Agent-based modeling of language evolution

Torvald Lekvam   **Björn Gambäck**   **Lars Bungum**
Department of Computer and Information Science, Sem Sælands vei 7–9
Norwegian University of Science and Technology, 7491 Trondheim, Norway
`torvald@lekvam.no`   `{gamback,larsbun}@idi.ntnu.no`

Abstract

Agent-based models of language evolution have received a lot of attention in the last two decades. Researchers wish to understand the origin of language, and aim to compensate for the lacking empirical evidence by utilizing methods from computer science and artificial life. The paper looks at the main theories of language evolution: biological evolution, learning, and cultural evolution. In particular, the Baldwin effect in a naming game model is elaborated on by describing a set of experimental simulations. This is on-going work and ideas for further investigating the social aspects of language evolution are also discussed.

1 Introduction

What is language? It is interesting how we can take a train of thought and transfer this into another 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. While animals also use language, even with a degree of syntax (Kako, 1999), spoken human language exhibits a vastly more complex structure and spacious variation.

To understand how language works — how it is used, its origin and fundamentals — our best information sources are the languages alive (and some extinct but documented ones). Depending on definition, there are 6,000–8,000 languages worldwide today, showing extensive diversity of syntax, semantics, phonetics and morphology (Evans and Levinson, 2009). Still, these represent perhaps only 2% of all languages that have ever existed (Pagel, 2000). As this is a rather small window, we want to look back in time. But there is a problem in linguistic history: our reconstruction techniques can only take us back some 6,000 to 7,000 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 (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 Africa (Klein, 2008; Tattersall, 2010).

The rest of this introduction covers some key aspects of language evolution. Section 2 then focuses on computational models within the field, while Section 3 describes a specific naming game model. Finally, Section 4 discusses the results and some ideas for future work.

1.1 Theories of origin: the biological aspect

There are two main ideas in biological evolution as to why humans developed the capacity to communicate through speech. The first states that language (or more precisely the ability to bear the full structure of language) came as an *epiphenomenon*, a by-product (spandrel) of an unrelated mutation. This theory assumes 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. Thus there should exist an innate capacity in the form of a *universal grammar* (Chomsky, 1986), which can hold a finite number of rules enabling us to carry any kind of language.

According to the second idea, language emerged in a strictly adaptational process (Pinker and Bloom, 1990). That is, that language evolution can be explained by natural selection, in the same way as the evolution of other complex traits like echolocation in bats or stereopsis in monkeys. Both ideas — innate capacity vs natural selection — have supporters, as well as standpoints that hold both aspects as important, but at different levels (Deacon, 2010; Christiansen and Kirby, 2003).

49

1.2 Theories of origin: the cultural aspect

Biology aside, the forces behind the emergence of human language are not strictly genetic (and do not operate only on a phylogenetic time scale). Kirby (2002) argues that, in addition to biological evolution, there are two more complex adaptive (dynamical) systems influencing natural language; namely *cultural evolution* (on the glossogenetic time scale) and *learning* (which operates on a individual level, on the ontogenetic time scale).

In addition, there is the interesting Darwinian idea that cultural learning can guide biological evolution, a process known as *the Baldwin effect* (Baldwin, 1896; Simpson, 1953). This theory argues that culturally learned traits (e.g., a universal understanding of grammar or a defense mechanism against a predator) can assimilate into the genetic makeup of a species. 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 a learned trait would become innate. On the other hand, there is a cost of genetic assimilation as it can prohibit plasticity in future generations and make individuals less adaptive to unstable environments.

There has been much debate recently whether language is a result of the Baldwin effect or not (Evans and Levinson, 2009; Chater et al., 2009; Baronchelli et al., 2012, e.g.), but questions, hypotheses, and simulations fly in both directions.

2 Language evolution and computation

Since the 90s, there has been much work on simulation of language evolution in bottom-up systems with populations of autonomous agents. The field is highly influenced by the work of Steels and Kirby, respectively, and has been summarized and reviewed both by themselves and others (Steels, 2011; Kirby, 2002; Gong and Shuai, 2013, e.g.).

