectrum of the effective use of language by individuals, from those who use "tangled syntax" to those who "speak with elegance", which to some degree must be affected by genetics, and beyond this "normal" range there exist genetically transmitted grammatical deficits.

- Language comprehension abilities do not have to be in perfect synchrony with language production abilities, as is evident in the fact that a non-Italian speaker could rely on cognates and general context to understand an Italian newspaper article to a reasonable degree, allowing possessors of a genetic innovation to retain the ability to communicate with their peers while also possessing an advantage.

- A selective pressure for an organism to increase its size may be so small that from one

generation to the next, it can't be measured against the noise of individual variation, but this does not mean that given enough time the organism would not grow; a reproductive advantage does not have to be large to ultimately have an effect on an evolutionary timescale, as long as it is persistent.

- The absence of language faculty in closely-related species such as chimpanzees does not exclude the possibility that it may have evolved in some of our ancestor species that diverged from the chimpanzee-human common ancestor, or the fact that chimpanzees themselves have evolved since the divergence.

Currently the only mechanism that is capable of describing the design of a complex organ such as the LAD is natural selection, and the implication that its complexity or uniqueness exempts it from this explanation is false. Pinker writes in [20]: "if human language is unique in the modern animal kingdom, ... , the implications for a Darwinian account of its evolution would be as follows: none. A language instinct unique to modern humans poses no more of a paradox than a trunk unique to modern elephants."

2.1.3   Baldwin Effect

In discussing the Continuist approach it is important to also discuss a phenomenon known as the Baldwin Effect. The Baldwin Effect, first written about by J. M. Baldwin in 1896 [21], describes the way in which a learned behaviour can guide the course of an organsim's evolution. Superficially this seems similar to a Lamarckian process by which information is passed from the phenotype to the genotype, e.g. an organism transmitting a behaviour that it learned during its lifetime to its children genetically, which is known to be impossible, but it is in fact distinct. Although the Baldwin Effect itself is hard to conclusively observe, there are a number of observable mechanisms which can be thought to demonstrate it.

Smoothing of the Fitness Landscape

One such mechanism is observed in a simulation carried out by Hinton and Nowlan [4], which describes how learning carried out through an organism's lifetime, i.e. changes to its phenotype, can guide the evolution of subsequent generations, without communicating any information back to the genotype. Consider a population of organisms, each with a neural network of 20 connections, which can have a value of either 0, 1. The neural network confers a large fitness advantage if the connections are configured in one specific way and no advantage otherwise. In the case where the values of the connections are all genetically determined, there would be no incremental evolutionary path to the right configuration and as such the evolutionary search would be to try possibilities at random, with only 1 in every 20^2 organisms having the correct one. This extremely unlikely case is what is proposed to have happened by Discontinuists in human evolution.
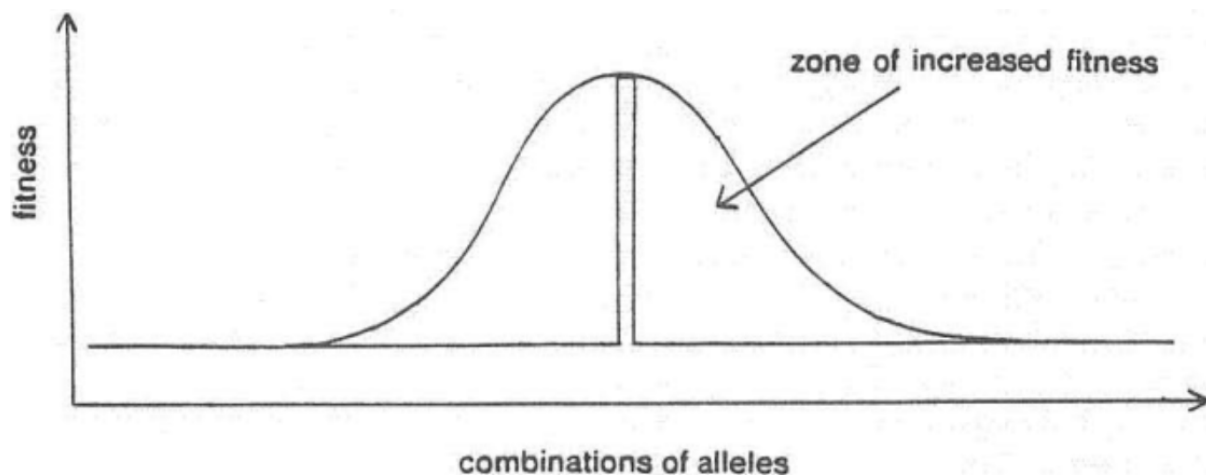
Figure 2.1: Effective smoothing of the fitness landscape when phenotypic plasticity is enabled. Adapted from [4].

What if instead the organism's genotype contained three possible alleles for each connection determining which value that connection should take; 0 or 1, indicating the connection takes a value of 0 or 1 respectively, or ?, indicating that the value of the connection can be determined by trial and error over an organism's lifetime, i.e. that it can be learned. An organism with ten correct alleles and ten ? alleles would stand a reasonable chance of learning the correct configuration during its lifetime, and thus of gaining the massive fitness advantage, which in turn would make it more likely to pass on its genes to the next generation. The phenotypic plasticity allowed by the ? allele has the effect of smoothing the fitness landscape (see Figure 2.1), creating a fitness gradient from organisms with few correct alleles to organisms with many correct alleles. In their trials, Hinton et al. find that the correct configuration appears far earlier than the 2^{20} organisms required for the evolutionary search.

This simulation is well regarded as a solid demonstration of the Baldwin Effect and is even cited as evidence in [19].

Niche Construction

Another such mechanism is that of niche construction, whereby an organism's behaviour and subsequent change of its environment can modify the selection pressures present [22]. All living organisms interact with their environment, by taking energy and resources from it, excreting byproducts, and building structures like nests and shelters to name a few examples. In doing so they will often have an effect on the environment around them at different scales, from local to global levels, which can lead to the modification of existing or emergence of new selection pressures, and, if the behaviour is persistent enough, subsequent generations can then adapt to this niche in the modified environment. Thus the learned behaviour of that organism can guide its evolution, i.e. the Baldwin Effect.

An example of a niche constructing behaviour from human history is the link between dairy

farming and lactose tolerance. Humans may have begun harvesting milk when they first domesticated cattle but wouldn't have been able to take full advantage of the new food source as lactose tolerance generally declines with age after childhood. Nonetheless, this new behaviour would have resulted in the modification of the selection pressures, whereby it would have been advantageous for people to retain lactose tolerance into adulthood and avail of an additional food source. Given that the behaviour was persistent and dairy foods continued to be available, individuals with an increased tolerance to lactose would have also experienced increased fitness relative to their lactose intolerant peers, and the cycle continued until the present day. There is evidence from genetic studies such as [5] that in areas where dairy farming has been practised for longer, there are also higher levels of lactose tolerance (see Figure 2.2).

2.2   Simulating the Emergence of Language

Many different types of language games exist for simulating the emergence and evolution of languages, with perhaps the most common being known as the Naming Game. One of the primary contributors to the various iterations of NG over the years is Luc Steels, with the first such model being laid out in 1995 [23].

2.2.1   Naming Game

All NGs follow a basic structure with other variants adding complexity on to the same underlying framework. The goal in a NG is usually for a group of initially linguistically unskilled agents to establish a shared vocabulary for referring to some objects or concepts in their world. Each step involves two agents, a speaker and a listener, who communicate a word or name and update their internal vocabularies should a new word be encountered. Specifically the steps are as follows:

1. A speaker and listener are selected from the pool of agents, along with an object they are to communicate about.

2. The speaker utters the name of this object in their vocabulary and communicates it to the listener.

3. The listener looks up this name in their vocabulary and communicates the corresponding object to the speaker.

4. If the object the listener understood is the same as the object the speaker communicated about, then the naming game is a success.

5. Otherwise the game is a failure.

Agents construct vocabularies of name-object pairs using the following criteria:
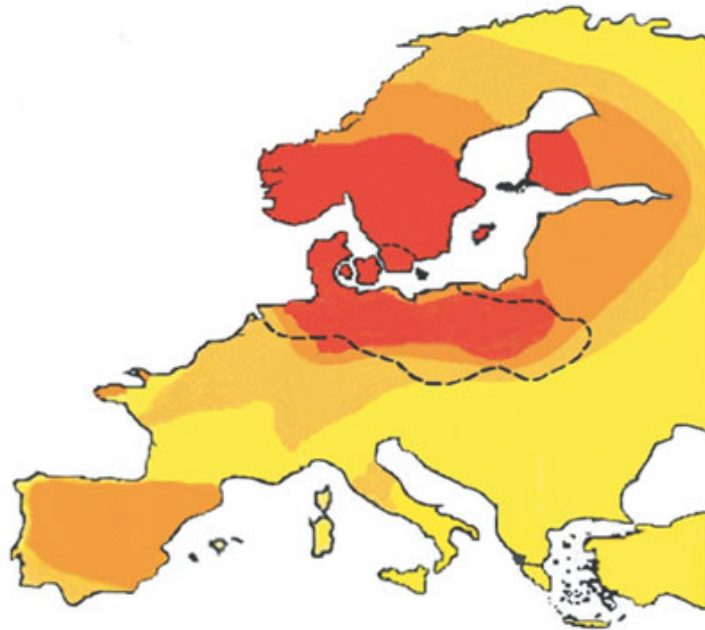
(a) Measured diversity in milk producing alleles of cattle breeds from low (yellow) to high (red)
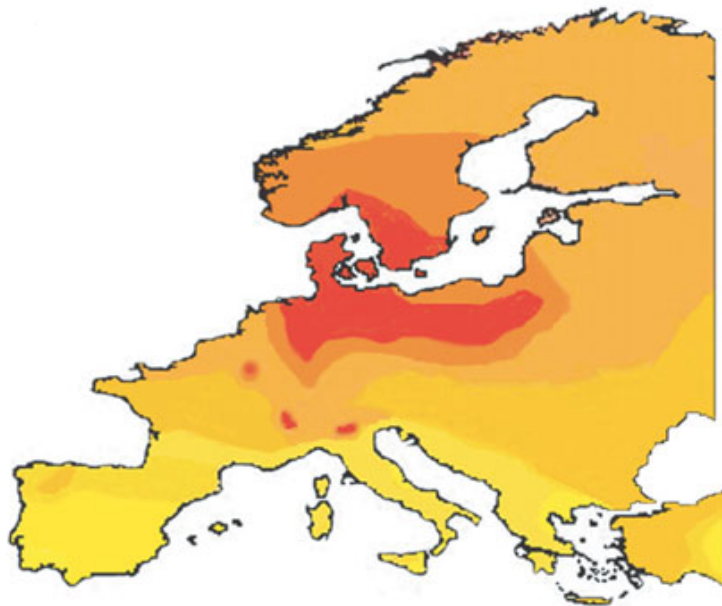
(b) Incidence of lactase persistence allele in humans from low (yellow) to high (red)

Figure 2.2: Figure 2.2a shows the genetic diversity of cattle, with areas of high genetic diversity indicating that domestication occurred longer ago an vice versa. Figure 2.2b shows the frequency of the allele which allows the production of lactase into adulthood. The high degree of correlation suggests that the cultural behaviour of dairy farming and consumption lead to the evolution of lactose tolerance in humans, a possible example of the Baldwin Effect in action. Adapted from [5].

