BUCLD 38 Proceedings
To be published in 2014 by Cascadilla Press
Rights forms signed by all authors

Emergence of Natural Language Lexicons: Empirical and Modeling Evidence from Homesign and Nicaraguan Sign Language

Russell Richie, Marie Coppola and Charles Yang

1. Introduction

Where do we get language from? Clearly, it requires our human brains: a chimp exposed to a lifetime's worth of language will not surpass a child who has only had a few years of exposure. But it also clearly requires something from the environment: language-deprived children clearly don't spring forth speaking Hebrew, or Greek, or Sanskrit. What is it about the human learner, and about the environments (environment broadly construed, as we will see) that they are typically exposed to, then, that give rise to language? The answer, put deceptively simply, is that language emerges when many individual humans interact; that is, language self-organizes in a population. In this paper, we consider how conventional lexicons self-organize within populations. We present data from two classes of naturally emerging signed communication systems, and we use an agent-based computational model (that is, a model that captures individuals, their interactions, and their emergent behavior) to explore the role that social networks play in the rates of conventionalization.

Language emergence has thus far mostly been studied with computational, or, to a lesser extent, experimental methodologies (Experimental Semiotics). The experimental work brings human participants into the lab to accomplish some task in pairs or groups, deprived of familiar channels of communication (speech, writing, gesture). Thus, participants must create a new communication system. The systems that emerge in these settings are surprisingly language-like, possessing conventions (Galantucci, 2005), compositionality at the sublexical (Galantucci, Kroos, & Rhodes, 2010) and lexical levels (Selten & Warglien, 2007), and form-meaning mappings that become more arbitrary with use (Theisen, Oberlander, & Kirby, 2010). While examining different aspects of

* Corresponding author, russell.richie@uconn.edu, University of Connecticut; Marie Coppola, University of Connecticut; Charles Yang, University of Pennsylvania. We are very grateful to the homesigners and their family members and friends, and the Nicaraguan Sign Language participants, for allowing us to learn from them. We thank Julia Fanghella for assistance with coding, and we especially thank Dr. Ann Senghas and the Language Acquisition and Development Laboratory for providing the Nicaraguan Sign Language videos. We also acknowledge NIH grant P30 DC010751 to Marie Coppola and Diane Lillo-Martin, and NIH/NIDCD R01 DC005407 to Ann Senghas.

language with different methodologies, what these studies have in common is that they all emphasize the role of interaction in the emergence of the communication system.

A computational modeling literature, mostly unintegrated with these experimental studies, has also examined how language might emerge from interactions among individuals. Barr (2004) examined how conventions could emerge from egocentric agents that think and act locally; perhaps surprisingly, he found that conventionalization is, in some circumstances, most likely and most efficient when agents update their behavior based on local, rather than global, system-level information. de Boer (2000) showed how symmetric and dispersed vowel systems (in which vowels are maximally distinctive), which are characteristic of the world's languages, can emerge from interactions among agents that do not explicitly attempt to optimize their vowel systems.

While focusing on interaction proper, these studies have mostly examined the effects of dyadic interaction, rather than the effects on language emergence of more global properties of a community. One of the very few exceptions is Gong, Baronchelli, Puglisi, and Loreto (2012), who used simulations of agent-based models to ask how a community's social network influences the rate at which the community conventionalizes labels for categories over a perceptual continuum (color). Among other things, they found that the social network that conventionalized most quickly was a star network (in which one agent is connected to every other agent, but these agents aren't connected to each other), followed by fully connected networks (in which every agent is connected to every other agent). This suggests that the amount and distribution of interaction within a community can influence the rate of conventionalization specifically and language emergence generally. While suggestive, these conclusions must be taken cautiously given that the simulations are not connected to empirical data, either experimental or naturalistic.