Computational research in this field is limited to modeling very simplified features of human language in isolation, such as strategies for naming colors (Bleys and Steels, 2011; Puglisi et al., 2008), different aspects of morphology (Dale and Lupyan, 2012), and similar. 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.

A language game simulation (Steels, 1995) is a model where artificial agents interact with each other in turn 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 not given any prior knowledge of how language should look like or how consensus should unfold. No centralized anchors are involved: a simulation is all self-organized.

Agents are chosen (mostly at random) as hearer and speaker, and made to exchange an utterance about a certain arbitrary concept or meaning in their environment. If the agents use the same language (i.e., 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 a failure. If an agent wants to express some concept without having any utterances 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 fairly direct connection between utterance and actual meaning in language game models (emotions and social situations do not bias how language is interpreted).

A simple language game normally is characterized by many synonyms spawning among the agents. As agents commence spreading their own utterances around, high-weighted words start to be preferred. Consensus is reached when all agents know the highest weighted word for each concept. Commonly, the agents aim to reach a single coherent language, but the emergence of multilingualism has also been simulated (Lipowska, 2011; Roberts, 2012). Cultural evolution can be captured by horizontal communication between individuals in the same generation or vertical communication from adults to children. The latter typically lets the agents breed, age and die, with the iterated learning model (Smith et al., 2003) being popular.

A variety of language games exist, from simple naming games, where the agents' only topic concerns one specific object (Lipowska, 2011), to more cognitive grounding games (Steels and Loetzsch, 2012). There have also been studies on some more complex types of interaction, such as spatial games (Spranger, 2013), factual description games (van Trijp, 2012) and action games (Steels and Spranger, 2009), where the agent communication is about objects in a physical environment, about real-world events, and about motoric behaviors, respectively.

3 The Baldwin effect in a naming game

Several researchers have created simulations to investigate the Baldwin effect, starting with Hinton and Nowlan (1987). Cangelosi and Parisi (2002) simulate agents who evolve a simple grammatical language in order to survive in a world filled with edible and poisonous mushrooms. Munroe and Cangelosi (2002) used this model to pursue the Baldwin effect, with partially blind agents initially having to learn features of edible mushrooms, but with the learned abilities getting more and more assimilated into the genome over the generations. Chater et al. (2009) argue that only stable parts of language *may* assimilate into the genetic makeup, while variation within the linguistic environment is too unstable to be a target of natural selection. Watanabe et al. (2008) use a similar model, but in contrast state that genetic assimilation not necessarily requires a stable linguistic environment.

Lipowska (2011) has pursued the Baldwin effect in a simple naming game model with the intention of mixing up a language game in a simulation that incorporates both learning, cultural and biological evolution. The model places a set of agents in a square lattice of a linear size L, where every agent is allowed — by a given probability p — to communicate with a random neighbor.

At each time step, a random agent is chosen and p initially decides whether the agent is allowed to communicate or will face a "population update". Every agent has an internal lexicon of N words with associated weights ($w_j : 1 \leq j \leq N$). Whenever a chosen speaker is to utter a word, the agent selects a word i from its lexicon with the probability $w_i / \sum_{j=1}^{N} w_j$. If the lexicon is empty ($N = 0$), a word is made up. A random neighbor in the lattice is then chosen as the hearer. If both agents know the uttered word, the dialog is deemed a success, and if not, a failure. Upon success, both agents increase the uttered word's weight in their lexica by a learning ability variable. Each agent k is equipped with such a variable l ($0 < l_k < 1$). This learning ability is meant to, in its simplicity, reflect the genetic assimilation.

Instead of engaging in communication, the chosen agent is occasionally updated, by a probability $1 - p$. Agents die or survive with a probability p_s which is given by an equation that takes into account age, knowledge (lexicon weights in respect to the population's average weights), and simulation arguments. If the agent has a high-weighted lexicon and is young of age, and therefore survives at a given time step, the agent is allowed to breed if there are empty spaces in its neighborhood.

