

Did Social Networks Shape Language Evolution? A Multi-Agent Cognitive Simulation

David Reitter

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA, USA
reitter@cmu.edu

Christian Lebiere

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA, USA
cl@cmu.edu

Abstract

Natural language as well as other communication forms are constrained by cognitive function and evolved through a social process. Here, we examine whether human memory may be uniquely adapted to the social structures prevalent in groups, specifically small-world networks. The emergence of domain languages is simulated using an empirically evaluated ACT-R-based cognitive model of agents in a naming game played within communities. Several community structures are examined (grids, trees, random graphs and small-world networks). We present preliminary results from small-scale simulations, showing relative robustness of cognitive models to network structure.

1 Introduction

A language, even if shared among the members of a community, is hardly static. It is constantly evolving and adapting to the needs of its speakers. Adaptivity in natural language has been found at various linguistic levels. Models of *dialogue* describe how interlocutors develop representation systems in order to communicate; such systems can, for instance, be observed using referring expressions such as *the wall straight ahead* that identify locations in a maze. Experiments have shown that communities converge on a common standard for such expressions (Garrod and Doherty, 1994).

Models of the horizontal *transmission of cultural information* within generations show on a much larger scale how beliefs or communicative standards spread within a single generation of humans. Recently, language change has accelerated through the use of communication technologies, achieving changes that used to take generations in years or even months or weeks. However, the

structure of electronic networks mimics that of more traditional social networks, and even communication via mass media follows a power-law-driven network topology.

The individual agents that are effecting the language change depend on their cognitive abilities such as memory retrieval and language processing to control and accept novel communication standards. Do the local, cognitive constraints at the individual level interact with the structure of large-scale networks? Both social structure and individual cognitive systems have evolved over a long period of time, leading to the hypothesis that certain network structures are more suitable than others to convergence, given the specific human cognitive apparatus. Some properties of human cognition are well established, e.g., in cognitive frameworks (Anderson et al., 2004). Was human cognition shaped by social networks? Why are memory parameters the way they are? Social network structures may hold an answer to this question. If so, we should find that naturally occurring networks structures are uniquely suited to human learning, while others will perform less well when human learners are present.

The environment may have been influenced by individual cognition as well. Why are social networks structured the way they are? Human memory and possibly human learning strategies are the result of an evolutionary process. Social network structures can be explained by models such as *Preferential Attachment* (Barabasi and Albert, 1999), yet, even that is tied to evolved distributions of preferences in human agents. Dall’Asta et al. (2006) argue that the dynamic of agreement in small-world networks shows, at times, properties that ease the (cognitive) memory burden on the individuals. It is possible that the human memory apparatus and social preferences governing network structures have co-evolved. Such a theory would, again, suggest the hypothesis underlying

ing this study: that network structure and human memory are co-dependent.

2 Modeling Language Change

Network structure, on a small scale, does influence the evolving patterns of communication. The dichotomy between individual and community-based learning motivated experiments by Garrod et al. (2007) and Fay et al. (2010), where participants played the *Pictionary* game. In each trial of this naming game, each participant is paired up with another participant. One of them is then to make a drawing to convey a given concept out of a small set of known concepts; the other one is to select the concept from that list without engaging in verbal communication. Over time, participants develop common standards codifying those concepts: they develop a system of meaning-symbol pairs, or, *signs*. We take this system as the lexical core of the shared language. The convergence rate and the actual language developed differed as a function of the structure of the small participant communities: Fay (2010) either asked the same pairs of participants to engage in the activity repeatedly, or matched up different pairs of participants over time. Fay and Garrod's *Pictionary* experiments served as the empirical basis for a cognitive process model developed by (Reitter and Lebiere, 2009). Our model has agents propose signs by combining more elementary signs from their divergent knowledge bases, and also adopt other agent's proposals of signs for later reuse. The model, designed to match Fay's communities, was studied in a condition involving groups of eight agents, with two network structures: maximally disjoint with the same pairs of agents throughout the simulation, and maximally connected, with interactions between all possible pairs of agents.

Reitter and Lebiere's (2009) cognitive model reflects the *Pictionary* game. The model explains the convergence as a result of basic learning and memory retrieval processes, which have been well understood and made available for simulation in a cognitive modeling framework, ACT-R Anderson et al. (2004). Thus, properties of human memory and of the agent's learning strategies dictate how quickly they adopt signs or establish new signs: processes such as learning, forgetting and noise together with their fundamental parameters that are within well-established ranges provide strong con-

straints on the behavior of each agent and in turn the evolution of their communication within the network. This approach acknowledges that cultural evolution is constrained by individual learning; each agent learns according to their cognitive faculty (cf., Christiansen and Chater, 2008). With non-cognitive models, language change has been simulated on a larger scale as well (e.g., Kirby and Hurford, 2002; Brighton et al., 2005).

It is because adaptation according to experience is determined by human learning behavior that simulation in validated learning frameworks is crucial. Griffiths and Kalish (2007) for instance model language evolution through iteration among rational learners in a Bayesian framework; the purpose of the present project is to tie the simulation of language evolution to a concrete experiment and a more process-oriented cognitive architecture than the Bayesian framework. *ACT-R*'s learning mechanisms extend the Bayesian view with at least a notion of recency. Work on language processing has pointed out its relationship to memory retrieval from within the ACT-R framework, both for language comprehension (Budiu and Anderson, 2002; Lewis and Vasishth, 2005; Crescentini and Stocco, 2005; Ball et al., 2007) and for language production (Reitter, 2008). The individual language faculty as a result of biological evolution and adaptation to cultural language has been the focus of psycholinguistic models proposing specialized mechanisms (the Chomskian viewpoint); our model does not propose a specialized mechanism but rather declarative memory as store for lexical information, and procedural cognitive processes as regulators of certain communicative functions. Our multi-agent model sees part of the linguistic process as an instantiation of general cognition: the composition and retrieval of signs follows general cognitive mechanisms and can be formulated within cognitive frameworks such as ACT-R (Anderson et al., 2004) or SOAR (Laird and Rosenbloom, 1987).

