

Modeling the Emergence of Natural Language Lexicons

David Russell Mahood Richie

B.A., Bates College, 2009

A Thesis

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Master of Arts

At the

University of Connecticut

2013

APPROVAL PAGE

Masters of Arts Thesis

Modeling the Emergence of Natural Language Lexicons

Presented by

David Russell Mahood Richie, B.A.

Major Advisor _____
Marie Coppola

Associate Advisor _____
Letitia Naigles

Associate Advisor _____
Whitney Tabor

University of Connecticut

2013

Abstract

It is largely acknowledged that natural languages emerge from not just human brains, but also from rich communities of interacting human brains (de Boer, 2000; Galantucci et al., 2012; Senghas, 2005). Yet the precise role of such communities and such interaction on emergence of core properties of language has largely gone uninvestigated in naturally emerging systems, leaving the few existing computational investigations of this issue (de Boer, 2000, i.a.) somewhat ungrounded. Here we take a step towards investigating the precise role of community structure in the emergence of linguistic conventions with both naturalistic empirical data and computational modeling. We first show conventionalization of lexicons in two different classes of naturally emerging signed systems: (1) protolinguistic “homesigns” invented by linguistically isolated Deaf individuals, and (2) a natural sign language emerging in a recently formed rich Deaf community. We find that the latter conventionalized faster than the former. Second, we model conventionalization as a population of interacting individuals who adjust their probability of sign use in response to other individuals’ actual sign use. Simulations suggest that a richer social network, like that of natural (signed) languages, conventionalizes faster than a sparser social network, like that of homesign systems. We discuss our behavioral and computational results in light of other work on language emergence, other work of collective human behavior on complex networks, and other work on language development.

Keywords: lexicon; homesign; conventionalization; language emergence; language development; agent-based computational modeling; sign language; social networks.

Acknowledgments

I would like to thank our deaf and hearing Nicaraguan participants and the members of the Coppola lab, the Sign Language Reading Group at the University of Connecticut, audiences at the Texas Linguistics Society Conference in 2012 and the Cognitive Science Conference, and other venues where this work was presented. Special thanks go to Julia Fanghella for coding the homesign data, to Ann Senghas and the Language Acquisition and Development Research Lab for providing the NSL data, to Charles Yang for his work in devising the computational model presented here, and to my committee members Letty Naigles and Whitney Tabor, whose feedback on this work enhanced and sharpened it immeasurably. Extra special thanks go to my advisor on this work, Marie Coppola, for her guidance, patience, and understanding. I would also like to thank three anonymous reviewers of the version of this work presented at the meeting of the Cognitive Science Society in 2013. This research was supported by NSF IGERT grant #1144399 to UConn, NIH grant P30 DC010751 to Marie Coppola and Diane Lillo-Martin, NIH T32 and F32 fellowships to Marie Coppola, NSF grant BCS 0547554 to Diane Brentari, and NIH grant R01 DC00167 to Elissa L. Newport.

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, as philosophers and kings of old suspected. 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? We could easily of course say that a baby acquires human language when her parents expose them to human language. But where did those parents get their language? Repeating this cycle a few times, we are led to ask, how does language emerge in the first place, when none already exists? The answer, put simply (very deceptively so), is that language emerges when many individual humans interact; that is, language self-organizes in a population. In this paper, we will consider how conventional lexicons self-organize within populations. We will present data on conventionalization from two classes of emerging signed communication systems, and we will use an agent-based computational model (model that captures individuals, their interactions, and their emergent behavior) to explore the role these systems' social networks play in their respective rates of conventionalization.

Before delving further into this phenomenon, we take a moment to consider whether we have left the domain of developmental psychology, and entered sociolinguistics, artificial life, or something else entirely. While this line of questioning may also be relevant to these fields, we believe the concerns of developmental psychology are still front and center. Though we have started talking about how language emerges in groups of humans, the mind of the individual is critical to this enterprise. As we said, language requires human brains, and so any model of language emergence will have to posit a learner who, when put into a group of similar learners,

is capable of giving rise to human language. We will propose such a learner in this work. Of course, what exactly human learners must be endowed with, so that they can ultimately attain adult human behavior at any level, be it group or individual, is precisely what developmental psychology is concerned with. Indeed, this is essentially what the nature-nurture debate is about (Spelke & Newport, 1998). We will return in the General Discussion to a more in-depth treatment of the relation of this work and language emergence more generally to traditional issues in developmental psychology. These include the issue of critical period effects in language emergence, and comparison of the emergence of the lexicon in typical acquisition and in emerging languages.

As stated, the phenomenon of language emergence can't be investigated in typical language acquisition situations – the input children receive from those around them overrides whatever they might come up with on their own. Instead, research on this question has thus far mostly been experimental, or, to a lesser extent, computational. In the experimental work (a field now known as Experimental Semiotics), human participants are brought into the lab to accomplish some task in pairs or groups, while familiar channels of communication (speech, writing, gesture) are deprived, thus requiring participants to 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 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. Selten & Warglien (2007) found that conventionalization within a pair was fastest when one partner made

many more changes to their word-forms than did the other. Garrod, Fay, Lee, Oberlander, and MacLeod (2007), emphasizing the role of feedback and common ground, found that form-meaning mappings became arbitrary most quickly when participants could give each other feedback on each other's forms.

A computational modeling literature mostly unintegrated with this experimental work 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 (that is, systems where vowels are maximally distinctive), 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, the above studies have mostly examined the effects of dyadic interaction, rather than the effects of more global properties of a community on language emergence. One of the very few exceptions is Gong, Baronchelli, Puglisi, and Loreto (2012), who used simulations of agent-based models to investigate how the social network of a community influences the rate at which the community conventionalizes labels for categories over a perceptual continuum (e.g., color). They found that the fastest social network was a star network (one agent is connected to every other, who aren't connected to each other), followed by fully-connected networks (every agent connected to every other agent), small-world networks (a network with mostly local connections and a few long-range connections), and scale-free networks (a few highly-connected nodes and many not-well-connected nodes; more technically,

a network whose degree distribution follows a power law), followed by lattices (network whose drawing forms a regular tiling, as in a checkerboard or a chicken-wire fence) and ring networks (every node connects to exactly two other nodes, forming a ring). This would suggest that the amount and distribution of interaction within a community can influence the rate of conventionalization specifically and language emergence generally: specifically, that conventionalization is fastest in social networks with higher average node degree (i.e. networks with more highly connected nodes, e.g. fully-connected, small-world, scale-free), lower average shortest path length (as measured by number of connections traversed) between any two nodes (fully-connected, small-world, scale-free), and greater variability in node centrality (roughly, the relative importance of a node in a network; star network). While suggestive, these conclusions must be taken somewhat cautiously given that the simulations are not connected to any empirical data, experimental or naturalistic.

