

Learning a Common Language Through an Emergent Interaction Topology

Samarth Swarup
Dept. of Computer Science,
University of Illinois at
Urbana-Champaign,
Urbana, IL, USA
swarup@uiuc.edu

Kiran Lakkaraju
Dept. of Computer Science,
University of Illinois at
Urbana-Champaign,
Urbana, IL, USA
klakkara@uiuc.edu

Les Gasser
Graduate School of Library
and Information Science, and
Dept. of Computer Science,
University of Illinois at
Urbana-Champaign,
Urbana, IL, USA
gasser@uiuc.edu

ABSTRACT

We study the effects of various *emergent* topologies of interaction on the rate of language convergence in a population of communicating agents. The agents generate, parse, and learn sentences from each other using recurrent neural networks. An agent chooses another agent to learn from, based on that agent's *fitness*. Fitness is defined to include a frequency-dependent term capturing the approximate number of interactions an agent has had with others—its “popularity” as a teacher. This method of frequency-dependent selection is based on our earlier *Noisy Preferential Attachment* algorithm, which has been shown to produce various network topologies, including scale-free and small-world networks. We show that convergence occurs much more quickly with this strategy than it does for uniformly random interactions. In addition, this strategy more closely represents choice preference dynamics in large natural populations, and so may be more realistic as a model for adaptive language.

Categories and Subject Descriptors

E.1 [Data]: Data Structures—*Graphs and networks*; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms

Algorithms, Design, Experimentation, Languages

Keywords

Scale free networks, language evolution, preferential attachment

1. INTRODUCTION

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

One of the main problems in language evolution is to get a population of agents to converge to a common language, without globally imposing some kind of hierarchy on the population. The general problem can be stated as “Given a landscape of states and possible moves between them, and a search strategy under which agents learn from encounters with others, what constraints on the agent interaction topology will ensure the most rapid convergence to some highest-fitness state?” In effect, how should agents make choices of learning partners to achieve the fastest collective convergence to the best objective state? This is the problem we study, in the context of developing a common and high-quality language.

The “fitness” of a state in this formulation is defined by both an intrinsic measure (the objective quality of the state), and by a frequency-dependent measure (the number of agents that are in that state).

A simple strategy for rapid convergence would be to designate a special agent from which each agent learns its own language. This corresponds essentially to a pre-imposed or designed language, which may or may not be of the highest objective quality. Such a centralized system is brittle in practice because a) the teacher agent has to be responsible for adapting the language to keep up with changing tasks, environments and needs of all the agents, b) communicative load on the teacher increases at least linearly with population size, reducing scalability, and c) the centralized teacher is a single failure point. Multi-agent systems are generally distributed and open, which means that there is no central control point, and agents may enter and leave the population at any time. This means that although desirable for its speed, uniformity, and certainty, a pre-imposed interaction topology with a centralized teacher is both undesirable and unrealistic for a general multi-agent system.

2. RELATED WORK

Although there has been much work on how the topology of interaction graphs affect the convergence of a population of agents, none has focused on emergent topologies.

Delgado[2], extending the model of Shoham and Tennenholtz [5], examined the emergence of social conventions in a population of agents. His key result showed that social conventions could emerge much more quickly amongst a population of agents with small-world or scale-free interac-




Figure 1: The learning curves for frequency-dependent selection.

tion graph topologies rather than regular interaction graphs. The interaction graph topology was prespecified though, and there was no learning in the model.

Lee et. al [4] empirically studied the effect of interaction topology on the critical learning fidelity in the Komarova et. al. model [3]. Although several different interaction topologies were used, the topologies were not emergent and were prespecified by the creator of the experiments.

Our agents use simple recurrent neural networks to generate, parse, and learn languages from each other. Further, the interaction topology emerges in a distributed manner as a consequence of a specific rational agent strategy, based on our own Noisy Preferential Attachment (NPA) algorithm [6]. These two mechanisms, taken together, capture the desirable features of *agent autonomy* (learning and choice), *realism* (emergent global interaction patterns), and *flexibility* (continuous adaptation of the collective language space).

3. EXPERIMENTAL FRAMEWORK

Fitness: In our model, Each agent has a fitness, defined as follows:

$$f_i(t) = \alpha L_i(t) + (1 - \alpha) K_i(t). \quad (1)$$

$L_i(t)$ is the intrinsic fitness of the language used by agent i at time t , which may correspond to its expressivity or how good it is at describing a task. $K_i(t)$ is the number of times agent i has communicated a set of messages to another agent up to time-step t . $K_i(t)$ is normalized over the population. α is a constant representing a tradeoff between these two components of fitness.

Interaction topology: At each time-step, an agent chooses another agent with whom to interact, with probability proportional to the fitness of that agent. With a small probability p , however, it might switch to another agent chosen uniformly randomly. This “noise” is helpful in keeping the topology from getting too centralized.

The intuition is that if a large number of agents are listening to a particular agent, then it makes sense to pay more attention to that particular agent because its language is probably shared by many agents. The key thing here is that the central teacher agent *emerges* from the interactions between agents, rather than being prespecified.

The parameter p also models incomplete knowledge of the




Figure 2: The learning curves for uniformly random interaction.

interaction history of the entire population. The assumption is that, since there is no central controller telling each agent the fitness of all the other agents, each agent must maintain an estimate of the fitness of the other agents. Since the agents may not be able to observe all the other agents all the time, this estimate will have some error. It turns out that we can take advantage of this error to produce scale-free topologies. In the absence of error, i.e. if $p = 0$, frequency-dependent selection would result in a pin-wheel topology. This is because initially none of the nodes in the graph have any links, and so the first node to acquire a link acquires all the links. However, when $p \neq 0$, but is small, it can be shown that the resulting topologies have power-law like or exponential degree distributions. This model for determining interactions is derived from a version of the preferential attachment algorithm [1], which we call *Noisy Preferential Attachment* (NPA), that is based on the replicator-mutator equation of biological evolution.

