

LANGUAGE AND FRIENDSHIPS: A CO-EVOLUTION MODEL OF SOCIAL AND LINGUISTIC CONVENTIONS

MARTIN BACHWERK & CARL VOGEL

*Computational Linguistics Group
School of Computer Science and Statistics
Trinity College, Dublin 2, Ireland
{bachwerm,vogel}@tcd.ie*

Human social networks are dynamic, yet the majority of computational models of language evolution operate with either fixed or random interaction patterns between agents. This paper presents a model of lexicon formation in which agents can adjust their social ties based on the success of performed interactions. Simulation results show that different levels of ‘friendship’ between agents co-evolve with a system of linguistic conventions, with the observed social network formation processes largely resembling those of modern humans.

1. Introduction

In most computational models of language evolution, agents can be represented as vertices that are connected by predefined, and usually equally weighted edges. In some cases, these agents would be divided into two or more ‘age’ groups, whereby the vertices in some groups, e.g. the ‘child’ agents would not be connected among themselves (see Gong, 2010, for a review of such models). However, the common trait of nearly all such models is that the social ties between agents remain fixed throughout the simulation, i.e. agents can neither become ‘friendlier’ with some agents, nor break off all contact with others. This lack of dynamic social networks has been identified by Gong and Wang (2005) as one of the major issues with the majority of computational models of language evolution to date. In order to address this, Gong and Wang (2005) and Bachwerk and Vogel (2010) have introduced varying inter- and intra-group communication rates into their models of syntax and lexicon formation, respectively. Allowing dynamics in interaction patterns has proven to affect most aspects of the emerging communication system.

However, even the introduction of more sophisticated social structures into models of language evolution is not sufficient as the ties between (groups of) agents in such models remain fixed from the start of a simulation, whereas it has been observed by e.g. Panzarasa, Opsahl, and Carley (2009) that humans constantly re-evaluate their relationships with their interlocutors depending on the obtained feedback and the partner’s reciprocity. In effect, apart from contributing

to the question of the emergence of language *per se*, a computational model with dynamic ties between agents should also shed light on how agents' decisions during individual interactions influence the overall social structure of the population, i.e. how the very first language-induced friendships were made in human evolution. The following section 2 will provide a short overview of the model that was employed for the presented experiment on dynamic social ties in language evolution, the setup of which is described in section 3. We conclude with a discussion of obtained results in section 4 and suggestions for future work in section 5.

2. Model Overview

The experiments presented in this paper have been performed with the help of the Language Evolution Workbench (LEW) (for a more detailed description of the model see Bachwerk & Vogel, 2010). This workbench provides over 30 adjustable parameters and makes as few *a priori* assumptions about the agents' cognitive skills and their awareness of the possibility of a conventionalized communication system as possible. The few cognitive skills that are assumed can be considered as relatively widely accepted (see Christiansen & Kirby, 2003, for a review) as the minimal prerequisites for the emergence of language. These include the ability to observe and individuate events, the ability to engage in a joint attention frame fixed on an occurring event, and the ability to interact by constructing words and utterances from abstract symbols^a and transmitting these to one's interlocutor.^b Finally, the model assumes an infinite space of both meanings and utterances, recursively formed from a finite set of events/entities and symbols respectively.

The simulations of the LEW consist of consecutive rounds of interactions between pairs of agents, during which one of the agents is assigned the intention to comment on an observed event, while a second agent assumes that the topic of the utterance relates in some way to the event and attempts to decode the meaning of the encountered symbols accordingly. The agents' goal is ultimately to agree on a sufficiently broad and reliable lexicon that will enable them to consistently understand each other in the future. Importantly, there is no telepathy built into the model, neither in the form of direct meaning-form transmission, nor in the form of explicit feedback (see Smith, 2005, for further discussion). The lack of a telepathic ability means that the agents need to resort to alternative information sources in order to learn the meaning-form mappings that are being utilized by others. Here we make the assumption that if communication evolved out of the need for cooperation in certain tasks (see Bickerton & Szathmáry, 2011) then successful communication would lead to clearly perceivable payoffs during such

^aWhile we usually refer to such symbols as 'phonemes', there is no reason why these should not be representative of gestural signs.