This lack of connection of computational models to empirical data, and naturalistic data in particular, is common in the literature of the role of interaction in language emergence¹. 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. First, to see the linguistic

¹ The disconnect between experimental and computational approaches is a general concern for research on collective and cooperative behavior (see Goldstone & Gureckis, 2009 for review).

structures that emerge *without a rich community of users engaging in rich interactions*, researchers have investigated homesign systems, the gestural communication systems invented by linguistically-isolated, lone Deaf individuals interacting with hearing family and friends (e.g., Coppola & Newport, 2005; Coppola & Senghas, 2010). Structure surprisingly like that in natural languages emerges in these systems: a noun-verb distinction (Goldin-Meadow, Butcher, Mylander, & Dodge, 1994), pronominal points (Coppola & Senghas, 2010), phonological complexity in handshapes that more closely resembles that of Deaf signers than that of hearing gesturers (Brentari, Coppola, Mazzoni, & Goldin-Meadow, 2012), and the grammatical relation of subject (Coppola & Newport, 2005). We say such structure emerges in homesign without a rich community of users engaging in rich interactions for the following reasons. First, the communities tend to be rather small, with perhaps a dozen individuals using the system with any regularity. Second, the hearing family members and friends of the homesigner do not use the system with each other, as they can and do use 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, more like small-world or scale-free networks). Third, the contexts of use of the system and the subjects discussed using the system are generally rather restricted, often focusing on achieving a concrete goal (i.e., making food in the kitchen) where things under discussion can often be pointed to, rather than abstract goals (e.g., discussing the rights of Deaf individuals). Last, spontaneous conversations between homesigners and their partners tend not to be long and bidirectional, and are often characterized by frustration on both ends; in an experimental setting where context is removed, partners often have great difficulty understanding homesigners' descriptions of simple events (Carrigan & Coppola, 2012).

The second population that researchers investigate are recently formed, rich, integrated, 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 homesign systems, emerging Deaf communities are much like the sociolinguistic communities of any established language: they have large numbers of users, who all use the system with many other individuals, and in diverse contexts and about diverse subjects. By comparing the systems of the *earliest*² 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.

While this tack has been taken (but not often), the effects of rich interaction and community on language emergence that *have* been identified in this literature tend to be aspects of language that are arguably not central to language. Emerging sign languages, in contrast to homesign systems, have a count list (Flaherty & Senghas, 2011; Spaepen et al., 2011), *more* pronominal points, and greater integration of these points in their grammars (Coppola & Senghas, 2010), and more overt noun phrases for the first mention of a character in discourse (Coppola, Gagne, & Senghas, 2013). One exception to this trend is from Osugi, Supalla, and Webb (1999), who, like Gong et al. (2012) investigated the effect of patterns of social interaction on conventionalization (convention being a core property of language) among 21 deaf and hearing individuals in the geographically and genetically isolated Koniya region of Amami Island south of Japan. Showing images of objects and eliciting gestures for them, they showed that individuals (both Deaf and hearing) 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

² It is worth emphasizing that we are considering only the earliest members of such communities here. If we look at the structure emerging among later members of such communities, we would be looking not only at the effect of rich interactions and communities, but also at the effect of intergenerational transmission and acquisition (e.g. Senghas, 2003).

greater conventionalization (to a point³). However, Osugi et al. (1999) suffers from the opposite problem of 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 would lead to greater conventionalization and thus language emergence, they are limited by opposite problems. A more complete investigation, then, 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 this data that replicates the effects. This project represents such an effort. We first present new longitudinal data on conventionalization from naturally emerging homesign systems. We then compare this data to prior, indirect reports and new cross-sectional data of lexical consistency in Nicaraguan Sign Language (NSL), a natural sign language emerging in a vibrant Deaf community (Senghas & Coppola, 2001; Senghas, 2003). 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. Finally, we implement the homesign-type social network and the NSL-type social network in the model to explore the effects of these networks on rates of conventionalization. To preview our results, our empirical results show that homesign families conventionalize more slowly than did NSL; our simulations suggest that this may be because NSL has a more inter-connected network of users than do homesign systems.

³ Gong et al., quite surprisingly, found that the star network was faster than a fully-connected network of the same number of nodes, despite the former having lower average degree, higher average shortest path length, and lower clustering than the latter.

Natural language lexicons

In the empirical portion of our study, we (a) examine conventionalization over a 9-year period in form-meaning mappings for basic objects and concepts among deaf Nicaraguan homesigners and their family and friends, and (b) compare this to prior, indirect reports of conventionalization in the first cohort of Nicaraguan Sign Language users, and new cross-sectional data of lexical consistency in the same NSL users.

Study 1 -- Homesign

Method

Participants Participants were four deaf Nicaraguan homesigners [3 male; aged 11 to 33 years (M=24) at various times of testing] and nine of their hearing family members and friends [4 male; aged 10 to 59 (M=30) at various times of testing; we henceforth refer to these family and friends as communication partners]. 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, these homesigners have 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 have typical lives. See Table 1 for relations between the homesigners and their family members.

Family 1	Family 2	Family 3	Family 4
Homesigner	Homesigner	Homesigner	Homesigner
Mother	Mother	Mother	Younger brother
Older brother	Younger brother		Younger sister
Friend	Younger sister		

Table 1. The four homesign groups, and the relations the hearing communication partners bear to their respective homesigner.