This lack of connection of computational models to empirical data, and naturalistic data in particular, pervades the literature on the role of interaction in language emergence (as well as research on collective behavior generally, see Goldstone & Gureckis, 2009). The disconnect between computational models and naturalistic data is largely due to the general paucity of naturalistic data on language emergence in the first place (most natural languages have existed for millennia), and the tendency of what little work exists to focus on the role of intergenerational transmission (Senghas, 2003), critical period effects (Senghas & Coppola, 2001), or the contribution of the language learner (vs. linguistic input) to language emergence (Coppola and Newport, 2005).

The few studies that have investigated the effects of interaction and community on language emergence have compared two rare populations of Deaf individuals, neither of which has access to accessible (i.e., visual), conventional linguistic input. First, to see the linguistic structures that can emerge in an individual *without a rich community of users engaging in rich interactions,* researchers have investigated homesign, the gesture communication systems invented by linguistically isolated, single Deaf individuals interacting with their hearing family members and friends, none of whom knows a sign language (e.g., Coppola & Newport, 2005; Coppola & Senghas, 2010). Structure surprisingly like that found in established languages emerges in these systems: a noun-verb distinction (Goldin-Meadow, Butcher, Mylander, & Dodge, 1994 (in a child homesigner); Goldin-Meadow, Brentari, Coppola, Horton, and Senghas, under review (in adult homesigners), pronominal points (Coppola & Senghas, 2010), distribution of phonological complexity in handshapes that more closely resembles that of established sign languages than that of hearing gesturers (Brentari, Coppola, Mazzoni, & Goldin-Meadow, 2012), and the grammatical relation of subject (Coppola & Newport, 2005). There are many reasons to consider homesign systems as lacking a rich community of users (see Senghas, 2005), but we focus here on the effect of the structure of a community's social network. That is, the hearing family members and friends of the homesigner do not use the homesign system among themselves, but rather, their spoken language. This means that every homesign conversation involves the homesigner (as in a star network), making the creation and use of the system highly centralized (cf. typical sociolinguistic communities, which are much more decentralized).

The second population that researchers investigate are recently formed, rich Deaf communities (i.e., communities of very many Deaf individuals) from which new natural sign languages emerge (Senghas, 2005; Meir, Sandler, Padden, & Aronoff, 2010). In contrast to the relatively limited linguistic community found in homesign systems, the sociolinguistic settings of emerging Deaf communities share one important property with those of established languages – both consist of many individuals who use the system with many other individuals. Further, those other individuals also use the system as a primary language. Keep in mind that neither homesigners nor the initial group of users of a new community sign language have access to pre-existing linguistic input; the main difference between them is with whom and how they interact. By comparing the systems of the *founding* members of such communities to homesign systems, we can determine which aspects of linguistic structure require (or at least benefit from) a community of users.

One earlier study (Osugi, Supalla & Webb, 1999) has investigated the effect of patterns of social interaction on conventionalization among 21 deaf and hearing individuals in the isolated Japanese Amami Islands. They showed that individuals' lexical items were consistent with each other to the extent that they interacted, thus obtaining a result somewhat similar to Gong et al. (2012)'s: more connections between agents/signers leads to greater conventionalization (to a point[1]). However, Osugi et al. (1999) suffers from the opposite problem of

[1] Surprisingly, Gong et al., found that the star network was faster than a fully connected network (with the same number of nodes), despite having fewer connections per node, longer paths between a given pair of nodes, and lower clustering (fewer neighbors of a given node are neighbors themselves) than the fully connected network.

Gong et al. (2012): it postulated no mechanism of conventionalization among a population of agents that would predict such effects of patterns of interaction.