- If the speaker contains no name in its vocabulary for the selected object, then they will invent a new one.

- If the listener does not know the name uttered by the speaker, they will add the name and object to their vocabulary.

- If the listener and the speaker use the same word to refer to different objects, the game is a failure and the listener will add the speaker's name-object pair to their vocabulary.

The interesting thing to note is that after an initial period of linguistic chaos where agents are trying to establish the "correct" names to use, they will eventually converge on a shared vocabulary without any rules as to what that vocabulary should look like. A consensus emerges without any governing body deciding which names should be used for which objects. This type of NG exhibits "horizontal" communication, that is communication that occurs between agents in the same generation, and is generally useful when trying to study the cultural evolution of a language or the diffusion of new words throughout a population.

2.2.2 Alignment

The vocabularies established by agents in a traditional NG are in no way optimal and often will contain a large number of synonyms and homonyms, resulting in very large vocabulary sizes as agents try to account for every name of every object possessed by every agent. Steels et al. [24] describe the process of alignment as a way of addressing this issue. Each name-object pair is associated with a score which is incremented on successful communication attempts using that name, and decremented on unsuccessful attempts. When a vocabulary entry reaches a score of 0 it is discarded. This leads to much more optimised vocabulary sizes and therefore reduces the computational complexity of such an NG when dealing with sufficiently large numbers of agents.

2.2.3 Evolutionary Naming Game

NGs can be combined with evolutionary algorithms, where agents are made to reproduce, die, and mutate. This allows the study of both "horizontal" communication and "vertical" communication between generations. As the agents in a traditional NG will be the same agents each iteration, there is not much interesting behaviour to be observed once a shared vocabulary has been established. The agents of an ENG on the other hand will constantly be changing as old agents die and new agents replace them, allowing us to observe both the horizontal and vertical evolution of culture, along with biological evolution on the vertical axis.
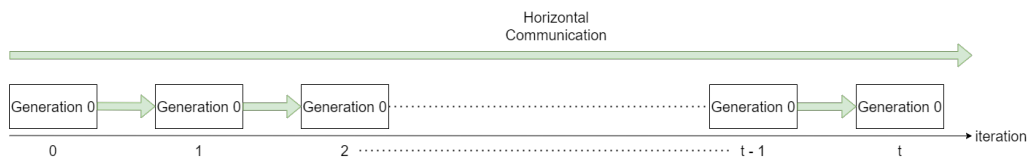
(a) NG
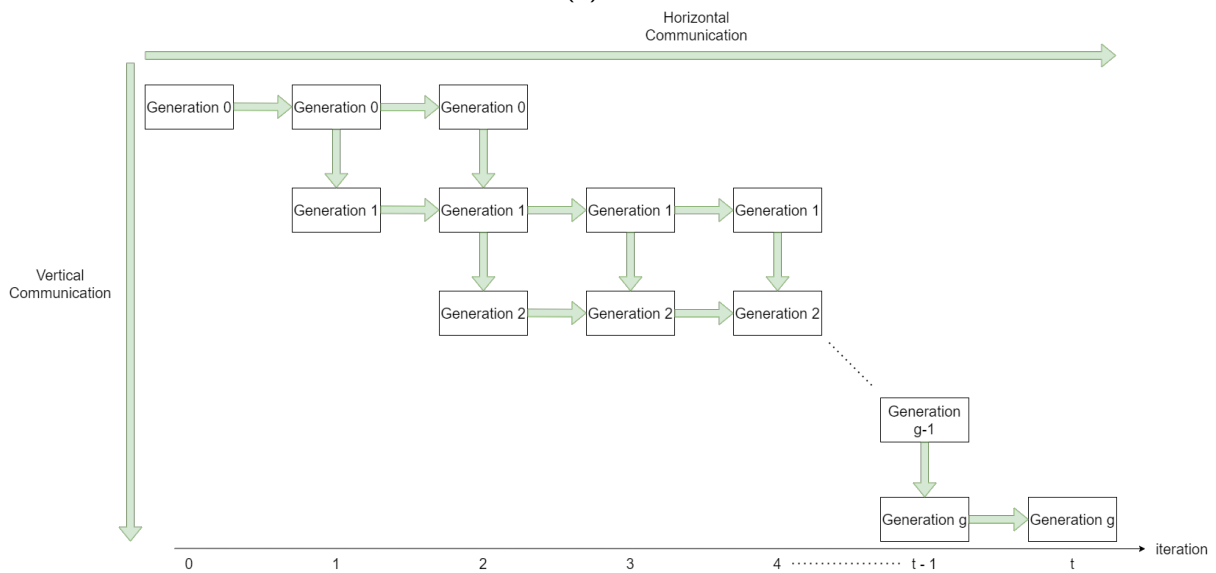
(b) ENG

Figure 2.3: Figure 2.3a demonstrates the horizontal communication that takes place between agents in the same generation in a traditional NG. Figure 2.3b demonstrates the horizontal and vertical communication taking place between agents in the same generation and subsequent generations in an ENG.

3 Lipowska ENG

This chapter covers where the bulk of the work on this project was spent. It begins with a description of Lipowska's ENG model as outlined in [6, 7], followed by detailing the implementation of the model, and finally a discussion and comparison of the results to those reported in the original papers.

3.1 Model Description

The Lipowska ENG, as the name suggests, combines elements of evolutionary naming games with alignment. Agents occupy a grid and can only interact with the eight grid cells adjacent to their own, either through communication or reproduction. The ENG continues until the specified number of iterations has been reached or there are no longer any living agents: extinction. A number of user-defined parameters control the model: world size L, communication probability p, mutation probability p_{mut}, age weighting parameter A, knowledge weighting parameter B, maximum lexicon size N.

Agents are placed in a square grid of size $L * L$. Each agent k has an internal lexicon λ_k of size N with associated weights $w_j : 1 \leq j \leq N$, and an inherited learning ability l_k reflecting the ease with which words are learned, and an age a_k, which is incremented at each timestep t.

Each timestep consists of $L * L$ individual updates (see Figure 3.1 for a flow diagram of one individual update), such that on average each agent is chosen to speak once and hear once per t. Communication follows roughly the same process as a traditional NG, as described in Section 2.2, with the added mechanic of alignment in the form of the word weights.
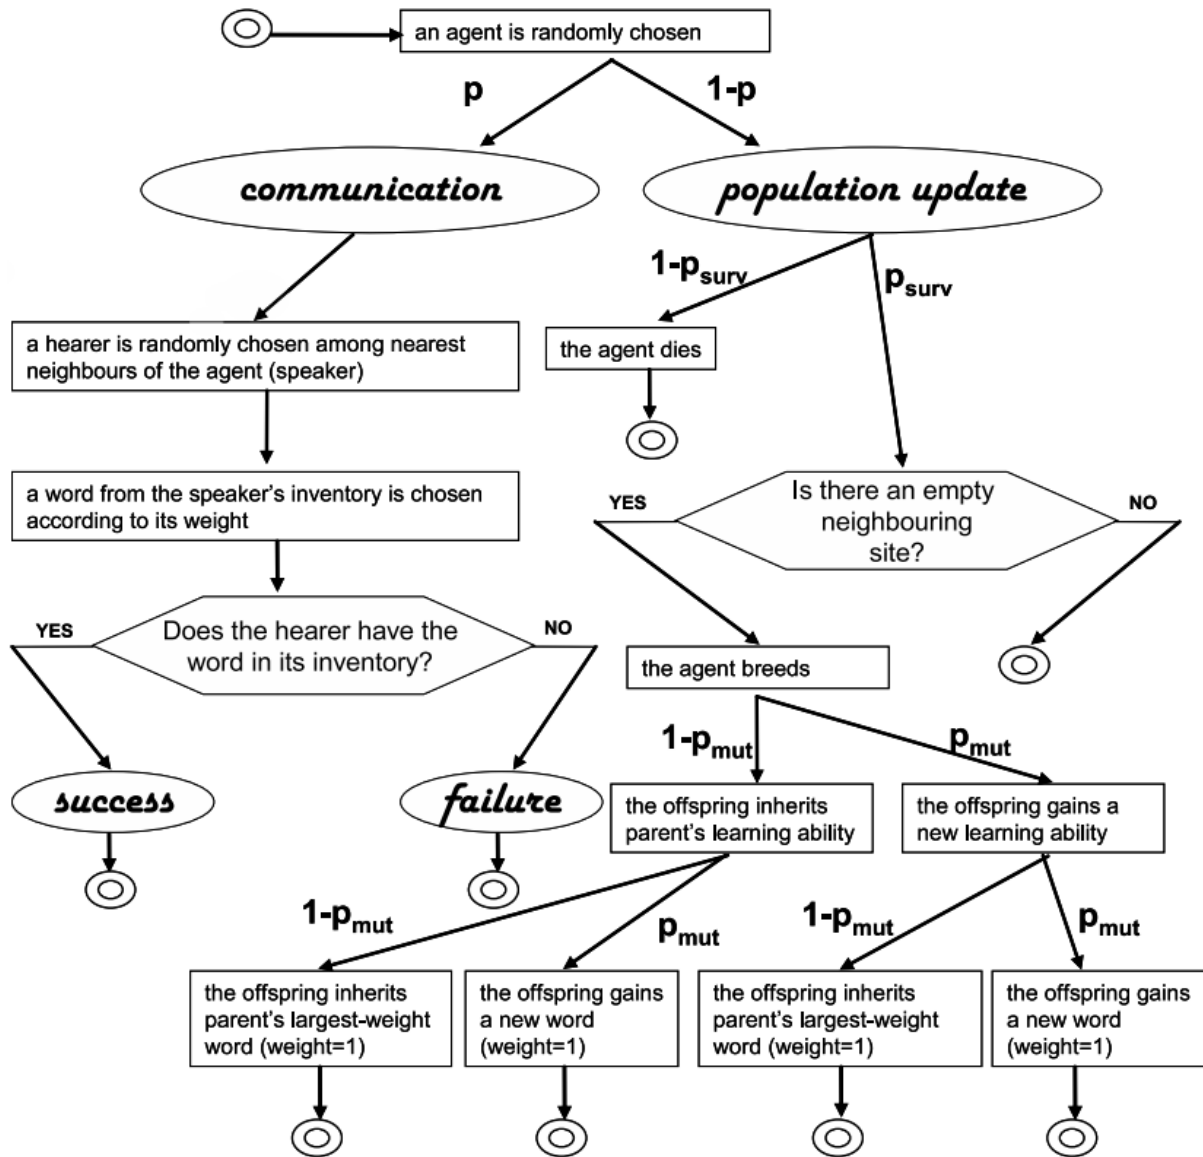
Figure 3.1: Flow diagram describing an individual update of one agent in the Lipowska ENG. Reproduced with permission from [6].