Stimuli Stimuli were images of 22 basic objects and concepts. All items were familiar to participants. Nineteen of these objects and concepts were taken from Osugi et al. (1999), which itself was derived from Swadesh (1971). The stimulus items were: boy, cat, cold, cook, cow, dog, egg, fire, fish, flower, ice, girl, hot, moon, orange, palm tree, potato, rain, snake, stones, and sun.

Procedure In 2002, 2004, 2006, and 2011, participants were shown images of the objects and concepts described above. Participants were tested individually. 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.

Coding Participants' responses were coded by a research assistant in consultation with the author. A majority of responses contained more than one gesture (2 gestures: 40%, 3 gestures: 15%, 4 gestures: 4%, and 5 gestures: 2%); 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)⁴.

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 the absence. The distance between two responses to the same object is thus a measure of

⁴ We have also coded every gesture for its formal components (e.g., handshape, location, movement), but this coding does not bear on the current analysis, and so we do not discuss it further.

conventionalization. We define distance here 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. See Table 2 for sample calculations)^{5,6}. For a given object in a given year, we calculated this distance between each homesigner's response and that of each homesigner's communication partner's responses. For example, we calculate the distance between Homesigner 1's 2011 response to 'cow' and his mother's 2011 response to 'cow', as well as their 2006, 2004, and 2002 responses to 'cow'. For each homesigner-partner pair and year, we average these distances across all tested objects, yielding an overall measure of lexicon distance or conventionalization between a pair. Results are summarized in Figure 1, which shows decreases in lexicon distance across partners. To give a sense of the scale of weighted distance, consider a partner that with probability P will agree with a homesigner in the usage of a CC. Simulations show that a partner agreeing with a homesigner 92.5% of the time gives a weighted distance of .069, and agreeing 96% of the time gives a weighted distance of 0.036, that is, a ~50% reduction in error. This is roughly the change a typical communication partner (the mother in family 1, indicated by a solid blue line) undergoes from 2002 to 2011.

'cow'	HORNS	MILKING	DRINK	Distance (from HS)
Homesigner	1	1	0	
Sister	0	1	1	.6
Mother	1	0	0	.8
<i>Frequency</i>	2	2	1	

Table 2. Sample calculations of frequency-weighted hamming distance. Distance is the number of Conceptual Components that differ between a pair of responses, weighted by each Conceptual Components' frequency divided by the token frequency of all conceptual components. For example, the Homesigner-Sister distance is calculated as follows: $2/5*1 + 2/5*0 + 1/5*1 = .6$.

⁵ CC's used more frequently offer more opportunities for convergence, and so should arguably be weighted more heavily in calculating distance.

⁶ This measure ignores the ordering of gestures and is thus an overestimate of conventionalization.

We ran two tests to establish that (1) communication partners gradually converge with their respective homesigners, but that (2) even in 2011, convergence was not complete (where distance would be zero). To investigate our first question, we first extracted, for every partner, slopes of the linear regressions predicting homesigner-partner distance from year of testing. A one-tailed, one-sample Wilcoxon Signed Rank test on the nine slopes indicated that the median of this sample was significantly below 0 ($W=0, p < .01$), confirming the gradual convergence between homesigners and partners. To investigate our second question, we ran a series of one-tailed, one-sample Wilcoxon Rank-Sum tests on the 2011 homesigner-communication partner distances. We found that these distances, despite decreasing over time, are still significantly greater than 0; all 9 of 9 such tests are highly significant (W 's $\geq 91, p$'s $\leq .001$).

To give an interim summary, we found that homesigners are converging with individual communication partners, but that even in 2011, the last year we tested them, they have not converged completely. We now turn our attention to our second study, on Nicaraguan Sign Language.




Figure 1: Average distances, across objects/concepts tested, between a partner's lexicon and his/her associated homesigner's lexicon, per year. Partners converge with their respective homesigners. Y = Younger; O = Older. Line type (full, dotted, short dashed, long dashed) indicates homesign group membership.

Study 2 – Nicaraguan Sign Language

Method

Participants

Participants were eight (2 males; 21-32 years, $M=27$) Deaf users of Nicaraguan Sign Language. 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 from the first cohort of NSL users, and formed NSL *de novo*.

Stimuli

Because these data were collected independently of the homesign data, different elicitation materials were used. First, and most importantly, overlap in objects/concepts between the two

NSL and homesign lists was not perfect. Only nine items were shared between the two lists: cat, dog, cow, rain, sun, ice, egg, fish, and orange.

Procedure

Same as Study 1.

Coding

Coding was carried out by the author, and was otherwise identical to Study 1.

Results

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 used an analysis procedure slightly modified from Study 1: 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 2011 Homesigner-Communication Partner distances for the same set of objects (see Figure 2): a one-sample Wilcoxon Signed Rank test determined that the median homesigner-partner distance was significantly above the NSL average distance ($W=36, p < 0.01$). Given this, and that comparable periods of time had passed since the beginning of NSL and the various homesign systems, this suggests that the first cohort NSL users conventionalized faster than the various homesign systems.




Figure 2. Average distances, across objects/concepts tested. On the left, distances are between a partner’s lexicon and his/her associated homesigner’s lexicon, in 2011. On the right is the average distance among lexicons of all pairs of NSL signers tested, in 2003. Given that the median of the homesign-partner distances was significantly greater than the NSL average distance (Wilcoxon Signed Rank Test, $W=36$, $p < 0.01$), and that comparable periods of time had passed since the beginning of NSL and the various homesign systems, we conclude that NSL has conventionalized faster than the various homesign systems.

Discussion

We showed above that deaf homesigners slowly converge on form-meaning mappings with their hearing communication partners, but that convergence is not complete, even in 2011, the latest year in which we collected data. We then showed that the NSL community converged on a basic lexicon in less than 25 years. By 2011, all of the present homesigners had been using their respective systems for well more than 15 years, yet none of them had converged completely with any of their communication partners. 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, the hearing users only use the system to interact with that deaf

user. In NSL and other deaf community sign languages, however, all users of the system interact with other users of the system *using the system*. In other words, the homesign interactive structure is a star network, while the NSL/Deaf community structure is roughly a fully-connected network⁷. We now turn to our model, which replicates this observed convergence, and allows us to test these predictions.