Thus, while Osugi et al. (1999) and Gong et al. (2012) both suggest that, for the most part, richer patterns of interaction lead to greater conventionalization and thus language emergence, they are both limited. A more complete investigation of the effects of interaction and community structure on conventionalization would present empirical evidence of such effects, and then adduce a computational model motivated by these data that replicates the effects. This project represents such an effort. We first compare conventionalization within homesign family groups and in Nicaraguan Sign Language (NSL), a natural sign language emerging in a vibrant Deaf community (Kegl & Iwata, 1989; Senghas & Coppola, 2001; Senghas, 2003; Polich, 2005). We then present a general framework for studying conventionalization that incorporates elements of learning and social interactions. A specific implementation with reinforcement learning (Yang, 2002) appears to capture the observed trends of conventionalization. We then implement the homesign-type social network and the NSL-type social network in the model to explore the networks' effects on rates of conventionalization. Finally, we report a preliminary investigation of the effects on the rate of conventionalization of two basic structural properties of networks – path lengths and clustering – that differ between the homesign- and NSL-type networks. As a preview, our empirical results show that individual homesign family groups conventionalized more slowly than did NSL. Our simulations suggest that this may be due to the richer interaction among users of NSL relative to homesign family groups. Another round of simulations suggests that degree of clustering is at least partially responsible for this difference.

2. Study 1 – Natural language lexicons

In the empirical portion of our study, we compared deaf Nicaraguan homesigners and their family and friends with first cohort users of Nicaraguan Sign Language on the extent of conventionalization of signs referring to basic objects and concepts. Crucially, the homesign family groups and the Nicaraguan Sign Language signers had used their respective systems for roughly the same amount of time (about 25 years).

2.1. Method

2.1.1. Participants

For the homesign portion of our data, participants were four deaf Nicaraguan homesigners [3 male; aged 24 to 33 years (M=29)] and nine of their hearing family members and friends [4 male; aged 17 to 59 (M=30)]. We henceforth refer to these family members and friends as communication partners. In each family group, we tested the homesigner's mother and one or two additional communication partners, who were either siblings (n = 4) or a friend (n = 1). The homesigners have minimal or no interaction with other deaf individuals, including each other, and have minimal or no knowledge of Nicaraguan Sign Language or spoken or written Spanish. Instead, they have each been using their respective invented gestural homesign systems all their lives. Despite their lack of linguistic input, they socialize with others, hold jobs, have families, and otherwise lead typical lives. For the Nicaraguan Sign Language portion of our data, participants were eight Deaf users of Nicaraguan Sign Language (2 males; 21-32 years, M=27). These signers had little to no interaction with other Deaf individuals prior to 1978, when they arrived at the Center for Special Education in Managua and formed the Deaf community there. Thus, these individuals were among the first cohort of NSL users, who began to form NSL *de novo*.

2.1.2. Stimuli

Stimuli were images of 9 familiar and basic objects and concepts: cat, dog, cow, rain, sun, ice, egg, fish, and orange.

2.1.3. Procedure

For homesign families, data were collected in 2011. For NSL users, data were collected in 2003. Participants were individually shown images of the objects and concepts described above. Using gesture and facial expressions, we elicited participants' gestural responses to these images. Hearing participants were asked to use only their hands to respond, and all were easily able to do the task. Participants responded to the camera, not to each other, and were not allowed to see each other's productions. All responses were videotaped for later analysis**.**

2.1.4. Coding

Participants' responses were coded by a research assistant and RR. A majority of responses produced by homesigners or their communication partners contained more than one gesture (62%, cf. 10% of responses by NSL signers); we coded every gesture individually for its *Conceptual Component* (CC), or aspect of the item's meaning that the gesture iconically represented. For example, a response to 'cow' might contain two gestures, one iconically representing horns (its CC is thus HORNS) and another iconically representing milking (its CC is thus MILKING).

2.2. Results

Treating every CC as a dimension in a combinatorial space, every response can be represented as a binary-valued vector, with 1 representing the presence of a given CC and 0 its absence. The distance between two responses to the same

Table 1. Sample calculations of frequency-weighted hamming distance

'cow'	HORNS	MILKING	DRINK	Distance (from Homesigner)
Homesigner	1	1	0	*n/a*
Sister	0	1	1	2/5*1 + 2/5*0 + 1/5*1 = .6
Mother	1	0	0	2/5*0 + 2/5*1 + 1/5*0 = .4
Frequency	*2*	*2*	*1*	

object is thus a measure of conventionalization. We define distance as the number of vector values by which two responses differ, and weight more heavily those vector values corresponding to CC's used more frequently (i.e., disagreement on the use of the CC ROUND will lead to a greater distance than disagreement on the infrequent CC MILKING.[2] See Table 1 for sample calculations.