3.2 Implementation

This section contains pseudocode descriptions of the algorithms described in Section 3.1. Algorithm 1 describes the main simulation loop, and Algorithms 2, 3, and 4 describe the communication, population update, and reproduction respectively. These pseudocode descriptions were developed with reference to the original descriptions of the model in [6, 7] and in correspondence with the author of the papers, Dorota Lipowska, to minimise the possible sources of error should others wish to recreate the model for further research. The model was implemented in the Python programming language using Jupyter Notebooks.

An agent's knowledge is defined as the sum of word weights in their lexicon, and is reflective of their linguistic ability. Agents who have participated in many successful communications

will tend to have higher knowledge relative to those who haven't. Additionally, agents with a higher learning ability will tend to accumulate knowledge more quickly than their counterparts with low learning ability.

W is the average knowledge over all agents (see Equation 3.2.1), where K is the number of agents currently alive, and is used in calculating an agent's individual survival probability p_{surv} (see Equation 3.2.2), also known as fitness, which is a decreasing function of the agent's age and an increasing function of their knowledge. The parameters A and B are control parameters used to weight the importance of age and knowledge respectively when computing p_{surv}. While performing preliminary experiments with different values of A and B, it was found that the model is quite sensitive to these parameters, and any change resulted in extinction within very few iterations.

Words are generated by randomly selecting a sequence of four lowercase letters, such that the probability of two agents inventing the same word independently is very small. The maximum lexicon size was kept fixed at $N = 1$ for all experiments, as it was in Lipowska's own experiments.

$$W = \sum_{k=1}^{K} \sum_{j=1}^{N} (w_j \in \lambda_k)/K \tag{3.2.1}$$

$$p_{surv} = \exp(-Aa_k) \times [1 - \exp(-B \sum_{j=1}^{N} w_j/W)] \tag{3.2.2}$$

3.3   Results

The interesting observation made in the results of the original papers is that of the "biolinguistic" transition. This transition is reported to occur when the communication probability p exceeds a threshold value ($p \approx 0.23$ in the original, $p \approx 0.12$ in the implementation), and has two distinct but related components; the biological and the linguistic. Moreover this behaviour is concluded to be a demonstration of the Baldwin Effect.

Figure 3.2 shows a comparison of the original results of the model published in [7] with those obtained from the recreated model implemented as described in the previous section. This experiment is intended to be a very obvious example of the biolinguistic transition; the simulation begins with a low value of $p = 0.1$ and increases abruptly to $p = 0.98$, almost the maximum value, at $t = 8000$. Shortly after, the success rate climbs to nearly 1.0, indicating that the vast majority of communications are successful and that the agents have established a shared language, which has become diffused throughout the population. After a transitory decrease, the learning ability also rises, implying that the selection pressure for a high

---

Algorithm 1 Lipowska ENG: Main simulation loop

Input: L, world grid size
Input: p, communication probability
Input: p_{mut}, mutation probability
Input: A, age weighting parameter
Input: B, knowledge weighting parameter
Input: N, maximum lexicon size
Input: *agents*, a list of all living agents

  1: **for** timestep $t = 1, 2, ...$ **do**
  2:     **for** update $u = 1, 2, ..., L * L$ **do**
  3:         Select random agent $k_{speaker}$ from *agents*
  4:         Select random number r from $[0, 1]$
  5:         **if** $r < p$ **then**   ▷ with probability p
  6:             **if** all neighbouring tiles are empty **then**
  7:                Skip to *population_update*
  8:             **else**
  9:                Select random agent $k_{listener}$ from neighbours of $k_{speaker}$
10:                *communication_update($k_{listener}$, $k_{speaker}$)*
11:             **end if**
12:         **else**   ▷ with probability $1 - p$
13:             *population_update($k_{speaker}$)*
14:         **end if**
15:     **end for**
16:     **for** agent k in *agents* **do**
17:         $a_k + = 1$
18:     **end for**
19: **end for**

Algorithm 2 Lipowska ENG: Communication update

Input: $k_{speaker}$, agent who is speaking
Input: $k_{listener}$, agent who is listening

  1: **if** $\lambda_{speaker}$ is empty **then**
  2:     Generate random new word i
  3:     Add i to $\lambda_{speaker}$ with weight $w_i = 1$
  4: **end if**
  5: Select random word i from $\lambda_{speaker}$ with probability $w_i / \sum_{j=1}^{N} w_j$
  6: **if** $\lambda_{listener}$ contains i **then**
  7:     Increment w_i from $\lambda_{speaker}$ by $l_{speaker}$
  8:     Increment w_i from $\lambda_{listener}$ by $l_{listener}$
  9:     **return** *success*
10: **else**
11:     Decrement w_i from $\lambda_{speaker}$ by $l_{speaker}$
12:     Add i to $\lambda_{listener}$ with weight $w_i = 1$
13:     **return** *failure*
14: **end if**

Algorithm 3 Lipowska ENG: Population update

Input: k, agent undergoing the population update
  1: Compute average lexicon weight W
  2: Compute survival probability p_{surv}
  3: Select random number r from $[0, 1]$
  4: **if** $r < p_{surv}$ **then**   ▷ with probability p_{surv}
  5:     Agent k survives
  6:     **if** a neighbouring tile is empty **then**
  7:         *reproduce()*
  8:     **end if**
  9: **else**   ▷ with probability $1 - p_{surv}$
10:     Agent k dies
11:     Remove agent k from list of all agents
12: **end if**

Algorithm 4 Lipowska ENG: Reproduction and inheritance

  1: Select random number r from $[0, 1]$
  2: **if** $r < p_{mut}$ **then**   ▷ with probability p_{mut}
  3:     $l_k = l_{parent}$
  4:     Select highest weighted word i from λ_{parent}
  5:     Add i to λ_k with weight $w_i = 1$
  6: **else**   ▷ with probability $1 - p_{mut}$
  7:     Select random number l from $[0, 1]$
  8:     $l_k = l$
  9:     Generate random new word i
10:     Add i to λ_k with weight $w_i = 1$
11: **end if**
12: Add new agent k to list of all agents
---

learning ability has increased greatly. The rate begins to ease after roughly 1000 iterations, as the proportion of mutations which are higher than the current average learning ability becomes increasingly smaller, giving the curve a distinct sigmoid shape. The constituents of the biolinguistic transition are evident in Figure 3.3 and Figure 3.4, where the languages and learning abilities of individual agents are plotted for both the small-p and large-p phases. Also notable is the correlation between agents with the same languages and learning abilities in Figures 3.3b and 3.4b. This is likely a side effect of the word inheritance mechanic, further discussed in Section 4.1.3.

The Baldwin Effect is exhibited in the fact that the change of communication intensity, in effect a cultural change in the learned behaviour (language faculty) of the agents, leads to the natural selection of a high learning ability universally, a biological change. In the small-p phase, the clusters of agents using the same language are small and having a large learning ability is not advantageous as communications are more often than not unsuccessful. Once the communication intensity increases, the clusters start to grow and a new niche presents itself which directs the evolution of subsequent generations. Agents with a high learning ability are now at an advantage as they are more likely to be successful in communicating and thus to increase their knowledge, which in turn increases their fitness.

Figure 3.5 demonstrates another interesting effect of changing p. p is increased continuously from 0.1 to 0.5 over the course of the simulation, instead of the abrupt increase in the previous experiment. What follows is a multi-stepped behaviour of plateaus separated by rapid transitions, as opposed to the single sigmoid transition of Figure 3.2. The language and learning ability maps of the various plateaus are plotted in Figure 3.6.