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 & 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 conventionalization between homesign and NSL.

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

⁷ The NSL community is more likely a small-world network, as social networks tend to be of these types (Newman, 2001; 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 is not terribly inaccurate.

homesign data. For a specific meaning, agent i accesses a vector of probabilities $P_c = \{p_c^i\}$, defined over these signs ($j = 1, 2, \dots, C$) such that with probability p_c^i , the c th sign is used by agent i and with probability $(1 - p_c^i)$, the c th 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).

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. See Figure 3 for an example interaction (without the renormalization).




Figure 3. A sample conversation between two agents.

Social matrix: homesign vs language We consider two social networks of agents. In the first, analogous to the homesign situation, one individual, the deaf signer (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 typical sign language emergence (Woll & Ladd, 2003).

Results 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 3), 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.

Table 3: 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%

General Discussion

In the current work, we (1) presented longitudinal data showing conventionalization of lexicons among users of naturally emerging language-like systems (homesign gesture systems); (2) showed that conventionalization in these homesign systems proceeds more slowly than in Nicaraguan Sign Language (NSL), an emerging sign language used by a recently emerged Deaf community; (3) formulated a general framework and causal model of conventionalization, in the form of a multi-agent reinforcement learning model that obtains conventionalization; and (4) 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. We discuss implications our findings below, as well as open questions.

As mentioned earlier, how exactly language emerges from communities has been somewhat mysterious. Computational modeling (e.g. Gong et al., 2012), while suggestive, has not been well-connected to and well-motivated by empirical work, whether experimental or naturalistic, and thus has largely gone unvalidated. Experimental work, in contrast, has almost exclusively focused on dyadic interaction. Finally, the benefits of being part of a community that had been identified by work on naturalistic systems have only identified somewhat peripheral linguistic properties, with the lone exception being Osugi et al. (1999)'s work on conventionalization and community structure, which did not offer a computational or even verbal model of how conventionalization proceeds. Despite the lack of demonstrated critical differences between homesign and (emerging) natural languages, and consequently the lack of understanding of how language emerges from (naturally occurring) communities, researchers have only speculated on what properties of the social context of homesign, NSL, and other languages influence their emergence (e.g. size of community, number of generations, age of exposure, context for transmission to new generations; Senghas, 2005; Meir, Sandler, Padden, & Aronoff, 2010). Again, however, it had not been clear *how* exactly these differences might influence language emergence, largely owing to the small number of emerging languages from which to draw generalizations, and the great variety in questions, materials, and analytic tools among researchers (Meir et al., 2010). Our work—to our knowledge, the first longitudinal observation of lexicons emerging naturally, and the first comparison of empirical data of naturally emerging languages to computational models—clearly demonstrates one contribution of rich sociolinguistic communities to language emergence: faster conventionalization.

Developmental psychology and the present work

We first wish to discuss how our work is relevant to more traditional concerns of developmental psychology and cognitive science more broadly. It is true that we paid most attention here to how structure emerged at the *group* level, whereas developmental psychology and the rest of cognitive science typically focus on how structure or knowledge emerges within *individuals*. One thus might be tempted to say our work is more relevant to sociolinguistics than to developmental psychology. However, even if one's interest is in structure in the mind of the individual, the structure of the group and the way in which individuals *interact* both influence how structures arise in the mind of an individual (e.g. Tomasello, Kruger, & Ratner, 1993). To use an example from the present work, clearly conventionalization at the group level entails stabilization of the lexicon within the individual. This obtains in our model, and makes a good deal of intuitive sense: an individual will likely not change his/her own lexicon when they are already consistent with everyone around them, as doing so will drastically hurt their communicative success. And of course, the high conventionality of the lexicons in typical language communities is precisely why children in these communities stabilize on a lexicon so fast (Clark, 1980).

But even more to the point, work like the present can reveal what linguistic properties emerge from use, and so don't have to be innately specified, and what kind of learner is required for such emergence. For example, various accounts of vowel dispersion contain explicit, synchronic, teleological constraints for vowel dispersion. In other words, these accounts build constraints on vowel dispersion into the phonological grammar (e.g. Flemming, 2004). However, de Boer (2000) and Boersma and Hamann (2008) each show with simulations that dispersed vowel inventories can arise over many instances of communication between agents due to more

general constraints on perception and production (e.g. noisy perception/production, discarding vowels that lead to unsuccessful communication, merging vowel categories close together in acoustic or articulatory space, among others). Thus, it seems plausible that at least this linguistic property emerges from use and more general properties of agents, rather than innate, teleological, language-specific constraints. While conventions themselves are a rather different linguistic property, since ‘apple’ or ‘pomme’ can’t be innately specified, our work still tells us what representational repertoire (here, probabilities of gesture use) and learning regime (here, a linear-reward reinforcement learning regime) can plausibly account for the facts of conventionalization (i.e. that conventionalization happens at all, and differs between homesign and NSL). As mentioned in the introduction, the kind of learners necessary to produce human behavior at any level, be it group or individual, is precisely what developmental psychology is concerned with.

It may also be worth comparing and contrasting the homesign system lexicon conventionalization that we observed here with the lexicon acquisition observed in typical language acquisition. Of course, in both cases, individuals (and groups, in the case of conventionalization) are going from a state of having no lexicon to a state of having a lexicon. Because of this, we believe that conventionalization in language emergence is, grossly speaking, more like a baby learning language than like older children or adults coining new words (which are really drops in their lexical bucket). It’s also obvious that in both cases ultimate attainment of a lexicon depends on the quantity of input or interactions – more input or interactions leads to greater attainment of a lexicon (see Hart & Risley, 1995 for quantity of input effects in typical lexicon acquisition). However, what is equally apparent is that a stable, conventional lexicon arises much faster in typical language acquisition than in language emergence: a typical three-year-old child will know at least hundreds of words that are conventional within their community

(Hart & Risley, 1995), but the homesign families we tested apparently hadn't agreed how to sign even the 22 very basic objects we tested even after they had been using the homesign system for decades. What is uncertain, however, is whether lexicon conventionalization in emerging languages is more like lexical innovation among typical adults, where words are invented and spread more formally, i.e. with explicit consensus and in-situ feedback on a user's usage of the word (as happens in any class with technical terminology, be it sociology or linguistics), or more like lexical innovations among children, where words are invented and spread more informally.