For homesign families, we calculated the distance between each homesigner's response and that of each homesigner's communication partner's responses for a given object. For the NSL community, we computed the distance between each pair of NSL signers for a given object. For each pair (whether homesigner to partner or NSL signer to NSL signer), we average these distances across all tested objects, yielding an overall measure of lexicon distance or conventionalization between a pair. Because the NSL community has no user who is the obvious standard of form (unlike the homesign systems, where homesigners serve as the standards to which we compare each communication partner), we computed the weighted distances between all pairs of NSL users, and then averaged across all these pairs. This average came out to 0.0045. We then compared this value to the Homesigner-Communication Partner distances (see Figure 1): a one-sample Wilcoxon Signed Rank test determined that the median homesigner-partner distance was significantly greater than the NSL average distance ($W=36$, $p < 0.01$). Given this, and that NSL and the homesign systems had been used for comparable lengths of times, these results suggest that NSL conventionalized faster than the various homesign systems.

2.3. Discussion

We showed above that the NSL community conventionalized a basic lexicon to a greater extent than had any of the present homesign systems after comparable periods of use. What might explain this difference in rate of conventionalization between homesign and NSL? One possibility concerns the differences in patterns of interaction between users of homesign systems and users of NSL (and other Deaf community sign languages, Woll & Ladd, 2003). While the deaf user of a homesign system uses the system for all interactions, with other users of the system *using the system*. In other words, the homesign

[2] CC's used more frequently offer more opportunities for convergence, and so should arguably be weighted more heavily in calculating distance. In addition, this measure ignores the ordering of gestures and is thus an overestimate of conventionalization.

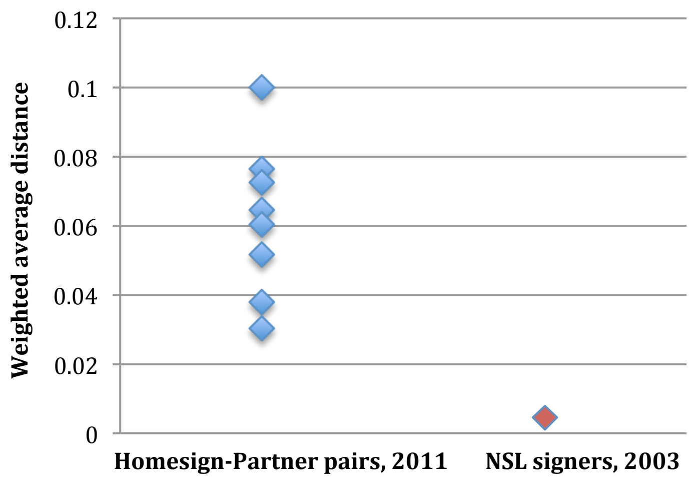

Figure 1. Average distances, across all objects and concepts tested. Users of Nicaraguan Sign Language conventionalized common forms for lexical items much faster than all Homesigner-Partner pairs tested.

interactive structure is a star network, while the NSL/Deaf community structure approximates a fully connected network[3]. We now turn to our model with which we test these predictions.

3. Study 2 – Modeling conventionalization

What are the conditions for conventionalization, whereby a shared lexicon emerges through strictly local linguistic interactions among linguistic individuals? At least two elements of the process suggest themselves. First, the individuals must be "lexicon ready". In the simplest case, they must be able to maintain a list of form-meaning pairings. Similar to our study of homesigns, the individuals must be capable of making combinatorial use of constitutive units as in our case of Conceptual Components. Second, the individuals must be capable of learning, or modifying their lexicon as the result of linguistic and social interactions. In this section, we first describe a general framework to study lexical conventionalization. We then study its dynamics through the use of reinforcement learning, where behaviors are rewarded or punished, making their probability of appearance in the future more or less likely, respectively (Bush &

[3] The NSL community is more likely a "small-world" network, as social networks tend to be of these types (Watts & Strogatz, 1998). However, fully connected networks and small-world networks share many structural properties, including low path lengths and high clustering, so approximating the NSL community as a fully connected network may not be terribly inaccurate.