^bThe hearers are not assumed to know the word boundaries of an encountered utterance. For simulations with synchronized transmission the reader is referred to Vogel (2010).

cooperation. By making this assumption, we can define a learning strategy for our agents, according to which an agent’s propensity $q_m(t)$ to utilize a lexical mapping m in an interaction at time t is updated for every lexical mapping m employed in an interaction according to the following rule (cf. Barrett & Zollman, 2009):^c

$$q_m(t + 1) = \begin{cases} q_m(t) + \pi(t) & \text{if mapping } m \text{ was utilized} \\ q_m(t) & \text{otherwise} \end{cases} \quad (1)$$

Where $\pi(t)$ represents the payoff of the interaction performed between two agents at time t and is the difference between the *actual communicative success of an interaction* $s(t)$ and the predefined *level of minimum success* s_{min} (see also Bachwerk & Vogel, 2011, for an alternative use of these parameters):^d

$$\pi(t) = s(t) - s_{min} \quad (2)$$

Communicative success of an interaction $s(t)$ is defined as the level of agreement between the meanings that the agents were observed to have denoted in what was spoken and what was heard (independently of whether they segmented the sound stream identically). In particular, for communicative agreement to be achieved, the agents need to agree on two things: what entities are being spoken about and what are the relations between these entities. Hereby, the relations between the entities are considered to be understood if the ordering of the meanings is sufficiently similar between the two agents. Note that communicative success can be meaningfully measured without subscribing to omniscience: cooperative tasks may be completed, for example, without knowing which parts of which expressions were understood as intended.^e

3. Experiment Design

The current experiment has been largely inspired by a model of language evolution presented by Gong and Wang (2005), in which the agents continuously updated the strengths of their relationships, i.e. the probabilities of future interactions with other agents based on the success of preceding communication bouts. Furthermore, if the tie connecting a pair of agents exceeded a certain weight threshold, the agents were considered as ‘friends’ (the labelling had no additional implications in the model). The results of experiments performed with this model have exhibited friendships to be established between agents, as well as a certain amount of clustering occurring throughout the population. Interestingly, the model introduced no limit on how many ‘friends’ an agent may have and considered agents

^cIn the LEW, q_m values represent the weights and thus the usage probabilities of competing meaning-form mappings in an agent’s lexicon, whereby $q_m(0)$ is effectively equal to 0.

^dAs defined, a negative value of $\pi(t)$ will result in a meaning-form mapping being punished.

^eThis approach agrees with Skyrms (2010) in saying that complete understanding is not a requirement for the emergence of conventionalized communicative mechanisms.

who were ‘friends’ with every other agent simply extremely popular. We believe that the ability of an agent to become friends with every other agent as described above is a severe limitation of the model by Gong and Wang (2005). It does not seem fitting that the definition of ‘friendship’ is local for every pair of agents and does not take into consideration the strengths of an agent’s ties with other agents.^f

In order to address the above issue, we have adjusted the equations (1,2) to perform updates of the strengths of agent relationships r for every agent $i, j \in \{1, \dots, n\}$ in a population of $n = 10$ agents based on the communicative success of their interaction, with the level of minimum success s_{min} kept at 0.25, as follows:^g

$$r_{ij}(t+1) = \begin{cases} r_{ij}(t) + \alpha\pi(t) & \text{if agents } i \text{ and } j \text{ were involved} \\ r_{ij}(t) & \text{otherwise} \end{cases} \quad (3)$$

Notably, when updating the connection strengths between agents, the payoff value is multiplied by a scaling parameter α that essentially defines the ratio between the speed of social and lexical learning (cf. Santos, Pacheco, & Lenaerts, 2006). In the presented experiments, the value of the social update rate α was varied between 1, 0.5 and 0.1, i.e. the agents were always at least as efficient in updating their lexical mappings as they were in adjusting their social ties.