The inverse biolinguistic transition is reported when p is decreased rather increased, which causes a population of agents with an initially high learning ability to naturally select for a lower one, however interestingly with a different threshold value of p than when increasing. This phenomenon was *not* observed in the recreated model and in fact the learning ability continued to increase as p decreased. A comparison of the hysteretic behaviour of the original model and the observed behaviour in the recreated model is shown in Figure 3.7.
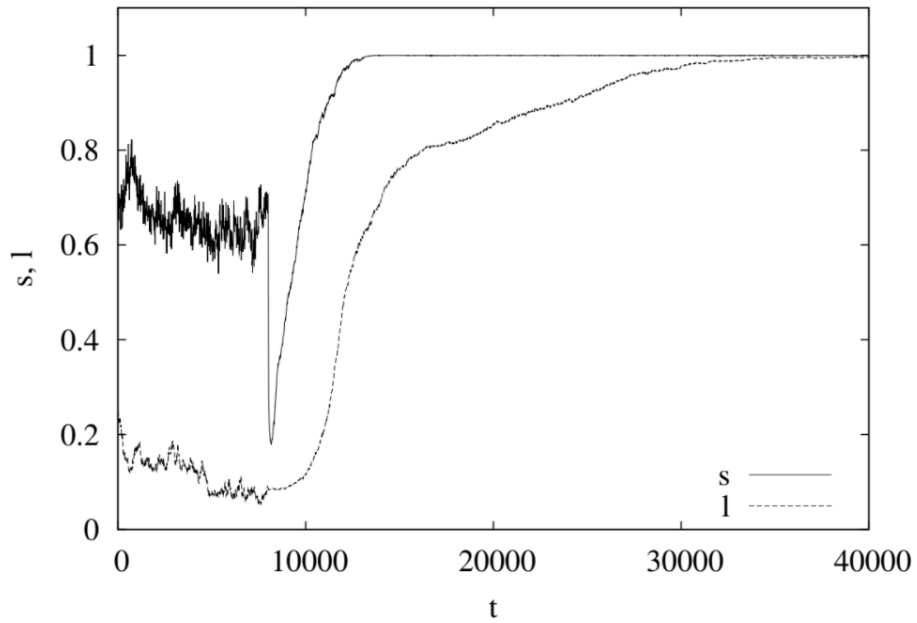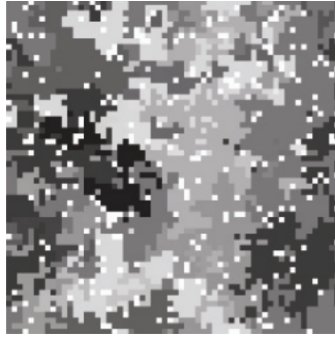
(a) Original
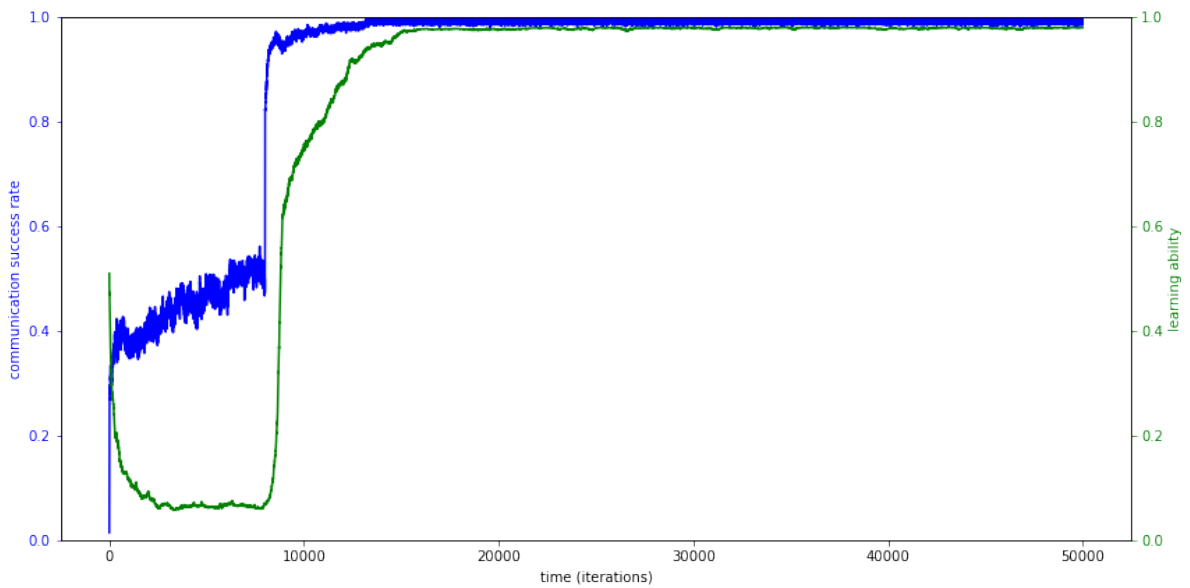
(b) Recreated

Figure 3.2: Experimental results demonstrating the natural selection of a high learning ability after an increase in communication intensity. The communication probability is initially set to $p = 0.1$ and increased to $p = 0.98$ at $t = 8000$. The remaining parameters for these experiments are the following: $L = 40$, $N = 1$, $p_{mut} = 0.001$, $A = 0.005$, $B = 5$. Plot of original results reproduced with permission from [7].

(a) Original: small-p phase
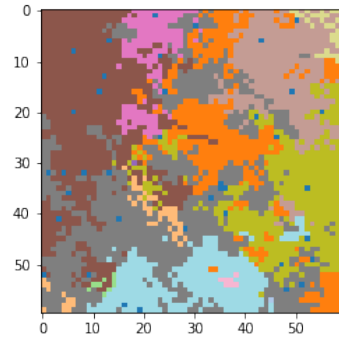
(b) Recreated: small-p phase
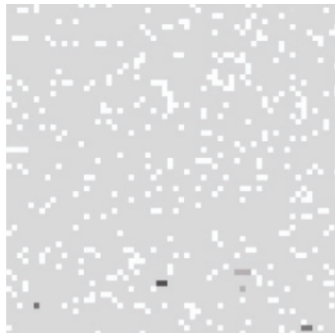
(c) Original: large-p phase
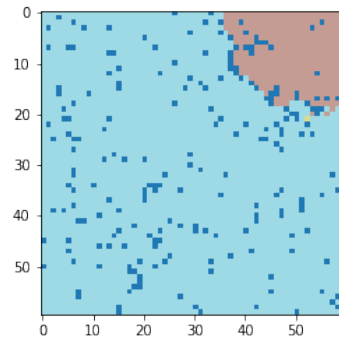
(d) Recreated: large-p phase

Figure 3.3: Illustration of the linguistic component of the "biolinguistic" transition. Figures 3.3a and 3.3c show the language spoken by each agent, with each greyscale level corresponding to a different language. In the small-p phase, the model is in a chaotic state, with many different languages spoken by small clusters of agents, which transitions to a homogeneous state in the large-p phase. Figures 3.3b and 3.3d show the same effect in the recreated model, with each colour representing a different language. The parameters for these experiments are the following: $L = 60$, $N = 1$, $p_{mut} = 0.001$, $A = 0.005$, $B = 5$. Plots of original results reproduced with permission from [6].

(a) Original: small-p phase
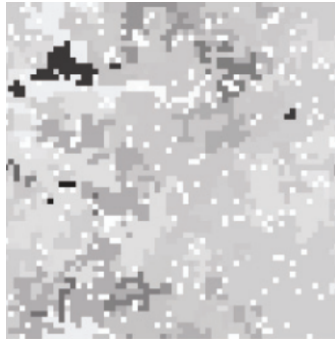

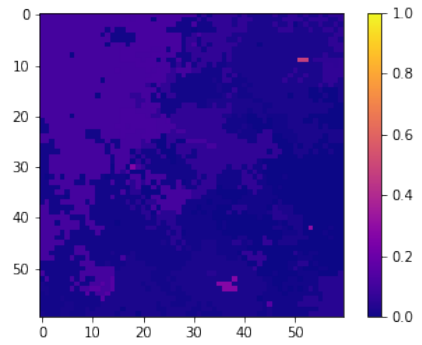(b) Recreated: small-p phase

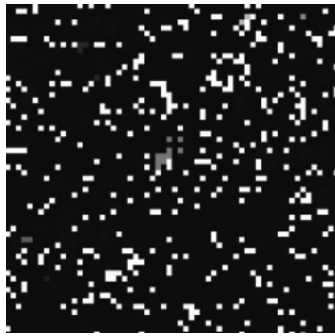(c) Original: large-p phase
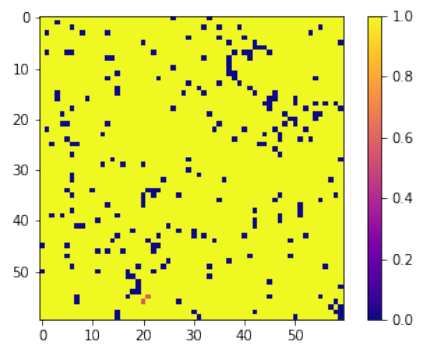
(d) Recreated: large-p phase

Figure 3.4: Illustration of the biological component of the "biolinguistic" transition. Figures 3.4a and 3.4c show the learning ability of each agent, with each greyscale level corresponding to a different learning ability, going from low (white) to high (black). The small-p phase is characterised by a low learning ability across the board except for some small clusters, and once again the large-p phase is more homogeneous, with most agents having a high learning ability. Figures 3.4b and 3.4d show the same effect in the recreated model. The remaining parameters for these experiments are the following: $L = 60$, $N = 1$, $p_{mut} = 0.001$, $A = 0.005$, $B = 5$. Plots of original results reproduced with permission from [6].
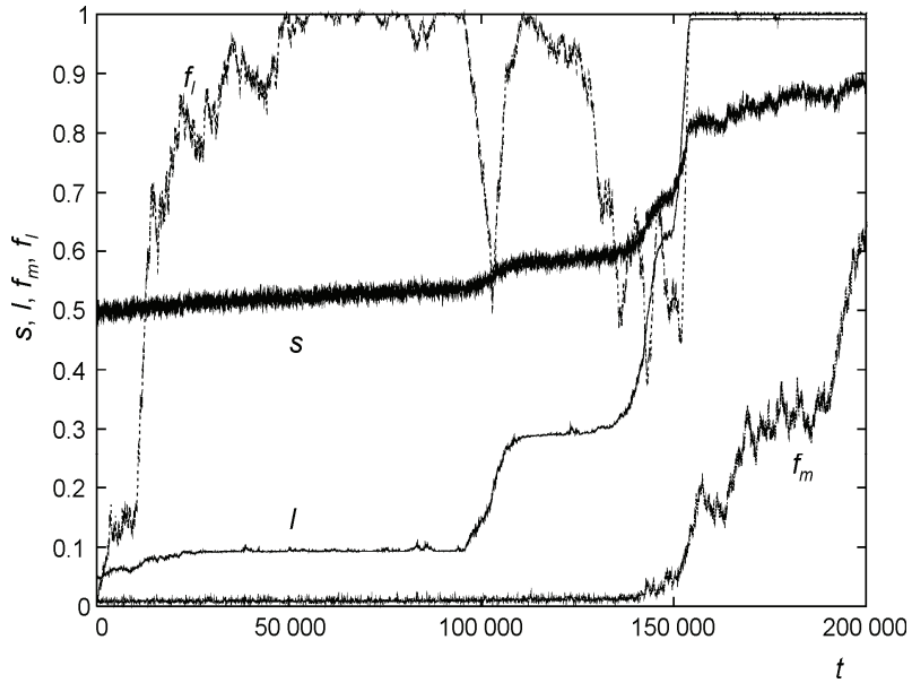
(a) Original
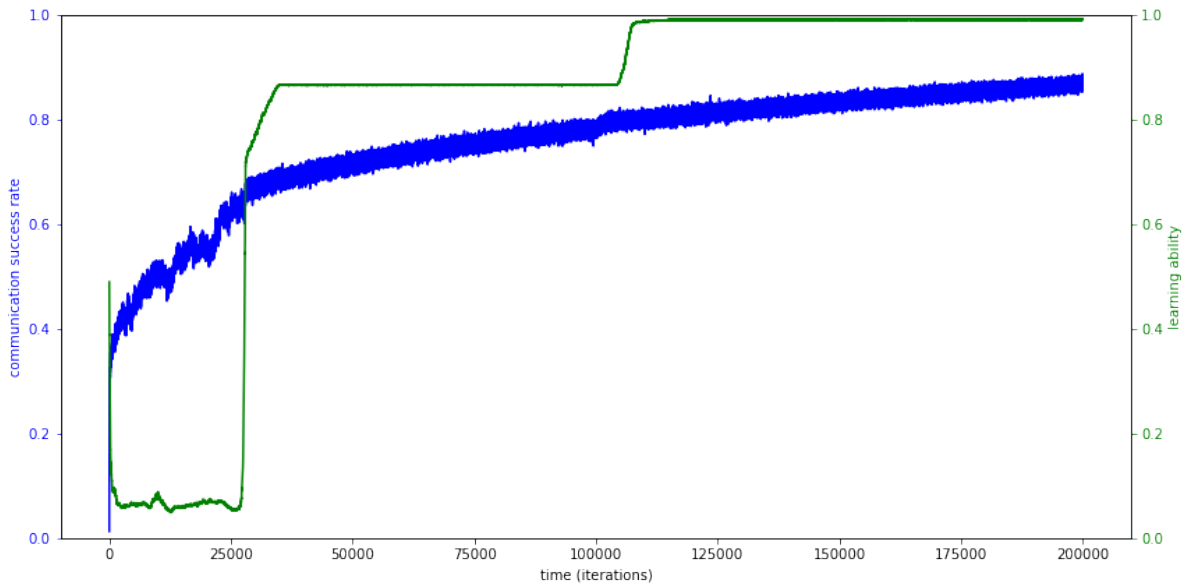
(b) Recreated