Mosteller, 1951; Yang, 2002) as a model of learning and social interactions. Last, we use the model to test the hypothesis regarding the difference in rate of conventionalization between homesign and NSL.

3.1. The framework

Consider a population of N agents communicating a set of meanings through the combinatorial use of C binary signs that are analogous to Conceptual Components in the homesign data. For a specific meaning, agent i accesses a vector of probabilities $P_c = \{p_c^i\}$, defined over these signs ($j = 1,2,\ldots,C$) such that with probability p_c^i, the cth sign is used by agent i and with probability ($1 - p_c^i$), the cth sign is not used. This representation can also be used to encode atomic use of signs, i.e., each meaning is expressed by one sign, in which case the vector $\sum_c p_c^i = 1$ (i.e., agent i has a probabilistic distribution of the signs and only one of them is chosen at each instance of use).

The central premise of the conventionalization model is that individuals adjust their choices of linguistic encoding in attunement with their communicative partners. To communicate a meaning, agent i instantiates a vector U_i of 0's and 1's according to P_i. Agent j, the listener, makes adjustments to P_j to agree with agent i by the use of some learning algorithm. The changes in the distance between P_j and P_i over time represent the extent of convergence or conventionalization.

Linguistic communications among agents may also have a social component. Consider a matrix $S = [s_{i,j}]$, which defines the probabilities of communication between agents i and j such that for all i, $\sum_j s_{i,j} = 1$. The social matrix provides a general platform to encode patterns of interactions among agents. A matrix with positive probabilities only among the neighboring agents, for instance, is a straightforward implementation of Schelling (1971)'s classic model of segregation. The matrix may be fixed or it may change as the result of communication. For instance, it seems reasonable that agents would modify their partner preferences based on past successes or failures of communication, which can be modeled as $s_{i,j}$ increasing if a successful communication has occurred between agent i and j, and decreasing upon failure.

As the result of the communicative interactions, the probability vectors for agents $\{P_i\}^t$ change over time, which characterizes the evolution of the lexicons in the population. In general, the dynamics of $\{P_i\}^t$ can be analyzed as a Markov Chain, first used by Niyogi & Berwick (1997) to study language learning and change. Different choices of the learning algorithm (L), which may be discrete or probabilistic (including Bayesian inference), the social matrix S (and its own evolution), together with the current values in $\{P_i\}^t$ define the transition matrix T^t at time t, which can be multiplied with $\{P_i\}^t$ to produce the next state of lexicon $\{P_i\}^{t+1}$. Similar models have been developed in the iterated learning framework (e.g., Kirby, Dowman & Griffiths, 2007).

3.2. Conventionalization through Reinforcement Learning

In what follows, we propose a specific learning model and consider several variant implementations relevant to the present study of sign convergence. The learning model is an instance of reinforcement learning (Bush & Mosteller, 1951), a simple, efficient and domain general model of learning now with considerable behavioral and neurological support (see Niv, 2009 for review), and one which has been used in computational and empirical studies of language acquisition (Yang, 2002). Let agent *j*'s current probability for sign *c* be *p*. Upon each communication, the listener *j* adjusts *p* to match agent *i*'s choices, following the Linear-Reward-Penalty (L_{RP}) scheme of Bush & Mosteller (1995) where the magnitude of change is a linear function of the current value of *p*:

- Agent *i* chooses 1: $p' = p + \gamma\,(1 - p)$
- Agent *i* chooses 0: $p' = (1 - \gamma\,)p$

where the learning rate γ is typically a small real number. All probabilities are subsequently renormalized.