The evaluation of experimental results was performed with an emphasis on two particular aspects: the communicative properties of the evolving language and the emergent social structures within the population. Even though the agents in the LEW are not endowed with any telepathic abilities, it is still possible for an observer to quantify their communicative success, e.g. by computing the precision and recall rates between a speaker’s intended meaning and the hearer’s decoded one. These two values can be then summarized with the help of a statistical measure that represents their harmonic mean, defined as $F1 = 2 * \frac{precision * recall}{precision + recall}$. In addition to the success measure, we also consider such statistics as lexicon use and lexicon precision, which give an indication of respectively how broad and reliable are the lexicons of the agents. Finally, we can also evaluate how some of the more general properties of the emergent lexicons compare to those of current human languages by assessing lexicon synonymy and homonymy levels.^h

The dynamic connection strengths between agents in the presented experiment have been introduced with the goal of observing what kinds of social ties will evolve in a population which is not bound by a predefined network structure, but where its members are capable of becoming more familiar with some, and more

^fAn agent who has a connection to every other agent with $w = 0.0001$ is as likely to communicate to any of these as an equally interconnected agent, but with $w = 1.0$, making the latter agent’s ‘friendship’ ties no different from the former agent’s weak ties.

^gIf an agent’s connection strength with another agent reaches 0, the former will avoid the latter in all future interactions.

^hModern languages seem to tolerate homonymy much more so than real lexical synonymy.

distant with others. In order to be able to evaluate the results on a per-agent basis, we define three types of social connections that an agent i might end up having with a second agent t based on the connection's relative weight $w_{it} = \frac{r_{it}}{\sum_{j=1}^n r_{ij}}$, i.e. the actual probability of i selecting t as a partner in an interaction: an acquaintance ($w_{it} \geq 0.1$), a friend ($w_{it} \geq 0.25$) and a best friend ($w_{it} \geq 0.5$).ⁱ In particular, we observed if and how many connections of any given type were established by the agents, given any particular value of α , with the goal of determining how the first language-induced friendships could have been established in human evolution.

4. Results and Discussion

In total, 600 simulations runs were performed for each of the three different values of α , with 500 rounds of 10 interactions each within every such simulation. The evolution of communicative success rates from these interactions are summarized in figure 1(a), which exhibits that a nearly equivalent amount of learning occurs for all four types of agents, regardless if they are dynamically adjusting their ties to others or not. The only significant exception are agents who are very efficient at adjusting their social ties ($\alpha = 1$) who not only achieve higher understanding levels overall ($t = 6.919$, $p < 0.0001$), but also have significantly less synonymy in their lexicons ($t = -7.745$, $p < 0.0001$), as can be seen from figure 1(b). However, the difference between the various agent types is not extremely large even for the two observational values mentioned above, while it does not reach significance levels at all for such properties as lexicon size, lexicon use, or lexicon precision.

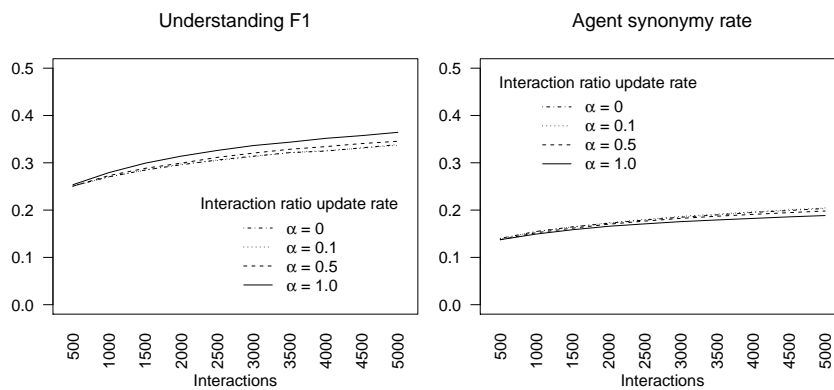


Figure 1. Evolution of (a) understanding (F1 measure) and (b) agent lexicon synonymy.

ⁱGiven this definition, the analytical limit for any particular agent is to have 2 best friends, 4 friends and 10 acquaintances, population size permitting.