One final point of particular developmental concern is the possibility that there may be a critical period for conventionalization, that is, that children may be more efficient at conventionalization than adults. Language acquisition of course exhibits a textbook critical period, with critical period effects observed in phonology (Flege, Yeni-Komshian, & Liu, 1999) and morphosyntax (Newport, 1990). Additionally, children have been accorded a special role in language emergence (Hudson Kam & Newport, 2009; Senghas & Coppola, 2001). The results of Hudson Kam & Newport (2009) are particularly suggestive here. They found that children regularize inconsistent determiner use in their input when learning an artificial language. Effectively, children boosted the frequency of use of the most commonly used determiner, and lowered the frequency of use of less commonly used determiners. It is plausible that something similar may happen in conventionalization. For a given object, child 'conventionalizers' may boost the frequency of use of the conceptual component(s) (or other formal primitive(s)) that is most frequently associated with that object, while dampening the frequency of use of all other conceptual components. Testing this question with naturalistic data may be difficult, as it would require finding and longitudinally observing child homesigners from a young age. Therefore, investigating this question experimentally may be more tractable.

Complex networks and human behavior

While little other work has looked at the effect of social network structure on language emergence, there has been somewhat more work on the influence of network topology on cooperative behaviors more generally (for a review, see Kearns, 2012). For example, our results accord with those of Judd, Kearns, and Vorobeychik (2010). In their game, networked human subjects each control their own choice of color, but must come to a consensus on a choice of color, but they can only see their own and their neighbors' colors. They found, as did we, that low average path lengths and high clustering coefficients (such as are found in the fully-connected NSL network) led to the fastest conventionalization. Judd et al. found the exact opposite, however, in the contrast coloring game, in which neighbors must choose different colors. While language emergence would seem to be a problem of consensus and not contrast among individuals, Judd et al.'s results point to the intriguing possibility that the effects of network topology on language emergence may depend on the domain of language under consideration (one possible problem of contrast among individuals might be users differentiating their idiolects to assert their identity). Perhaps consistent with this possibility, our results are rather different from those of Gong et al. (2012). They found that the low-clustering, high-average shortest path length star network (which our Homesign network) was superior to other high-clustering, low-average shortest path length networks—small-world, scale-free, and even fully-connected (identical to our NSL network)—in efficiency of conventionalization of categories and labels over a perceptual continuum. The precise reasons for differences in our simulation results are not clear, but as just noted, the answer very likely lies in the details of our respective models. There are at least a few possibilities. First, successful communication is a given in our model, but not in the Gong et al. Second, our meaning space has no topography; it's

just a list of objects to which one or more CC's are associated. In contrast, the Gong et al. meaning space is a one-dimensional perceptual continuum. Third, form-meaning mappings are probabilistic and compositional in our model, but discrete and holistic in the Gong et al. Whatever the reasons for the differences between our models, it will be critical to see how exactly the predictions of Gong et al. (2012)'s simulations bear out in empirical human data.

We also wish to briefly emphasize a point made by comparing our results to those of Gong et al. (2012), namely, that the positive effect of richer network structure is not obvious *a priori*. In addition to the Gong et al. results, there are other reasons to actually suspect that the star network, and other networks like it, would converge no more slowly and perhaps even more quickly than richer networks. That is, the hub of the star network could conceivably serve as the standard of form, to which all the other users adjust. No such obvious, single standard of form exists in an NSL-type network, where every node is identical in terms of its connectedness. In fact, experimental work has suggested that such role asymmetries help convergence. Selten and Warglien (2007), for example, in an experiment of language emergence, found that dyads that had a greater asymmetry of changes of forms (i.e. one partner made a greater proportion of all form changes) conventionalized to a greater degree and had greater communicative success⁸. It will be interesting to see, in both human data and computational models, how the benefits of role asymmetries, which we speculate emerge more easily in networks like the star network, trade off with the benefits of richer networks. We further speculate that the networks most efficient at conventionalization may actually be like scale-free networks or small-world networks. These

⁸ This finding may be similar to Gong et al. (2012)'s finding that greater variability in node centrality leads to faster conventionalization.

networks have low path lengths and high clustering, like fully-connected networks, yet have variability in node degree and centrality (and possibly consequently have role asymmetries)⁹.

Limitations

We acknowledge some limitations of the present work. First, the choice of parameters is perhaps not particularly well-grounded in the empirical data: no object we tested was described with 40 different conceptual components. Likewise, giving the agents one object to discuss is a simplifying if inaccurate assumption. And of course, the NSL community has far more than 5 signers. However, we strongly suspect that the parameter values chosen do not have much influence on the network effects we observed, and will soon be exploring the parameter space to test this. In addition, we have not explored the role of frequency on conventionalization (some objects are discussed more frequently than others), but the effects of frequency should be relatively straightforward – more frequently discussed objects are conventionalized more quickly.

Finally, additional explanations of the different rates of conventionalization, and of complexity in general, in homesign systems and NSL do of course exist. For example, the hearing users of the homesign system have a spoken language to communicate with, and are thus under less pressure to use and conventionalize the homesign system. This contrasts with the situation faced by the deaf homesigner and users of NSL, who can only use their signed communication system and are thus behooved to conventionalize at a greater rate. This particular explanation could also potentially be implemented in the current model by changing the learning rate γ . Likewise, other learning models, e.g. Bayesian, can be studied in the general dynamic framework of language emergence. However, in the absence of more data to test the unique

⁹ This of course would depend on each user being aware of their own and their neighbor's centrality, or at least node degree. This was only partially true in the Kearns work – users could see whether their neighbors were connected to each other, but not their neighbors overall centrality or degree.

predictions of different models, we opt here for one of the simpler possible models. These and other possibilities are not mutually exclusive and can be subject to future research.