3.3. Social matrix: Homesign vs. Language

We consider two social networks of agents. In the first, analogous to the homesign situation, one individual, the deaf user (say agent 1), communicates with all other (hearing) individuals, who do not use signs to communicate with each other. The matrix is initialized such that $s_{i,j} = 1 / (N - 1)$ where N is the total number of agents, $s_{i,1} = 1$ ($i \neq 1$) and $s_{i,j} = 0$ ($i, j \neq 1$). We also consider what can be referred as the language matrix, where all agents are Deaf and use signs to communicate with each other ($s_{i,j} = 1 / (N - 1)$, $i \neq j$), which corresponds more closely to the sociolinguistic settings of sign language emergence in Deaf communities (Woll & Ladd, 2003).

3.4. Results and Discussion

In our simulations, we consider a population of $N = 5$ agents, who discuss 1 object using some combination of 40 conceptual components. The choice of these particular parameter values was somewhat arbitrary; further exploration of the parameter space is planned. For each sign, we initialize the values in P_i for each agent randomly between 0 and 1; they start out preferring either the use or the non-use of each sign with random probabilities. The learning rate γ is set to 0.01 and is used for the adjustment of both P_i's. For each simulation, we run the simulations over 2 million instances of communications; in the case of convergence, i.e., all *N* agents in complete agreement with respect to sign usage (all P_i's above .98 or below .02), we record the number of iterations required for convergence. The main result is as follows: there is a significant difference in convergence time between the homesign-type model and the language-type model ($p<10^{-12}$; see Table 2). These results were robust under variations in the

Table 2 Average number of iterations to convergence (percentage of simulations reaching convergence in 2 million iterations)

	Homesign	Language
Average number of interactions to convergence	698K	260K
Percent simulations converging within 2M interactions	80%	100%

number of agents, objects, and conceptual components, indicating the importance of a mutually engaged community for the rapid emergence of a true linguistic system, and offering a potential explanation for the difference in rates of conventionalization between homesign and Nicaraguan Sign Language.

What remains unanswered so far, however, is which aspects of the network structural properties are responsible for this difference. The two networks differ in several ways: the average node degree (number of connections per node) is higher in the NSL network, the NSL network is more clustered (neighbors of a node are neighbors themselves), the average shortest path from one node to another is lower in the NSL network, and the NSL network is less centralized than the homesign network. We thus conducted another series of preliminary simulations to investigate which network properties might be responsible for hastening conventionalization.

Our first hypothesis was that low path lengths were contributing most to the high rate of conventionalization, as low path lengths allow information—here, the form-meaning mappings—to propagate quickly through a network. To manipulate this property while leaving constant other properties aside from clustering coefficient, we created regular lattice networks, where nodes are connected in a ring, and then also connected to their k nearest neighbors. We then varied across networks the probability of randomly rewiring a given connection between nodes, as in Watts and Strogatz (1998), who showed that, as the rewiring probability increases, path length and then clustering decrease. Importantly, note that this process leaves the *average* node degree unchanged. We then ran simulations on these networks varying in rewiring probability, path length, and clustering. The number of agents was set to 300, each agent was connected to their 10 nearest agents, the number of objects was set to 1, number of conceptual components to 1, and gamma to .1. The number of agents needed to be high (cf. the 5 agents from our earlier simulations, motivated by the size of the homesign communities) to replicate the networks of Watts and Strogatz (1998), while the other parameters were set to their values to keep simulation time reasonable—the robustness of the difference between full and star networks to differences in these parameters suggested to us that this was a reasonable decision. We found that time to conventionalization was remarkably stable as the rewiring probability increased and path length dropped drastically—time to conventionalization only increased when clustering dropped to near 0 (see Table 3). In fact, the only significant difference in conventionalization time across

Table 3. Simulations testing effect of path length and clustering coefficient

		Rewiring probability, *r*			
		0	0.01	0.1	1
Time to convergence	Mean	1.3E+04	1.3E+04	1.3E+04	1.6E+04
	SD	*2.8E+03*	*2.4E+03*	*2.6E+03*	*3.8E+03*
Average Clustering Coefficient	Mean	6.7E-01	6.5E-01	4.9E-01	3.0E-02
	SD	*2.2E-16*	*4.9E-03*	*1.1E-02*	*2.3E-03*
Average shortest Path length	Mean	1.5E+01	6.4E+00	3.5E+00	2.7E+00
	SD	*1.3E-14*	*6.9E-01*	*5.2E-02*	*2.9E-03*

rewiring probabilities was between $r = 1$ (a completely random network) and the other values of r (p-values less than 10^{-4}). From this we tentatively conclude that clustering coefficient, but not path length, has a significant effect on rate of conventionalization, contrary to our initial hypothesis.