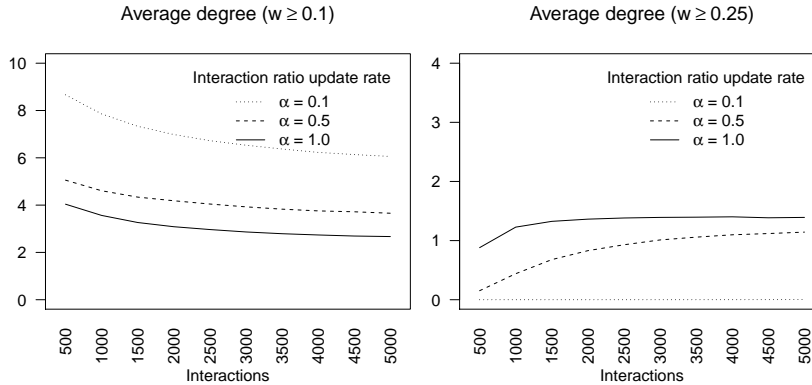


Figure 2. Evolution of average number of agent social ties with (a) $w \geq 0.1$ and (b) $w \geq 0.25$.

Having observed that agents do not seem to perform too differently even when allowed to adjust their interaction preferences throughout the simulations, we ask ourselves if the social structure of the corresponding groups is significantly evolving at all. The connection strengths between every pair of agents i, j were set to $r_{ij} = 1$ at the beginning, i.e. every agent started out with nine social ties whose relative weights $w_{ij} = 1/9$, i.e. $w_{ij} > 0.1$. From figure 2(a), one can see that this does not remain the case for long as the average number of connections per agent with $w \geq 0.1$ steadily decreases throughout the course of the simulations. In particular, for higher values of α , one observes a radical drop in the number of ‘acquaintances’ from the initial 9 to 5 for $\alpha = 0.5$ and 4 for $\alpha = 1$ by the end of the first 500 interactions (i.e. after around 100 interactions per agent).

The shrinking number of ‘acquaintances’ observed for all values of α suggests that the interaction weights are being redistributed more densely between the more popular communication partners of an agent, which leads to the establishment of first ‘friendships’. Figure 2(b) confirms that this is the case for agents with $\alpha \geq 0.5$ who on average make at least 1 ‘friend’ by the end of the simulations runs. However, agents with $\alpha = 0.1$ do not achieve a connection weight $w \geq 0.25$ with any of their conspecifics due to their ‘unwillingness’ to give up on other partners too hastily. In effect, the different social update rates model different human characters, ranging from patient pragmatics who are always cautious and thus end up having very practical friendships, if any at all, to the more impulsive types who are ready to become friends with anyone quickly. Finally, figure 3(a) shows that more than half the agents with $\alpha = 1$ end up having a ‘best friend’ by the end of the simulations, which is also the main reason why these agents outperform other agents in terms of communicative success.

When comparing the social structures that emerge in our simulations to empirical studies of human networks, we have observed two striking similarities in their

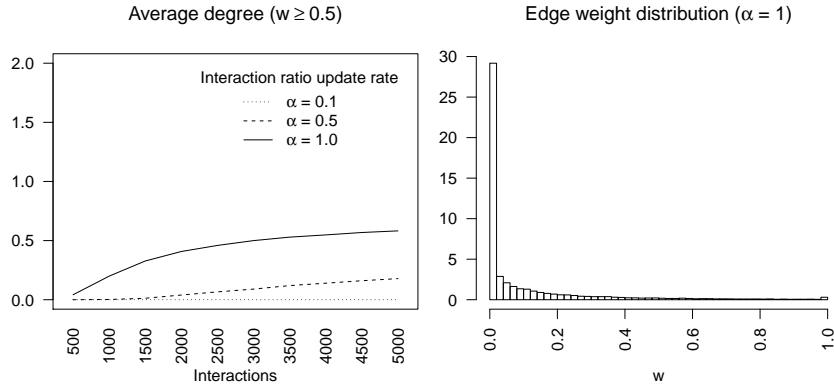


Figure 3. (a) Evolution of average number of agent social ties with $w \geq 0.5$ and (b) frequency distribution of edge weights for $\alpha = 1$.

construction and organization. First of all, the two-stage process of establishing social relationships that was observed for agents with a higher social update rate (in particular for $\alpha = 1$) has been also reported by Panzarasa et al. (2009) among others, who observed that social relationships were established at the very onset of group formation and were then slowly reinforced afterwards. Additionally, it appears that for these agents, the frequency distribution of edge weights, i.e. the connection strengths between agents, depicted in figure 3(b), has a distinct scale-free character (with the exception of an increase in $w = 1$ ties at the tail), which Barabási (2009) considers to be typical of most dynamic networks.