All in all, each time step can terminate with eight different scenarios: in addition to the two communication scenarios (success or failure), the scenario where the agent dies, as well as the one where the agents lives but only has non-empty neighbors (so that no change is possible), there are four possibilities for breeding. If the agent breeds, the off-spring either inherit the parent's learning ability or gain a new learning ability, with a probability p_m. With the same mutation probability, the off-spring also either gains a new word or inherits the parent's highest-weight word.

This model was implemented with the aim to reproduce Lipowska's results. She argues that her model is fairly robust to both population size and her given arguments; however, our experiments do not support this: as the Baldwin effect unfold, it does not follow the same abrupt course as in Lipowska's model. This could be due to some assumptions that had to be made, since Lipowska (2011), for instance, presents no details on how age is calculated. We thus assume that every time an agent is allowed to communicate, its age gets incremented. Another possibility could be to increment every agent's age at every time step, so that agents get older even if they do not communicate. Furthermore, the initial values for learnability are not clearly stated. Lipowska uses several different values in her analysis. We have used 0.5, which makes a decrease in learnability a part of the evolutionary search space as well.

Simulations with parameters similar to those used by Lipowska (2011) [*iterations* $= 200,000$, *mutation*$_{chance} = 0.01, L = 25, p = 0.4, l = 0.5$], produce results as in Figure 1, showing the highest weighted word per agent after 50k and 150k time steps, with each agent being a dot in a "heat map"; black dots indicate dead agents (empty space). The number of groups are reduced over time, and their sizes grow, as more agents agree on a lexicon and as favorable mutations spread through the population, (as indicated by agent learnability; Figure 2). Even after 200k iterations, consensus is not reached (which it was in Lipowska's simulation), but the agent population agrees on one word if the simulation is allowed to run further. It is natural to assume that the difference lays in the details of how age is calculated, as noted above.
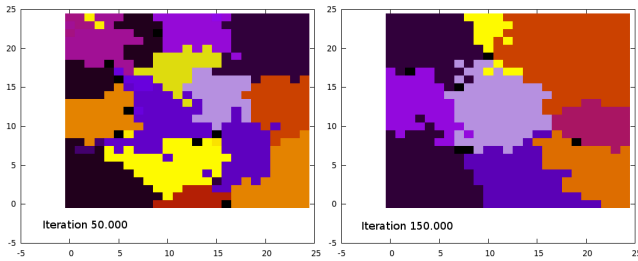
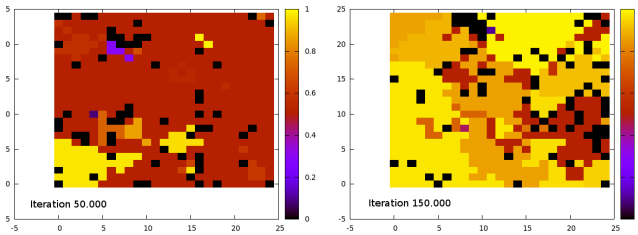Figure 1: Ca 16 different words dominate the population at iteration 50k and nine at iteration 150k.

Figure 2: Mutations favoring learnability at iteration 50k spread substantially by iteration 150k.
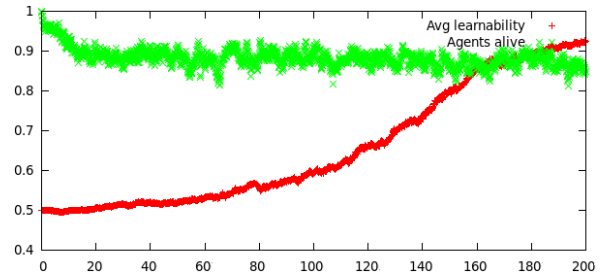

Figure 3: Fraction of agents alive in the lattice and average learnability in the population (s-shaped).
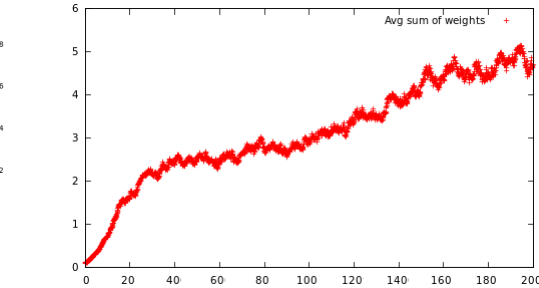
Figure 4: Average sum of weights in agent lexica.