Conventionalization and the rest of language emergence

Future work must also probe more deeply the question of how languages emerge from social interactions, and why and how particular kinds of communities are most favorable for language emergence. Here we focused on only one aspect of language: lexical conventions. Emergence of conventions at other levels of linguistic organization, e.g. in phonemic inventory (as in de Boer, 2000) or syntax, should of course also be investigated with closely integrated empirical and computational approaches. But there are other, arguably more central aspects of language, whose emergence is even less well-understood, which must be investigated. The emergence of grammatical, as opposed to lexical, linguistic forms could be one such component of language. The emergence of grammatical forms from lexical over historical time periods is of course the focus of the rich literature on grammaticalization (Bybee, 2006; Tabor, 1994), but this literature almost exclusively focuses on change in spoken languages with very long histories (but see Coppola & Senghas, 2010 on grammaticalization of deictics in Nicaraguan Sign Language). Moreover, focusing as it does on corpus analysis rather than computational modeling, this literature offers little in the way of explicit mechanistic explanations of grammaticalization¹⁰. Future work could therefore attempt to model data like that of Coppola & Senghas (2010), or perhaps better yet, data from Experimental Semiotics studies that elicit grammaticalization in the lab. A model of this phenomenon would show how interactions among agents, i.e. language *use*,

¹⁰ One exception may be Tabor (1994), who implemented a connectionist model of grammaticalization of degree modifier ‘kind-of’ from noun-prep ‘kind of’. However, this model only showed how a learner could reanalyze ‘kind of’ as a degree modifier when its syntactic distribution becomes more like that of other, *preexisting* degree modifiers. Thus, Tabor (1994) is not really a model of the emergence of grammatical forms entirely *de novo*.

could lead to the gradual emergence of grammatical forms from lexical forms. This will be a difficult but necessary task for a fuller understanding of language emergence.

Relatedly, future work could investigate how conventionalization of lexical items interacts with emergence of structural properties of language. For example, conventionalization of individual compounds (which our participants' productions may be, or may be evolving toward), and conventionalization of structural properties of compounds in general (e.g. head-modifier order), clearly can have a reciprocal influence: conventionalization of structural properties can hasten when enough individual compounds of similar properties have conventionalized, and conventionalization of individual compounds should be easier if a community agrees on the structural properties of compounds in general¹¹¹². What other interactions might there be? One possibility is that conventionalization/stabilization of a lexicon leads to increased automatization of lexical items, which in turn frees cognitive resources for use (and innovation?) of more grammatical linguistic elements (Lexical Competence Hypothesis; Hudson and Eigsti, 2003). Many grammaticalization scholars have noted a similar positive correlation between the frequency of a construction and its rate of grammaticalization (Bybee, 2006). It remains to be seen what the relationship is between the LCH hypothesis and the frequency-grammaticalization findings.

Conventionalization may also be a prerequisite for emergence of compositionality of word forms (Sandler, Aronoff, Meir, & Padden, 2011; Roberts & Galantucci, 2012). The traditional explanation of compositionality in word forms is that, as the size of a set of holistic word forms grows too large, forms become too difficult to differentiate and retrieve from the lexicon; introduction of forms composed of discrete phonemes solves this problem (Hockett, 1960;

¹¹ See Meir, Aronoff, Sandler, and Padden (2010) for an argument that structural patterns (e.g., head-modifier order) in the recently emerging Al-Sayyid Bedouin Sign Language compounds conventionalized before individual compounds.

¹² A link may be made here to syntactic bootstrapping, whereby structural properties assist in learning individual words (Naigles, 1990).

Pinker & Jackendoff, 2005). However, Sandler et al. (2011) found little evidence of compositionality of word forms in the recently emerging Al-Sayyid Bedouin Sign Language, which they argue serves the full needs of its community, which of course would require a sizeable lexicon. They argue that emergence of compositionality may require conventions (perhaps in addition to a large lexicon). According to their proposal, convention reduces reliance on iconicity for successful communication. When that happens, “formational elements themselves self-organize, under cognitive and motoric pressures for ease of articulation, formal symmetry, and the like” (pg. 573). Roberts and Galantucci (2012) found support for this proposal in an experimental semiotics study: degree of compositionality in a dyad’s invented communication system was strongly negatively correlated with its transparency to outsiders (a measure of iconicity) when set-size was partialled out. Interestingly, they found a non-significant relationship between set-size and compositionality (though they did not partial out transparency), contrary to the traditional hypothesis of emergence of compositionality. Both Sandler et al. (2011) and Roberts and Galantucci (2012) suggest that conventionality and set-size may be two different routes to compositionality, or that both may be required. While computational models of compositionality emerging from signal set-size overload have been proposed (e.g., Nowak, Krakauer, & Dress, 1999), a computational model of how compositionality might arise from decreasing iconicity and increasing conventionality has not been put forward. Developing such a computational model first, and then a model that captures both routes to compositionality, should clearly be an aim of future research.

In sum, to the extent that (1) richer networks are more favorable for conventionalization, and (2) conventionalization influences emergence of other properties of natural language, we may see that richer networks are more favorable for emergence of such properties as well.

Language emergence from communication and mere use

This discussion of what kinds of learners are necessary for language emergence brings us to another point: the extent to which language structure emerges from *successful* interaction/communication. Much of the work on how language emerges from interaction (including ours) assumes that much of language structure emerges from successful communication¹³. Our model, for example, stipulates that agents always know what ‘object’ is being talked about – this allows them to change their probability of gesture use for the object under discussion every time it is discussed. In fact, it is difficult to imagine how conventions could arise without users knowing what is under discussion. Otherwise, how could they figure out which form-meaning mappings to adjust? As knowing what object/meaning is under discussion arguably occurs if and only if communication is successful, perhaps in this sense, successful communication really does drive emergence of (at least lexical?) conventions. Successful communication is similarly key in the emergence of conventional, dispersed vowel inventories in de Boer (2000)’s simulations. For example, if a vowel’s success/use ratio drops below a certain threshold, it is discarded from the inventory. Grammaticalization would also seem to require conventionalization: it is difficult to imagine how a new, more abstract use of a form could survive if people did not agree that it could be used in that novel way. It is thus plausible that emergence of much of linguistic structure depends on successful communication.