Figure 3.5: Experimental results demonstrating the multi-stage transition from low to high learning ability as communication intensity increases continuously from $p = 0.1$ at $t = 0$ to $p = 0.5$ at $t = 200000$. The remaining parameters for these experiments are the following: $L = 60$, $N = 1$, $p_{mut} = 0.001$, $A = 0.005$, $B = 5$. Plot of original results reproduced with permission from [6].
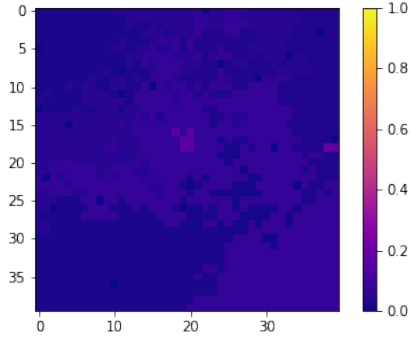
(a) Learning ability: small-p plateau
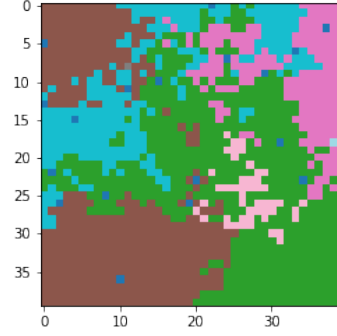
(b) Languages: small-p plateau
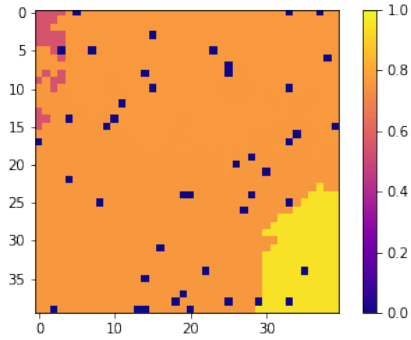
(c) Learning ability: middle-p plateau
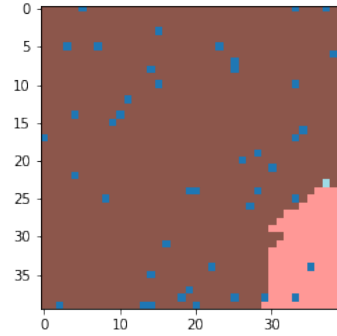
(d) Languages: middle-p plateau
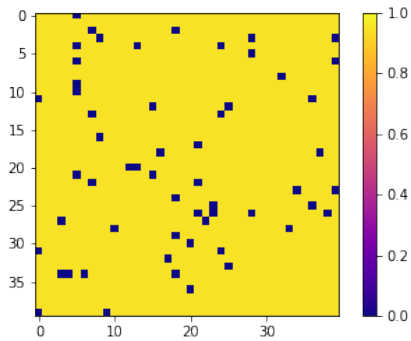
(e) Learning ability: large-p plateau
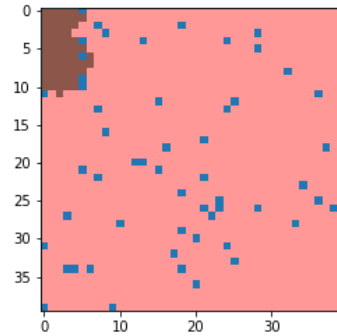
(f) Languages: large-p plateau

Figure 3.6: Learning ability and language plots for each plateau in the continuous p-increasing scenario. The parameters for this experiment are the following: $L = 60$, $N = 1$, $p_{mut} = 0.001$, $A = 0.005$, $B = 5$.
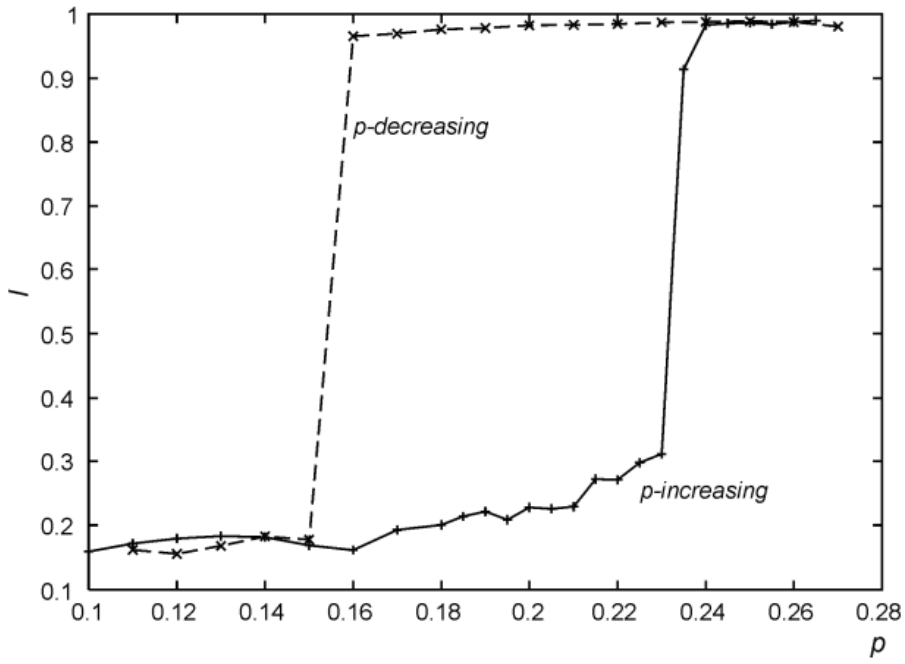
(a) Original
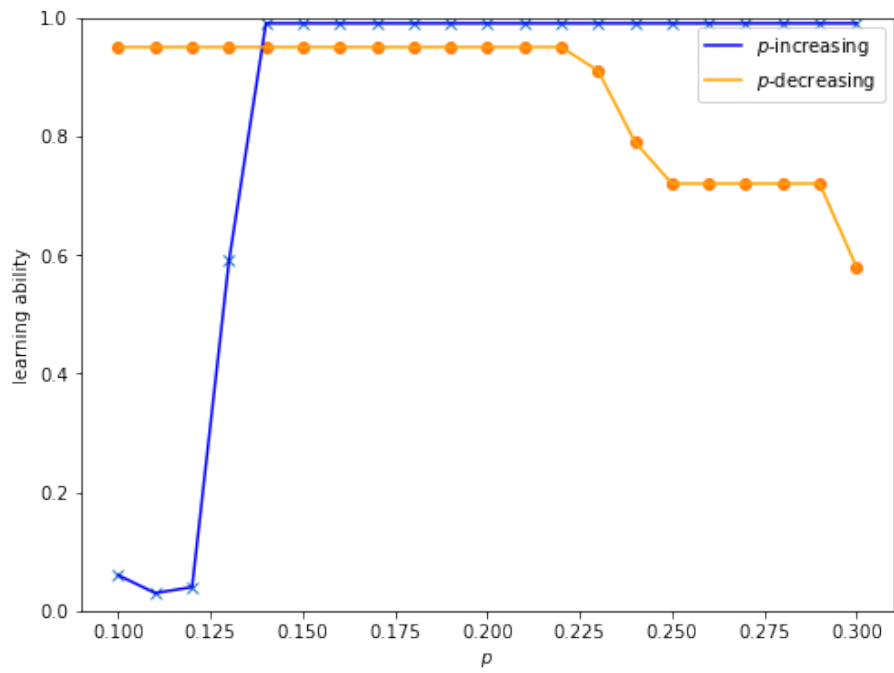
(b) Recreated

Figure 3.7: Comparison of average learning ability for different values of p when increasing and decreasing p. Plot of original results reproduced with permission from [6].

4 Modifications

This chapter discusses some modifications that were made to the Lipowska ENG and experiments carried out on the modified variants. First a description of the modifications is laid out, followed by the results obtained from these modified versions.

4.1 Description

In implementing the original Lipowska ENG, it seems some of the design choices were made arbitrarily. In order to better understand the effect of these choices on the behaviour of the model, some modifications are proposed and discussed in the following sections.

4.1.1 Survival Probability

Referring back to the calculation of survival probability using Equation 3.2.2, one choice stands out; the survival probability of an individual agent refers to a property of the entire population, namely W, the average knowledge over all agents. It seems an interesting choice that the fitness of one agent could be affected by a property of another. Instead shouldn't the first agent's fitness be independent of other agents?