Diverting from Lipowska's parameters and skewing towards faster turnover (higher mutation rate, higher possibility of survival with richer lexicon/higher age, etc.), gives behavior similar to hers, with faster and more abrupt genetic assimilation, as shown Figure 3. The upper line in the figure 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 death and birth are equally tensioned.

Agents with higher learnability tend to live longer, and the lower graph in Figure 3 shows the average learnability in the population. It is roughly sigmoid (S-shaped; cf. Lipowska's experiment) as a result of slow mutation rate in the first phase, followed by a phase with rapid mutation rate (ca 100k–170k) as the learnability also gets inherited, and decreasing rate towards the end when mutations are more likely to ruin agent learnability (when the learning ability l is at its upper limit). As can be seen in Figure 4, the agents rapidly get to a stable weighted lexicon before the Baldwin effect shows itself around time step 100k.

As mentioned, Lipowska's model did not reflect the robustness argued in her paper: for other values of p, the number of empty spots in the population lattice starts to diverge substantially, and for some values all agents simply die. As population sizes vary, the number of iterations must also be adjusted to get similar results. If not, the agents will not reach the same population turn-over as

for smaller population sizes since only one agent may be updated per iteration. Lipowska (2011) compensated with higher mutation rate on simulations with different population sizes; however, these could be two variables somewhat more independent of each other. The model would have been much more stable if it contained aspects of a typical genetic algorithm, where agents are allowed to interact freely within generations. This way, the model could be acting more upon natural selection (and in search of the Baldwin effect), instead of relying on well-chosen parameters to work.

4 Discussion and future work

Language is a complex adaptive system with numerous variables to consider. Thus we must make a number of assumptions when studying language and its evolution, and can only investigate certain aspects at a time through simplifications and abstractions. As this paper has concentrated on the agent-based models of the field, many studies reflecting such other aspects had to be left out.

In addition, there has lately been a lot of work studying small adjustments to the agent-based models, in order to make them more realistic by, for example, having multiple hearers in a language game conversations (Li et al., 2013), different topologies (Lei et al., 2010; Lipowska and Lipowski, 2012), and more heterogeneous populations (Gong et al., 2006).

In general, though, simulations on language evolution tend to have relatively small and fixed sizes (Baronchelli et al., 2006; Vogt, 2007) — and few studies seem to take social dynamics (Gong et al., 2008; Kalampokis et al., 2007) or geography into account (Patriarca and Heinsalu, 2009).

Further work is still needed to make existing models more realistic and to analyze relations between different models (e.g., by combining them). Biological evolution could be studied with more flexible (or plastic) neural networks. Cultural evolution could be investigated under more realistic geographical and demographical influence, while learning could be analyzed even further in light of social dynamics, as different linguistic phenomena unfold. Quillinan (2006) presented a model concerning how a network of social relationships could evolve with language traits. This model could be taken further in combination with existing language games or it could be used to show how language responds to an exposure of continuous change in a complex social network.

Notably, many present models have a rather naïve way of selecting cultural parents, and a genetic algorithm for giving fitness to agents in terms of having (assimilated) the best strategies for learning (e.g., memory efficiency), social conventions (e.g., emotions, popularity), and/or simple or more advanced grammar could be explored.

A particular path we aim to pursue is to study a language game with a simple grammar under social influence (e.g., with populations in different fixed and non-fixed graphs, with multiple hearers), contained within a genetic algorithm. In such a setting, the agents must come up with strategies for spreading and learning new languages, and need to develop fault-tolerant models for speaking with close and distant neighbors. This could be a robust model where a typical language game could be examined, in respect to both biological and cultural evolution, with a more realistic perspective.