4. General Discussion

In the current work, we (1) showed that conventionalization in homesign systems has proceeded more slowly than in Nicaraguan Sign Language (NSL), a sign language emerging from a recently formed Deaf community; (2) formulated a general framework and causal model of conventionalization, in the form of a multi-agent reinforcement learning model that obtains conventionalization; (3) showed that an NSL-inspired model where all agents interact with each other converges significantly faster than a homesign-inspired model in which one agent (i.e., a deaf individual) interacts with every other agent (i.e., hearing individuals), but these other agents interact only with the first agent; (4) showed that high clustering coefficients but apparently not low path lengths hasten conventionalization.

The empirical and computational results (1) - (3) converge on the following conclusion: richer networks hasten conventionalization. Our work is somewhat different in this respect from Gong, Baronchelli, Puglisi, and Loreto (2012), also a simulation study investigating the effects of social network structure on the emergence of language. In their model, agents had to carve a perceptual continuum into perceptual categories, and then agree upon labels that refer to one or more categories. In contrast to our model, where the communicating agents know the referent (e.g., moon) but adjust the probabilities of producing corresponding gestures, Gong et al.'s agents must infer the referent through perception. Whereas we found that full networks converged much faster than star networks, Gong and colleagues found that star networks offer comparable convergence properties as do fully connected networks. These modeling studies suggest that the dynamics of individual interactions in social networks are fairly complex a priori, and may be clarified by incorporating additional empirically grounded factors from either experimental (e.g., Selten & Warglien, 2009) or naturalistic considerations, as in the present study.

Our additional simulations (4) suggested that one particular component of the richer networks of NSL and other typical sociolinguistic communities that hastens their conventionalization is their higher degree of clustering, or degree to which neighbors of a given node are neighbors themselves. However, we can't completely rule out an extremely non-linear effect of path length on conventionalization, or an effect of other structural properties (average degree, the degree distribution generally, or centralization). Future work thus has a number of clear avenues for exploring the effect of network structures on performance of the model presented here, and in language emergence and change generally.

References

Barr, Dale. (2004). Establishing conventional communication systems: Is common knowledge necessary? *Cognitive Science*, *28*(6), 937–962. doi:10.1016/j.cogsci.2004.07.002

Bush, Robert, & Mosteller, Frederick. (1951). A mathematical model for simple learning. *Psychological Review*, *68*, 313-323.

Brentari, Diane, Coppola, Marie, Mazzoni, Laura, & Goldin-Meadow, Susan. (2012). When does a system become phonological? Handshape production in gesturers, signers, and homesigners. *Natural Language and Linguistic Theory*, *30*(1), 1-31.

Coppola, Marie, & Newport, Elissa. L. (2005). Grammatical Subjects in home sign: Abstract linguistic structure in adult primary gesture systems without linguistic input. *Proceedings of the National Academy of Sciences*, *102*(52), 19249–53. doi:10.1073/pnas.0509306102

Coppola, Marie, & Senghas, Ann. (2010). The emergence of deixis in Nicaraguan signing. In Diane Brentari (Ed.), *Sign Languages: A Cambridge Language Survey* (pp. 543–569). Cambridge, UK: Cambridge University Press.

de Boer, Bart. (2000). Self-organization in vowel systems, *Journal of Phonetics*, *28*(4), 441–465.