The proposed modification to Equation 3.2.2 is shown in Equation 4.1.1. Here s refers to the individual success rate of an agent, i.e. the ratio of the number of successful communications made by that agent to the total number of communication attempts initiated by that agent. The parameter C is a control parameter introduced to replace B in Equation 3.2.2. The value of C is chosen so as to make the fitness landscape as similar to that of the original model as possible. For a comparison of the two fitness landscapes plotted as surfaces see Figure 4.1.

$$p_{surv} = \exp(-Aa_k) \times s^C \tag{4.1.1}$$
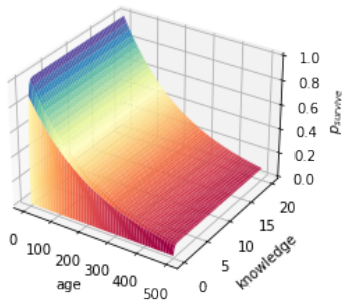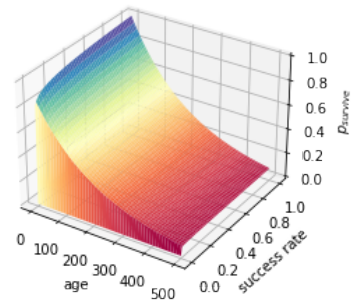
(a) Original | (b) Modified

Figure 4.1: Plotted here are the surfaces of Equations 3.2.2 and 4.1.1. Age and knowledge are plotted on the horizontal axes of Figure 4.1a, and age and success rate on the horizontal axes of Figure 4.1b. The vertical axes of both figures show the fitness, with the maximum (indicated with blue) at low age and high knowledge or low age and high success rate. The values of A and B are 0.005 and 5 respectively, and the value of C is set to 0.005, such that the surface is similar to that of the original function.

4.1.2 Word Acquisition

Agents can acquire new words in three ways:

1. Inventing a new word.

2. Learning a new word as listener after failed communication attempt.

3. Inheriting highest weighted word from parent.

As shown in Algorithms 2 and 4, the initial lexical weight given to a newly encountered word is 1, equivalent to the maximum possible learning ability. In the first case this makes sense, as an agent should know well how to use a word that they themselves have invented. However, in such a simulation where word learning is such an important mechanic, why is an agent's learning ability only used when they are hearing or otherwise using a word they already know, and not when learning it for the first time? The modified version of Algorithm 2 substitutes line 3 such that $w_i = l_{speaker}$ and line 12 such that $w_i = l_{listener}$ and the same applies for lines 5 and 10 of Algorithm 4 with $w_i = l_k$.

4.1.3 Word Inheritance

A similar point can be made about the inheritance of the parent's highest weighted word by the child. By simply implanting the word into the child's lexicon at birth, learning is not allowed to take place. Could newly born agents instead rely on the fact that they will communicate with other agents in the neighbouring tiles (including the parent), who are themselves likely to speak the same language as the parent because of their proximity? Theoretically, the

25

behaviour of agents within language clusters would likely remain similar to that observed in the Lipowska ENG, but there might be interesting observations to be made at the regions where multiple language clusters meet. Some child agents may prefer to use the language of other agents around them over their parents' language.

4.1.4 Non-exclusive Updates

Communication updates and population updates are exclusive in the Lipowska ENG, meaning that as p is increased and therefore $1 - p$ decreases, agents are much less likely to die simply because they are less likely to face a population update. Agents should not be less likely to die simply by the action of communicating more frequently instead of by any benefit that that more frequent communication may bring. This modification discards the parameter p and instead has agents perform *both* a communication and a population update at each step.

4.2 Results

Here the results of the various modifications are presented, with a more detailed discussion left for Chapter 5. The parameters for all these simulations unless otherwise stated are the following: $L = 40$, $N = 1$, $p_{mut} = 0.001$, $A = 0.005$, $B = 5$, $p_{start} = 0.1$, $p_{end} = 0.5$. p_{start} and p_{end} are the values of p at the first and last iteration respectively, with the value of p for any given iteration in between being a linear interpolation between these two values.

4.2.1 Survival Probability

In this case (see Figure 4.2), there is no biological transition to high learning ability despite the continuous increase of p over time. Convergence on a single language still occurs after 10-20k iterations, indicating that although agents have a low average learning ability, they are still able to communicate effectively, further evidenced by the high success rate.

4.2.2 Word Acquisition

As previously mentioned, there are three cases in which agents acquire new words. Figure 4.3 shows the results for a number of different scenarios. In two cases, the convergence on a high learning ability and a single language happens early on, even before p has increased significantly. The third case exhibits the expected multi-stepped behaviour, but with one notable difference; the maximum average learning ability is slightly lower than previously.
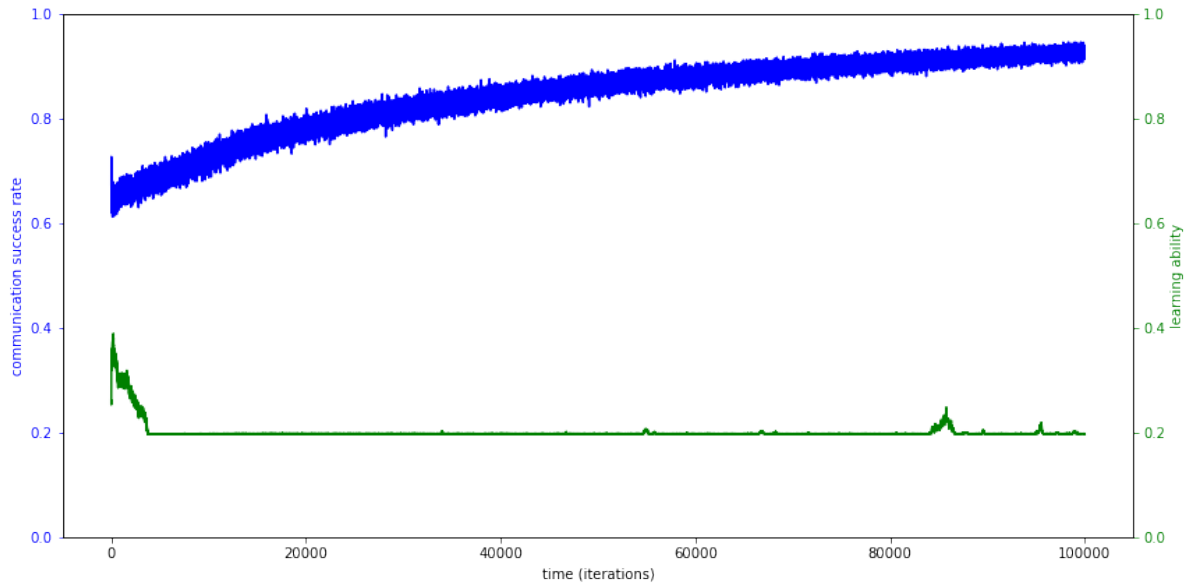
Figure 4.2: Results of simulation with new fitness function. Convergence on a learning ability of 0.2 occurs after 5k iterations and does not increase despite the continuous increase of *p* from 0.1 to 0.5.

4.2.3 Word Inheritance

Word inheritance seems to be very important to the original ENG's behaviour. Without it the simulation stays in a multilingual state as new languages are constantly being invented by new offspring and no large language clusters can form (see Figure 4.5). A very low average learning ability is also observed (see Figure 4.4).

An alternative would be to allow agents invent words only after a certain age before which they can only learn, further discussed in 6.1.

4.2.4 Non-exclusive Updates

Similarly to the alternative word learning scenarios, the simulation converges on a high learning ability early on (see Figure 4.6). Furthermore, the dominant language has also been established by 8000 iterations, which is slightly earlier than in other scenarios.

A more interesting experiment might be to genetically determine an agent's propensity for communicating, further discussed in Section 6.1.
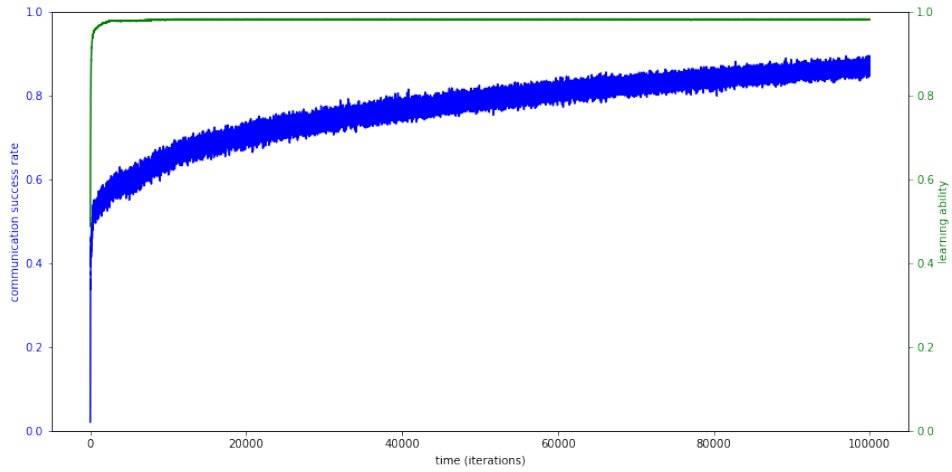
(a)
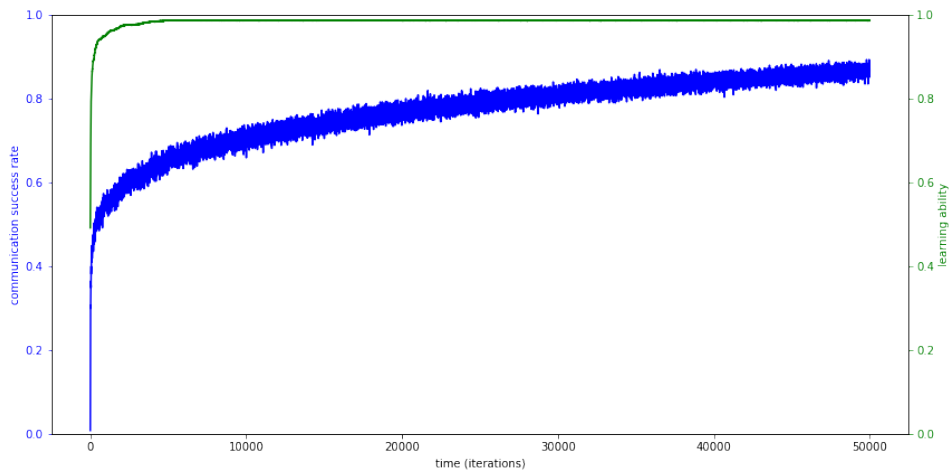
(b)
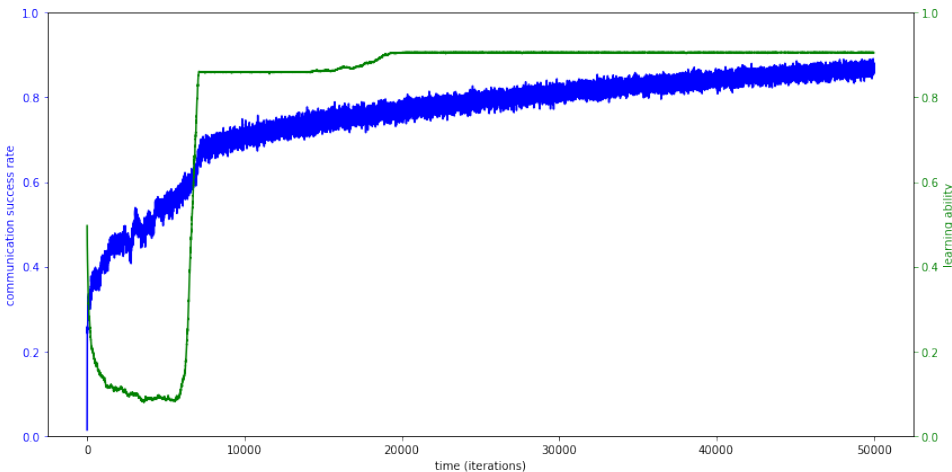
(c)