Acknowledgments

References

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

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

Andrea Baronchelli, Nick Chater, Romualdo Pastor-Satorras, and Morten H. Christiansen. 2012. The biological origin of linguistic diversity. *PLoS ONE*, 7(10):e48029.

Joris Bleys and Luc Steels. 2011. Linguistic selection of language strategies. In G. Kampis, I. Karsai, and E. Szathmáry, editors, *Advances in Artificial Life. Darwin Meets von Neumann*, volume 2, pages 150–157. Springer.

Angelo Cangelosi and Domenico Parisi. 2002. Computer simulation: A new scientific approach to the study of language evolution. In Angelo Cangelosi and Domenico Parisi, editors, *Simulating the Evolution of Language*, chapter 1, pages 3–28. Springer, London.

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

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

Morten H. Christiansen and Simon Kirby. 2003. Language evolution: consensus and controversies. *TRENDS in Cognitive Sciences*, 7(7):300–307.

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

Terrence W. Deacon. 2010. A role for relaxed selection in the evolution of the language capacity. *Proceedings of the National Academy of Sciences*, 107(Supplement 2):9000–9006.

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

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

Tao Gong, James W. Minett, and William S-Y Wang. 2006. Language origin and the effects of individuals popularity. In *Proceedings of the 2006 IEEE Congress on Evolutionary Computation*, pages 999–1006, Vancouver, British Columbia, Jul. IEEE.

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

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

Edward Kako. 1999. Elements of syntax in the systems of three language-trained animals. *Animal Learning & Behavior*, 27(1):1–14.

Alkiviadis Kalampokis, Kosmas Kosmidis, and Panos Argyrakis. 2007. Evolution of vocabulary on scale-free and random networks. *Physica A: Statistical Mechanics and its Applications*, 379(2):665 – 671.

Simon Kirby. 2002. Natural language from artificial life. *Artificial Life*, 8(2):185–215.

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

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

Bing Li, Guanrong Chen, and Tommy W.S. Chow. 2013. Naming game with multiple hearers. *Communications in Nonlinear Science and Numerical Simulation*, 18(5):1214–1228.

Dorota Lipowska and Adam Lipowski. 2012. Naming game on adaptive weighted networks. *Artificial Life*, 18(3):311–323.

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

Steve Munroe and Angelo Cangelosi. 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.

Mark Pagel. 2000. The history, rate and pattern of world linguistic evolution. In Ch. Knight, J.R. Hurford, and M. Studdert-Kennedy, editors, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, chapter 22, pages 391–416. Cambridge University Press.

Marco Patriarca and Els Heinsalu. 2009. Influence of geography on language competition. *Physica A: Statistical Mechanics and its Applications*, 388(2–3):174–186.

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

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

Justin Quillinan. 2006. Social networks and cultural transmission. Master of Science Thesis, School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh, Scotland, Aug.

Sean Geraint Roberts. 2012. *An evolutionary approach to bilingualism*. Ph.D. thesis, School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh, Scotland, Oct.

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

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

Michael Spranger. 2013. Evolving grounded spatial language strategies. *KI-Künstliche Intelligenz*, 27(2):1–10.

Luc Steels and Martin Loetzsch. 2012. The grounded naming game. In L. Steels, editor, *Experiments in Cultural Language Evolution*, pages 41–59. John Benjamins.

Luc Steels and Michael Spranger. 2009. How experience of the body shapes language about space. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 14–19, Pasadena, California, Jul. IJCAI.

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

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

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

Remi van Trijp. 2012. The evolution of case systems for marking event structure. In L. Steels, editor, *Experiments in Cultural Language Evolution*, pages 169–205. John Benjamins.

Paul Vogt. 2007. Group size effects on the emergence of compositional structures in language. In F. Almeida e Costa, L.M. Rocha, E. Costa, I Harvey, and A. Coutinho, editors, *Advances in Artificial Life: Proceedings of the 9th European Conference (ECAL 2007)*, pages 405–414, Lisbon, Portugal, Sep. Springer.

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