Galantucci, Bruno. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, *29*(5), 737–67. doi:10.1207/s15516709cog0000_34

Galantucci, Bruno, Kroos, Christian, & Rhodes, Theo. (2010). The effects of rapidity of fading on communication systems. *Interaction Studies*, *11*(1), 100–111. doi:10.1075/is.11.1.03gal

Goldin-Meadow, Susan, Brentari, Diane, Coppola, Marie, Horton, Laura, & Senghas, Ann. (under review). Watching language grow in the manual modality: How the hand can distinguish between nouns and verbs.

Goldin-Meadow, Susan, Butcher, Cynthia, Mylander, Carolyn, & Dodge, Mark. (1994). Nouns and verbs in a self-styled gesture system: What's in a name? *Cognitive Psychology, 27*, 259-319.

Goldstone, Robert. & Gureckis, Todd. (2009). Collective behavior. *Topics in Cognitive Science, 1*, 412-438.

Gong, Tao, Baronchelli, Andrea, Puglisi, Andrea, & Loreto, Vittorio. (2012). Exploring the roles of complex networks in linguistic categorization. *Artificial Life*, *18*(1), 107–21. doi:10.1162/artl_a_00051

Kegl, Judy, & Iwata, Gail. (1989). Lenguaje de Signos Nicaragüense: A pidgin sheds light on the "creole?" ASL. In R. Carlson & S. DeLancey & S. Gilden & D. Payne & A. Saxena (Eds.), *Fourth Annual Meeting of the Pacific Linguistics Conference* (pp. 266-294). Eugene, OR.

Kirby, Simon, Dowman, Mike, & Griffiths, Thomas. (2007). Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences*, *104*(12), 5241-5245.

Meir, Irit, Sandler, Wendy, Padden, Carol, & Aronoff, Mark. (2010). Emerging Sign Languages. In Marc Marschark and Patricia Spencer (Eds.) *Oxford Handbook of Deaf Studies, Language, and Education*, Volume 2. Oxford: Oxford University Press, (pp. 267-280).

Niyogi, Partha, & Berwick, Robert. (1997). Evolutionary consequences of language learning. *Linguistics and Philosophy, 20 (6),* 697-719.

Osugi, Yutaka, Supalla, Ted, & Webb, Rebecca. (1999). The use of word elicitation to identify distinctive gestural systems on Amami Island. *Sign Language and Linguistics*, *2*(1), 87–112.

Polich, Laura. (2005). *The Emergence of the Deaf Community in Nicaragua: "With Sign Language You Can Learn So Much."* Washington, DC: Gallaudet University Press.

Schelling, Thomas. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology, 1*, 143-186.

Selten, Reinhard, & Warglien, Massimo. (2007). The emergence of simple languages in an experimental coordination game. *Proceedings of the National Academy of Sciences*, *104*(18), 7361–6. doi:10.1073/pnas.0702077104

Senghas, Ann, & Coppola, Marie. (2001). Children creating language: how Nicaraguan sign language acquired a spatial grammar. *Psychological Science*, *12*(4), 323–8.

Senghas, Ann. (2003). Intergenerational influence and ontogenetic development in the emergence of spatial grammar in Nicaraguan Sign Language. *Cognitive Development*, *18*(4), 511–531. doi:10.1016/j.cogdev.2003.09.006

Senghas, Ann. (2005). Language emergence: Clues from a new Bedouin sign. *Current Biology*, *15*(12), 463–465.

Theisen, Carrie, Oberlander, Jon, & Kirby, Simon. (2010). Systematicity and arbitrariness in novel communication systems. *Interaction Studies*, 1971–1976.

Watts, Duncan & Strogatz, Steven. (1998). Collective dynamics of 'small-world' networks. *Nature, 393,* 440-442.

Woll, Bencie, & Ladd, Paddy. (2003). Deaf communities. In Marc Marschark & Patricia Spencer (Eds.), *Oxford Handbook of Deaf Studies, Language, and Education*, (pp. 151-162), Oxford: Oxford University Press.

Yang, Charles. (2002). *Knowledge and learning in natural language*. New York: Oxford University Press.