Figure 4.3: Figure 4.3a shows the results when all three methods of word acquisition set the word weight to the agent's learning ability. Figure 4.3b shows the results when word invention has a unit weight and word learning and inheritance use the agent's learning ability. Figure 4.3c shows the results when word invention and inheritance use a unit weight and word learning uses the agent's learning ability.
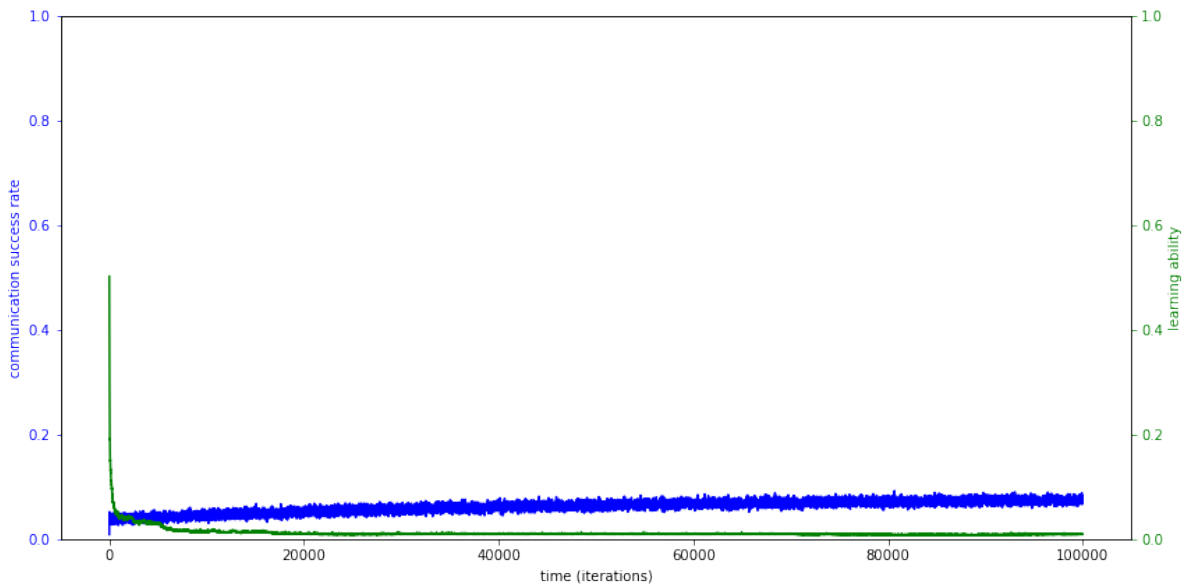
Figure 4.4: Results of simulation with no word inheritance. The learning ability very quickly converges on a value close to 0 and does not change throughout the duration of the simulation.
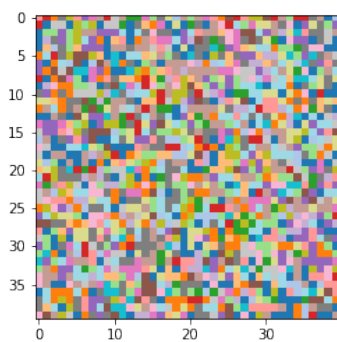
(a) $t = 50000$
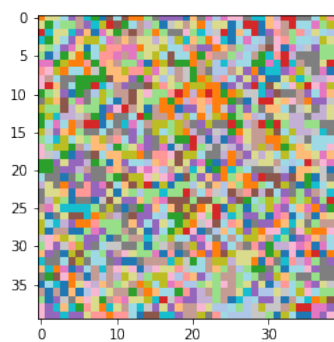
(b) $t = 100000$

Figure 4.5: Language plots for no word inheritance experiments. No language clusters are allowed to form as new words are constantly being invented.
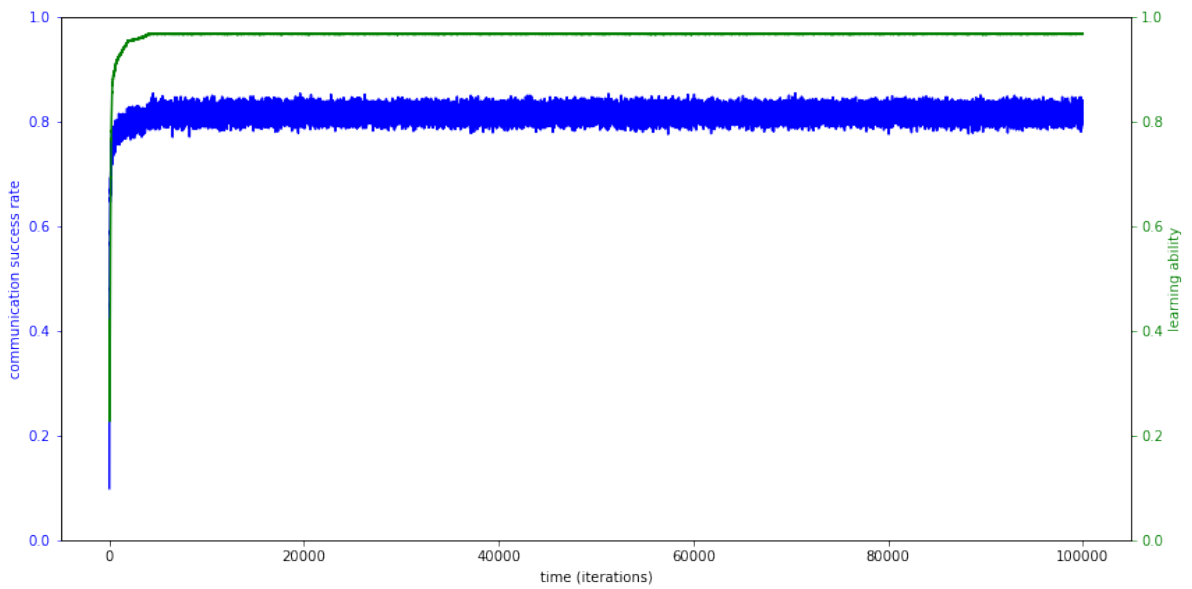
Figure 4.6: Results of simulation where every agent faces a communication update and population update at every individual update.

5 Discussion

This chapter contains a general discussion of the results of the models from the preceding Chapters.

5.1 Lipowska ENG

As stated in Section 1.1, one of the goals of this project is to implement and verify the Lipowska ENG. Contrary to what is reported by Lekvam et al. [25], the results of the Lipowska ENG are reproducible.

For those experiments described in the original papers [6, 7], the results obtained by the recreated implementation were mostly very similar and the same observations could be made. The biolinguistic transition dependent on the value of p was observed for a number of different population sizes ranging from $30 \times 30 = 900$ to $60 \times 60 = 3600$. The transition can also be seen as a demonstration of the Baldwin Effect by niche construction. The notable exception to this is in the case of p-decreasing experiments, where the hysteretic behaviour of the model was not observed.

In summary, two conclusions can be drawn from these results; firstly, the implementation of the Lipowska ENG described in Chapter 3 is correct and the pseudocode descriptions therein can serve as the basis for continued research. Secondly, the results reported by Lipowska in [6, 7] are verified to be correct, except for p-decreasing scenarios.

5.2 Modifications

Lipowska claims that the ENG is "to some extent robust with respect to some modifications of its rules. For example, qualitatively the same behavior is observed for modified parameters [A] and [B], different form of the survival probability p_{surv} (provided it is a decreasing function of [a] and an increasing function of $\sum_j w_j$), or different breeding and/or mutation rules." [7] Considering this, at least some of the modifications from Chapter 4 should produce the same or similar results as the original.

In reality, the model was found to be very sensitive to different values of A and B, with most values resulting in either immediate extinction of all agents or no biolinguistic transition at all. Consequently all of the results presented in this dissertation use the values $A = 0.005$ and $B = 5$, the same values suggested by Lipowska. The new fitness function outlined by Equation 4.1.1, although it is a decreasing function of age a and an increasing function of knowledge $\sum_j w_j$, resulted in completely different behaviour.

Among the modifications to the breeding and mutation rules are those introduced in the word inheritance and word learning sections. The removal of word inheritance has a great impact on the model's behaviour, where agents are essentially not able to establish a language to use. Instead many languages exist simultaneously and are continuously invented and forgotten as agents are born and die. As a result of the low proportion of successful communications, there is no advantage to having a high learning ability as agents are not likely to be able to increase their knowledge and thus their fitness. The modifications to word acquisition also present a departure from the model's expected behaviour. When all three methods of word acquisition are made to set the new word's weight to the agent's learning ability, it is immediately useful for the agents to have a high learning ability, and the same applies for when only word learning and word inheritance use the agent's learning ability, but word invention uses a unit weight. The third case, when the learning ability is only used in word learning, results in the multi-stepped biolinguistic transition that is characteristic of the continuous increase of p. This again suggests that the word inheritance mechanic is critical.