5. Conclusions and Future Work

In summary, it appears that in order for a non-random social structure to emerge, agents need to make their decisions in a decisive and resolute manner. In effect, the fact that social dynamics of agents in our model so closely resemble those of modern humans for certain parameter levels certainly suggests that, if our predecessors were anything like us socially, it is likely that one dominant characteristic of early hominids was a quickness to bond based on successful communication. However, one issue that arises with such individuals who become ‘friends’ with a few of their conspecifics and mainly ignore the majority of the population is that it is hard to imagine how they could have agreed on a highly conventionalized communal communication system such as language.

Future work should look further into the limits of lexicon formation in a population of socially aware agents, while making sure that the group does not evolve into pairs of exclusive ‘best friends’ who do not communicate with others. In particular, one should also investigate the possibility of the existence of either universally esteemed and popular agents or simply variably ‘social’ agents, as well as

the effects that the more influential agents could have on the decisions of others, including both their friend-making and language learning strategies.

References

- Bachwerk, M., & Vogel, C. (2010). Modelling Social Structures and Hierarchies in Language Evolution. In M. Bramer, M. Petridis, & A. Hopgood (Eds.), *Research and Development in Intelligent Systems XXVII* (pp. 49–62). Springer.
- Bachwerk, M., & Vogel, C. (2011). Establishing Linguistic Conventions in Task-Oriented Primeval Dialogue. In A. Esposito, A. Vinciarelli, K. Vicsi, C. Pelachaud, & A. Nijholt (Eds.), *Analysis of Verbal and Nonverbal Communication and Enactment: The Processing Issues* (pp. 48–55). Springer.
- Barabási, A.-L. (2009). Scale-Free Networks: A Decade and Beyond. *Science*, 325(5939), 412–413.
- Barrett, J. A., & Zollman, K. J. S. (2009). The Role of Forgetting in the Evolution and Learning of Language. *Journal of Experimental & Theoretical Artificial Intelligence*, 21(4), 293–309.
- Bickerton, D., & Szathmáry, E. (2011). Confrontational scavenging as a possible source for language and cooperation. *BMC Evolutionary Biology*, 11(1), 261.
- Christiansen, M. H., & Kirby, S. (2003). Language evolution: consensus and controversies. *Trends in Cognitive Sciences*, 7(7), 300–307.
- Gong, T. (2010). Exploring the Roles of Horizontal, Vertical, and Oblique Transmissions in Language Evolution. *Adaptive Behavior*, 18(3-4), 356–376.
- Gong, T., & Wang, W. S.-Y. (2005). Computational Modeling on Language Emergence: A Coevolution Model of Lexicon, Syntax and Social Structure. *Language and Linguistics*, 6(1), 1–42.
- Panzarasa, P., Opsahl, T., & Carley, K. M. (2009). Patterns and Dynamics of Users Behaviour and Interaction: Network Analysis of an Online Community. *Journal of the American Society for Information Science and Technology*, 60(5), 911–932.
- Santos, F. C., Pacheco, J. M., & Lenaerts, T. (2006). Cooperation Prevails When Individuals Adjust Their Social Ties. *PLoS Computational Biology*, 2(10), 1284–1291.
- Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.
- Smith, A. D. M. (2005). The Inferential Transmission of Language. *Adaptive Behavior*, 13(4), 311–324.
- Vogel, C. (2010). Group Cohesion, Cooperation and Synchrony in a Social Model of Language Evolution. In A. Esposito, N. Campbell, C. Vogel, A. Hussain, & A. Nijholt (Eds.), *Development of Multimodal Interfaces: Active Listening and Synchrony* (pp. 16–32). Springer Berlin Heidelberg.