However, successful communication can not be the whole story. Compelling evidence for this comes from Carrigan and Coppola (2012). They showed simple 1 or 2 argument videotaped events to the same four adult Nicaraguan homesigners in the current study, who then produced signed descriptions of the events. Each homesigner’s mother then watched the descriptions produced by her son or daughter and had to choose the picture (from an array of four) that

¹³ Indeed, so-called functionalist approaches to language structure are a large strand of linguistics research (e.g. Croft, 1995)

corresponded to the event described. Carrigan and Coppola found that the homesigners' mothers (the communication partners who have the longest history of communicating with the homesigners) did not understand the descriptions as well as did native ASL signers, who had never interacted with the homesigners. This suggests that there is comprehensible structure in the homesigners' productions, but since the mothers were not able to make effective use of this structure, it also suggests that successful communication with the mothers is not its source. However, we must make a subtle but very important distinction here. That is, while it seems unlikely that *successful* communication is the source of this structure, it seems much more plausible that mere *use*, or *attempted* communication, contributed to it—it seems unlikely that someone who never communicated with anyone (so, a feral child) would, in their first utterance ever, produce signs like the homesigners. Thus, one possible explanation of is that, over many attempts (successful or not) to communicate with their interlocutors, homesigners may 'discover' some way of structuring their signs that they believe will be effective at communication. Or, perhaps the structure that they alight upon through repeated use is somehow more functional for themselves. Or, perhaps the structure that they discover has no functional motivation, but still only arises after repeated use. Whatever the case, Carrigan and Coppola (2012)'s results and the current results suggest that emergence of particular properties of language likely require different ingredients: successful communication, or mere use (and, of course, human learners). Which properties require which ingredients should be one of the major questions for further research on language emergence.

Conclusion

In this work, we investigated how exactly language emerges from communities, focusing on the issue of how lexical conventions arise. We compared empirical data from four different

homesign communities to data from Nicaraguan Sign Language, showing that the latter conventionalized faster than the former. We then used an agent-based reinforcement learning model of conventionalization to show that the NSL community's richer social network (compared to that of homesign communities) may explain its faster conventionalization. Clearly, these results only scratch the surface of how language self-organizes within a community. We hope that the connections we have drawn between conventionalization and other linguistic phenomena, like grammaticalization and compositionality (key phenomena of the two general domains of language, syntax/semantics and phonology/phonetics), will prove fruitful in future research, and thus lead to a more comprehensive model of language emergence. Likewise, we hope that we have presented a compelling case as to why developmental psychologists and other cognitive scientists who are interested in the (development of) minds of individuals should still care about group-level phenomena. Finally, we hope that comparison of our work and work like it to Carrigan and Coppola (2012) and similar work makes clear that the emergence of language will require explanations that are functionalist, non-functionalist, and everything in between.

References

- Barr, D. (2004). Establishing conventional communication systems: Is common knowledge necessary? *Cognitive Science*, 28(6), 937–962. doi:10.1016/j.cogsci.2004.07.002
- Bush, R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, 68, 313-323.
- Boersma, P., & Hamann, S. (2008). The evolution of auditory dispersion in bidirectional constraint grammars. *Phonology*, 25(02), 217. doi:10.1017/S0952675708001474
- Brentari, D., Coppola, M., Mazzoni, L., & Goldin-Meadow, S. (2012). When does a system become phonological? Handshape production in gesturers, signers, and homesigners. *Natural Language and Linguistic Theory*, 30(1), 1-31.
- Bybee, J. (2006). From usage to grammar: the mind's response to repetition. *Language* 82(4): 711-733.
- Carrigan, E. & Coppola, M. (2012). Mothers do not drive structure in adult homesign systems: Evidence from comprehension. In N. Miyake, D. Peebles, & R. P. Cooper, (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 1398-1403). Sapporo, Japan: Cognitive Science Society.
- Clark, E. V. (1980) Convention and innovation in acquiring the lexicon. (Keynote Address, Twelfth Annual Child Language Research Forum, Stanford University, March 1980.) *Papers & Reports on Child Language Development [Stanford University]* 19, 1-20.
- Coppola, M., Gagne, D., Senghas, A. (2013). WHO chased the bird? Narrative cohesion in an emerging language. *Extended abstracts of the Linguistic Society of America*.
- Coppola, M., & Newport, E. 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 of the United States of America, 102(52), 19249–53.

doi:10.1073/pnas.0509306102

Coppola, M., & Senghas, A. (2010). The emergence of deixis in Nicaraguan signing. In D.

Brentari (Ed.), *Sign languages: A Cambridge Language Survey* (pp. 543–569). Cambridge,

UK: Cambridge University Press.

Croft, W. (1995). Autonomy and functionalist linguistics. *Language*, 71, 490–531.

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

Flaherty, M., & Senghas, A. (2011). Numerosity and number signs in deaf Nicaraguan adults.

Cognition, 121(3), 427–436. doi:10.1016/j.cognition.2011.07.007.Numerosity

Flege, J. E., Yeni-Komshian, G. H., & Liu, S. (1999). Age constraints on second language acquisition. *Journal of Memory and Language*, 41(1), 78–104. doi:10.1006/jmla.1999.2638

Flemming, Edward (2004). Contrast and perceptual distinctiveness. In Bruce Hayes, Robert

Kirchner, and Donca Steriade (Eds.), *Phonetically-Based Phonology*. Cambridge: Cambridge

University Press, pp. 232-276.