The final modification allows agents to both communicate and carry out the evolutionary processes of breeding and dying in the same update, in some ways analogous to setting $p = 1$. This means that large language clusters can form much faster than otherwise and having a high learning ability is immediately beneficial.

Seemingly the ENG is not as robust as is claimed by Lipowska, but further analysis would help to clarify this.

6 Conclusion

This chapter presents the closing remarks of this dissertation, with reference back to the goal and research questions stated in the beginning.

As stated in Section 1.1, the goal of dissertation is to investigate ways in which agent-based models can be used to simulate the emergence of language, looking specifically at the Lipowska ENG [6, 7]. The goal was further split into the following research questions:

1. Can the results of the Lipowska ENG be verified?

2. How robust is the Lipowska ENG to modifications?

The answers to those research questions are the following:

1. As concluded by the discussion in Section 5.1, the Lipowska ENG is verifiable. The results produced from recreating the original experiments were very close to those reported. This enabled the investigation of the second research question.

2. The Lipowska ENG is not robust to changes of the parameters A and B, nor is it robust to modifications of the fitness function as far as the one suggested in Section 4.1.1. Furthermore it is not robust to changes in the transmission of language between parent and child agents. The model is however robust to changes in population size.

In this respect the original goal was achieved and this project can be considered a success.

6.1 Future Work

A great effort was made to clarify some of the ambiguities of the original articles and to provide pseudocode descriptions of the ENG, such that others could use this model as the basis for further research. As such, this section will describe some suggestions for building on the work carried out during the course of this dissertation.

The word inheritance mechanic was found to be central to the behaviour observed in the Lipowska ENG, however it still leaves questions about why agents inherit language rather than learning it. As opposed to the modification in Section 4.1.3, future models could introduce a

33

period of childhood or infancy, during which agents can only learn language. This could allow language clusters to grow sufficiently large, before agents start inventing languages of their own. This could also be combined with an investigation into the Critical Period Hypothesis — a hypothesis claiming that there is a time window in childhood whereby language acquisition is associated with lower cognitive effort, after which it becomes much more difficult [26] — giving agents a boosted learning ability during childhood, which then fades as they age.

Another possibility could be a continuation of the non-exclusive update variant. What if, instead of every agent communicating at the same rate at every iteration, agents were given an individual genetically determined communication probability? Would the selection pressures of the environment favour highly social and linguistically skilled agents or vice versa?

Bibliography

[1] Maulucioni. Distribution of y-dna haplogroup r1a, 2021. URL `https://commons.wikimedia.org/wiki/File:Mapa_de_R1a.png`. (Accessed on 23/08/2021).

[2] R. Bouckaert, P. Lemey, M. Dunn, S. J. Greenhill, A. V. Alekseyenko, A. J. Drummond, R. D. Gray, M. A. Suchard, and Q. D. Atkinson. Mapping the origins and expansion of the indo-european language family. *Science*, 337(6097):957–960, August 2012. doi: 10.1126/science.1219669.

[3] R. Bouckaert, P. Lemey, M. Dunn, S. J. Greenhill, A. V. Alekseyenko, A. J. Drummond, R. D. Gray, M. A. Suchard, and Q. D. Atkinson. Inferred geographic origin of the indo-european language family, 2012. URL `https://science.sciencemag.org/content/337/6097/957/tab-figures-data`. (Accessed on 23/08/2021).

[4] Geoffrey Hinton and Steven Nowlan. How learning can guide evolution. *Complex Systems*, 1(3):495–502, 1987. doi: N/A.

[5] Albano Beja-Pereira, Gordon Luikart, Phillip R England, Daniel G Bradley, Oliver C Jann, Giorgio Bertorelle, Andrew T Chamberlain, Telmo P Nunes, Stoitcho Metodiev, Nuno Ferrand, and Georg Erhardt. Gene-culture coevolution between cattle milk protein genes and human lactase genes. *Nature Genetics*, 35(4):311–313, November 2003. doi: 10.1038/ng1263.

[6] Dorota Lipowska. Naming game and computational modelling of language evolution. 2011. doi: 10.12921/CMST.2011.17.01.41-51.

[7] Dorota Lipowska and Adam Lipwoski. Bio-linguistic transition and baldwin effect in an evolutionary naming-game model. *International Journal of Modern Physics C*, 19(03): 399–407, March 2008. doi: 10.1142/s0129183108012248.

[8] Iain Mathieson, Iosif Lazaridis, Nadin Rohland, Swapan Mallick, Nick Patterson, Songul Alpaslan Roodenberg, Eadaoin Harney, Kristin Stewardson, Daniel Fernandes, Mario Novak, Kendra Sirak, Cristina Gamba, Eppie R. Jones, Bastien Llamas, Stanislav Dryomov, Joseph Pickrell, Juan Luis Arsuaga, Jose Maria Bermudez de Castro, Eudald

Carbonell, Fokke Gerritsen, Aleksandr Khokhlov, Pavel Kuznetsov, Marina Lozano, Harald Meller, Oleg Mochalov, Vayacheslav Moiseyev, Manuel A. Rojo Guerra, Jacob Roodenberg, Josep Maria Verges, Johannes Krause, Alan Cooper, Kurt W. Alt, Dorcas Brown, David Anthony, Carles Lalueza-Fox, Wolfgang Haak, Ron Pinhasi, and David Reich. Eight thousand years of natural selection in europe. March 2015. doi: 10.1101/016477.

[9] Morten E. Allentoft, Martin Sikora, Karl-Göran Sjögren, Simon Rasmussen, Morten Rasmussen, Jesper Stenderup, Peter B. Damgaard, Hannes Schroeder, Torbjörn Ahlström, Lasse Vinner, Anna-Sapfo Malaspinas, Ashot Margaryan, Tom Higham, David Chivall, Niels Lynnerup, Lise Harvig, Justyna Baron, Philippe Della Casa, Paweł Dąbrowski, Paul R. Duffy, Alexander V. Ebel, Andrey Epimakhov, Karin Frei, Mirosław Furmanek, Tomasz Gralak, Andrey Gromov, Stanisław Gronkiewicz, Gisela Grupe, Tamás Hajdu, Radosław Jarysz, Valeri Khartanovich, Alexandr Khokhlov, Viktória Kiss, Jan Kolář, Aivar Kriiska, Irena Lasak, Cristina Longhi, George McGlynn, Algimantas Merkevicius, Inga Merkyte, Mait Metspalu, Ruzan Mkrtchyan, Vyacheslav Moiseyev, László Paja, György Pálfi, Dalia Pokutta, Łukasz Pospieszny, T. Douglas Price, Lehti Saag, Mikhail Sablin, Natalia Shishlina, Václav Smrčka, Vasilii I. Soenov, Vajk Szeverényi, Gusztáv Tóth, Synaru V. Trifanova, Liivi Varul, Magdolna Vicze, Levon Yepiskoposyan, Vladislav Zhitenev, Ludovic Orlando, Thomas Sicheritz-Pontén, Søren Brunak, Rasmus Nielsen, Kristian Kristiansen, and Eske Willerslev. Population genomics of bronze age eurasia. *Nature*, 522(7555):167–172, June 2015. doi: 10.1038/nature14507.

[10] Charles Perreault and Sarah Mathew. Dating the origin of language using phonemic diversity. *PLoS ONE*, 7(4):e35289, April 2012. doi: 10.1371/journal.pone.0035289.

[11] Q. D. Atkinson. Phonemic diversity supports a serial founder effect model of language expansion from africa. *Science*, 332(6027):346–349, April 2011. doi: 10.1126/science.1199295.

[12] Ian Tattersall. Becoming modern homo sapiens. *Evolution: Education and Outreach*, 2 (4):584–589, September 2009. doi: 10.1007/s12052-009-0164-x.

[13] When did we become fully human? what fossils and dna tell us about the evolution of modern intelligence. https://theconversation.com/when-did-we-become-fully-human-what-fossils-and-dna-tell-us-about-the-evolution- (Accessed on 11/08/2021).

[14] Andrea Parravicini and Telmo Pievani. Continuity and discontinuity in human language evolution: Putting an old-fashioned debate in its historical perspective. *Topoi*, 37(2): 279–287, October 2016. doi: 10.1007/s11245-016-9431-y.

[15] Noam Chomsky. *Syntactic Structures*. Mouton and Co., The Hague, 1957. ISBN 3112316002.

[16] Noam Chomsky. *Aspects of the Theory of Syntax*. The MIT Press, Cambridge, 1965. ISBN 0262530074.

[17] Michael Dobrovolsky. *Contemporary Linguistics. An introduction*. Longman, London and New York, 3rd. edition, 1996. ISBN 9780582246911.

[18] Noam Chomsky. *Knowledge of language : its nature, origin, and use*. Praeger, New York, 1986. ISBN 0275900258.

[19] Steven Pinker and Paul Bloom. Natural language and natural selection. *Behavioral and Brain Sciences*, 13(4):707–727, December 1990. doi: 10.1017/s0140525x00081061.

[20] Steven Pinker. *The language instinct : how the mind creates language*. HarperPerennial, New York, 1995. ISBN 0060976519.

[21] James Mark Baldwin. A new factor in evolution. *American Naturalist*, 30:441–451, 1896.

[22] F. John Odling-Smee, Kevin Laland, and Marcus Feldman. *Niche construction : the neglected process in evolution*. Princeton University Press, Princeton, 2003. ISBN 1400847265.

[23] Luc Steels. A self-organizing spatial vocabulary. *Artificial Life*, 2(3):319–332, April 1995. doi: 10.1162/artl.1995.2.3.319.

[24] Luc Steels and Martin Loetzsch. The grounded naming game. pages 41–59, 2012. doi: 10.1075/ais.3.04ste.

[25] Torvald Lekvam, Björn Gambäck, and Lars Bungum. Agent-based modeling of language evolution. 2014. doi: 10.3115/v1/w14-0510.

[26] Catherine E. Snow and Marian Hoefnagel-Hohle. The critical period for language acquisition: Evidence from second language learning. *Child Development*, 49(4):1114, December 1978. doi: 10.2307/1128751.