4. EXPERIMENTS AND RESULTS

We did an experiment with a population of 200 agents. The intrinsic fitness term for each agent, L_i , was set to zero, because in the current experiments the agents are not situated, i.e. they are not performing a particular task, nor are they receiving any external inputs except communication from other agents.

The parameter p , which is the probability of choosing a partner uniformly randomly, was set to 0.05. The maximum sentence length was set to 10. The number of agents sampled by the learner, m , was set to 1.

The set of terminals for all the languages was the same, and consisted of five symbols: **a**, **b**, **c**, **d**, **e**. Each agent’s SRN had five input nodes, six hidden layer nodes, and five output nodes. The symbols were encoded using a 1-of-n encoding, and all weights were initialized uniformly randomly in the range $[-0.5, 0.5]$.

The simulation was run for 2000 time steps, and coherence of the population was evaluated every 20 time steps. Coherence is evaluated as the average error of an agent on a set of sentences sampled from the entire population. An agent only took the last 300 time steps into account when selecting an agent to interact with. The simulation was run twenty times with identical parameters, and the results are shown in figures 1 and 2.

	a	b	c	d	e	none	mean frequency	variance
Freq-dep Interaction	5	0	13	1	0	1	.7715	.008
– Cluster A	1	0	5	0	0	0	.644	.005
– Cluster B	4	0	8	1	0	0	.811	.001
Uniform Interaction	0	0	0	0	0	20	nan	nan

Table 1: The dominating language in the 20 frequency-dependent interaction runs and the 20 uniformly random interaction runs. The letter dominance threshold was set to 0.75 and the language dominance threshold was set to 0.5. Mean frequency and variance is the average and variance of the proportions of the dominating language.

A key difference between the runs is that the networks in clusters A and B are dominated by a single node, whereas the outliers both have multiple nodes with high in-degree. This difference can be objectively measured by calculating the *betweenness centrality* for the networks. Betweenness centrality is a value that measures how close a network is to a pin-wheel structure. It is a number between 0 and 1, and a low value means that the network is closer in topology to a ring or a fully-connected network. As shown in table 2, the average betweenness centrality for clusters A and B is very close to 1, whereas the value for the outliers is much smaller.

Network	Average betweenness centrality
High outlier	0.5451
Cluster A	0.9348
Cluster B	0.9016
Low outlier	0.7881

Table 2: The betweenness centrality for a network measures to what extent it is dominated by a single node.

The surprising fact here is that betweenness centrality is smaller for both the high and the low outlier. This can be explained by examining the actual languages to which the populations converge.

An agent’s language is dominated by the letter **a** if the frequency of the letter **a** is higher than some *letter dominance threshold*. A run is dominated by letter **a** if the proportion of agents using **a** dominated languages is higher than the *language domination threshold*. For each run of the simulation, we found the dominant language, or whether the run was not dominated.

Table 1 shows the distribution of dominating languages. In nearly all the frequency-dependent interaction runs, there was a dominating language. The average of the proportion of the dominating language is .7715. The one run that did not converge to a dominating run is the high outlier in Figure 1. In this case, the language distribution was split between an **a** language and a **b** language, which is why there is a high error for this run. This run resulted in a network with low betweenness centrality, just like the low outlier, but the difference is that in this run, the two most connected nodes in the network correspond to agents with different languages. In the low outlier, the most connected nodes all correspond to agents with the same language.

5. DISCUSSION AND FUTURE WORK

We showed that by incorporating a frequency-dependent measure in the definition of fitness for an agent, we can

generate scale-free topologies, which lead to quick convergence. This definition of fitness makes intuitive sense: if many agents are talking to a particular agent, then it pays to talk to that agent more, since it probably speaks a widely-spoken and useful language.

However, the tradeoff in this case is convergence vs robustness. Networks with high betweenness centrality, as are generally observed in our simulations, are very susceptible to attack, because removing the dominant node will cause the network to fall apart. This suggests one direction for future work: how do we converge to a stable language if agents have finite lifetimes?

The other observation is that the population converges to a “simple” language, i.e. a language consisting entirely of strings of a single symbol. It may be the case that in the absence of some external pressure towards a complex language, populations always tend towards a maximally simple common language as it is the most easily learnable. This suggests another natural followup: how do we converge to a “good” language? The “goodness” of a language can be incorporated into our framework using the frequency-independent part of fitness. Thus, agents that are particularly good at solving some task and talking about it may have high frequency-independent fitness. Therefore we intend to study convergence in a situated population of agents, where the environment would provide the pressure towards a more complex language.

6. ACKNOWLEDGEMENTS

We thank the UIUC Language Evolution Group for very helpful and stimulating discussions. This work was supported in part by NSF grant IIS-0340996.

7. REFERENCES

- [1] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, October 1999.
- [2] J. Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141:171–185, 2002.
- [3] N. Komarova, P. Niyogi, and M. Nowak. Evolutionary dynamics of grammar acquisition. *Journal of Theoretical Biology*, 209(1):43–59, 2002.
- [4] Y. Lee, T. C. Collier, C. E. Taylor, and E. E. Stabler. The role of population structure in language evolution. In *Proceedings of the 10th International Symposium on Artificial Life and Robotics*, 2005.
- [5] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1-2):139–166, jul 1997.
- [6] S. Swarup and L. Gasser. Unifying network and evolutionary dynamics. Technical Report UIUCLIS-2006/1+AMAG.