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

Galantucci, B., Kroos, C., & Rhodes, T. (2010). The effects of rapidity of fading on

communication systems. *Interaction Studies*, 11(1), 100–111. doi:10.1075/is.11.1.03gal

Garrod, S., Fay, N., Lee, J., Oberlander, J., & Macleod, T. (2007). Foundations of representation: where might graphical symbol systems come from? *Cognitive science*, 31(6), 961–87.

doi:10.1080/03640210701703659

- Goldin-Meadow, S. (2003). *The Resilience of Language: what gesture creating in deaf children can tell us about how all children learn language*. New York: Psychology Press.
- Goldin-Meadow, S., Butcher, C., Mylander, C. & Dodge, M. (1994). Nouns and verbs in a self-styled gesture system: What's in a name? *Cognitive Psychology*, 27, 259-319.
- Goldstone, R. L. & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, 1, 412-438.
- Gong, T., Baronchelli, A., Puglisi, A., & Loreto, V. (2012). Exploring the roles of complex networks in linguistic categorization. *Artificial life*, 18(1), 107–21. doi:10.1162/artl_a_00051
- Hart, B., & Risley, R. T. (1995). *Meaningful differences in the everyday experience of young American children*. Baltimore: Paul H. Brookes.
- Hockett, C. (1960). The origin of speech. *Scientific American* 203: 89–96.
- Hudson, C. L., & Eigsti, I. (2003). The Lexical Competence Hypothesis: A cognitive account of the relationship between vernacularization and grammatical expansion in creolization. *Journal of Pidgin and Creole Languages*, 18(1), 1–79. doi:10.1075/jpcl.18.1.02hud
- Judd, S., Kearns, M., & Vorobeychik, Y. (2010). Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences of the United States of America*, 107(34), 14978–82. doi:10.1073/pnas.1001280107
- Kearns, M. (2012). Experiments in social computation. *Communications of the ACM*, 55(10), 56. doi:10.1145/2347736.2347753
- Kirby, S., Dowman, M., & Griffiths, T. L. (2007). Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences*, 104(12), 5241-5245.

- López Gómez, J., J., A. M. Perez Castellón, J. M. Rivera Rostrán, J. F. Baltodano Baltodano (Eds.) (1997). *Diccionario del Idioma de Señas de Nicaragua* [Dictionary of Nicaraguan Sign Language]. ANSNIC [National Deaf Association of Nicaragua]. Managua, Nicaragua.
- Meir, I., Aronoff, M., Sandler, W. & Padden, C. (2010). Sign languages and compounding. In S. Scalise & I. Vogel (Eds.): *Compounding*. John Benjamins, 301-322.
- Meir, I., Sandler, W., Padden, C., & Aronoff, M. (2010). Emerging Sign Languages. In M. Marschark and P. Spencer (Eds.) *Oxford Handbook of Deaf Studies, Language, and Education*, Volume 2. Oxford: Oxford University Press, 267-280.
- Naigles, Letitia. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, 17, 357-374.
- Newman, M. E. J. (2001). [The structure of scientific collaboration networks](#). *Proc. Natl. Acad. Sci.*, 98, 404-409.
- Newport, E. (1990). Maturation constraints on language learning. *Cognitive Science*, 14(1), 11–28. doi:10.1016/0364-0213(90)90024-Q
- Niyogi, P., & Berwick, R. (1997). Evolutionary consequences of language learning. *Linguistics and Philosophy*, 20 (6), 697-719.
- Nowak, M.A. and Krakauer, D.C. (1999) The evolution of language. *Proceedings of the National Academy of Sciences*, 96, 8028–8033.
- Osugi, Y., Supalla, T., & Webb, R. (1999). The use of word elicitation to identify distinctive gestural systems on Amami Island. *Sign language & linguistics*, 2(1), 87–112. Retrieved from <http://cat.inist.fr/?aModele=afficheN&cpsid=1554456>
- Pinker, Steven, and Ray Jackendoff. (2005). The faculty of language: What’s special about it? *Cognition* 95: 201–236.

- Polich, L. (2005). *The emergence of the deaf community in Nicaragua: With sign language you can learn so much*. Washington, DC: Gallaudet University Press.
- Roberts, G. and Galantucci, B. (2012) The emergence of duality of patterning: In-sights from the laboratory. *Language and Cognition*, 4(4): 297–318.
- Sandler, W., Aronoff, M., Meir, I., & Padden, C. (2011). The gradual emergence of phonological form in a new language. *Natural language & linguistic theory*, 29(2), 503–543.
doi:10.1007/s11049-011-9128-2
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1, 143-186.
- Selten, R., & Warglien, M. (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, A., & Coppola, M. (2001). Children creating language: how Nicaraguan sign language acquired a spatial grammar. *Psychological science*, 12(4), 323–8. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11476100>
- Senghas, A. (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, A. (2005). Language emergence: Clues from a new Bedouin sign. *Current Biology*, 15(12), 463–465. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0960982205006135>
- Senghas, R. J. (1997). *An 'unspeakable, unwriteable' language: Deaf identity, language & personhood among the first cohorts of Nicaraguan signers*. (unpublished doctoral

- dissertation). University of Rochester, NY.
- Spaepen, E., Coppola, M., Spelke, E. S., Carey, S. E., & Goldin-Meadow, S. (2011). Number without a language model. *Proceedings of the National Academy of Sciences*, *108*(8), 3163–8. doi:10.1073/pnas.1015975108
- Spelke, E.S., & Newport, E.L. (1998). Nativism, empiricism, and the development of knowledge. In R.M. Lerner (Ed.), *Handbook of Child Psychology, Vol. 1: Theoretical Models of Human Development*. 5th edition. (Editor-in-Chief: William Damon). New York: Wiley.
- Swadesh, M. (1971). *The origins and diversification of language*. Sherzer, J. (Ed.), Chicago: Aldine.
- Tabor, W. (1994). *Syntactic Innovation: A Connectionist Model*, Stanford University dissertation.
- Theisen, C., Oberlander, J., & Kirby, S. (2010). Systematicity and arbitrariness in novel communication systems. *Interaction Studies*, 1971–1976.
- Tomasello, M., Kruger, A.C., Ratner, H.H. (1993). Cultural learning. *Behavioral and Brain Sciences*, *16*, 495-552.
- Watts, D.J. & Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. *Nature*, *393*, 440-442.
- Woll, B., & Ladd, P. (2003). Deaf communities. In M. Marschark & P. Spencer (Eds.), *Oxford Handbook of Deaf Studies, Language, and Education*, (pp. 151-162), Oxford: Oxford University Press.
- Yang, C. (2002). *Knowledge and learning in natural language*. New York: Oxford University Press.