In this study, we adapted the 2009 model and simulated language convergence in several larger-scale networks. We investigate the relationship between human memory function in the retrieval of linguistic items and the structure of social networks on which humans depend to communicate.

3 Network structures

Differences in naturally occurring social networks are hardly as extreme as in Fay's experiment. Some agents will be connected to a large number of other ones, while many agents will have just a few connections each. Concretely, the number of interaction partners of a randomly chosen community member is not normally distributed and centered around a mean. It shows a (Zipfian) power law distribution, with a number of hubs attracting many network neighbors, and a long tail of subjects interacting with just a few other ones each. Social networks are *small world* networks: the average distance between any two nodes in the networks is low, since many of them are connected to hubs. Non-organically connected communication and command networks follow other normals—tree graphs for instance. However, natural communication standards develop in networks that have very specific properties that can be observed in most organically developed networks.

Realistic social networks commonly show very specific properties. Social networks, in which links symbolize communication pathways or some form of social acquaintance, frequently exhibit the *small world* property. The mean minimum distance between any two nodes is relatively low, and the clustering coefficient is high (Watts and Strogatz, 1998).

Other forms of networks include tree hierarchies with a constant or variable branching factor (directed acyclic graphs). Such networks resemble communication and command hierarchies in military or business organizations. N-dimensional grid networks have nodes with constant degrees, which are connected to each of their two neighbors along each dimension in a lattice.

Much work on information or belief propagation, or decision-making in networks has used large artificial networks modeled after social ones; nodes in such networks are commonly simple agents that make decisions based on input fed to them by their neighbor nodes and pass on information. These often state-less agents do not necessarily employ learning or adaptivity, and when they do, learning does not reflect known cognitive properties of human memory. The mechanisms governing learning and retrieval in human memory have been studied in detail, leading to formal models of process that detail the units that may be stored in and retrieved from memory, the retrieval

time and accuracy depending on the frequency and recency of prior rehearsals, on contextual cues that may facilitate retrieval, and on individual differences. Cognitive agents can serve as a more realistic basis for network simulations (Sun, 2001).

Frequency, recency, contextual cues and chunking of the stored information determine retrieval probability, which is crucial when novel idioms are required to express meaning in communication. The process leads to the choice of one of several available synonyms. Our model sees this decision-making process as a matter of memory retrieval: given the desired meaning, which sign (word or drawing, compound noun or drawings) can be used to express it. This process is implicit (not consciously controlled), and it follows recent suggestions from cognitive psychology: Pickering and Garrod's (2004) Interactive Alignment Model proposes that explicit negotiation and separate models of the interlocutor's mental state aren't necessary, as long as each speaker is coherent and adapts to their interlocutors, as speakers are known to do on even simple, linguistic levels (lexical, syntactic). This shifts the weight of the task from a sophisticated reasoning device to the simpler, more constrained implicit learning mechanism of the individual.

The social network controls the interactions that the agents can experience. Each interaction is an opportunity to develop new signs and adapt the existing communication systems. It can be shown that even separate pairs of agents develop specialized communication systems, both empirically (Garrod and Doherty, 1994; Reitter and Moore, 2007; Kirby and Hurford, 2002) and in the specific model used here. When communication partners change, convergence towards a common system and the final transmission accuracy is slower (Fay et al., 2008). At this point it is unclear how the structure of the communication network and the learning process interact. Given that some types of networks show a wide distribution of degrees, where some nodes communicate much more often and with a wide variety of neighbors, while others communicate less often, recency and frequency of memory access will vary substantially. Other communication networks may reflect command hierarchies in organizations, which are constructed to ensure, among other things, more predictable information propagation.

We hypothesize that the human memory ap-

paratus and preferred social network structures have co-evolved to be uniquely suited to create a macro-organism that adapts its communication structures and reasoning mechanisms to novel situations. There is limited opportunity to test such a hypothesis under controlled conditions with a sufficiently large human network; however, cognitive models that have been developed to explain and predict human performance in isolated cognitive situations can be leveraged to study the development of sign systems.

In a simulated network with cognitive models representing agents at the network nodes, and communication between agents along network links, we expect that the social network structures lead to better, if not optimal, adaptivity during the establishment of a communication system. We expect that scale-free small world networks do best, outperforming tree hierarchies, random networks and regular grids (lattices).

3.1 Architecture

ACT-R’s memory associates symbolic chunks of information (sets of feature-value pairs) with subsymbolic, activation values. Learning occurs through the creation of such a chunk, which is then reinforced through repeated presentation, and forgotten through decay over time. The symbolic information stored in chunks is available for explicit reasoning, while the subsymbolic information moderates retrieval, both in speed and in retrieval probability. The assumption of rationality in ACT-R implies that retrievability is governed by the expectation to make use of a piece of information at a later point. Important to our application, retrieval is further aided by contextual cues. When other chunks are in use (e.g., *parliament*), they support the retrieval of related chunks (*building*).

The properties of memory retrieval in terms of time and of retrieval success are governed by the *activation* of a chunk that is to be retrieved. Three components of activation are crucial in the context of this model: *base-level activation*, *spreading activation* and *transient noise* (ϵ). Base-level activation is predictive of retrieval probability independent of the concurrent context. It is determined by the frequency and recency of use of the particular chunk, with t_j indicating the time elapsed since use k of the chunk. d indicates a base-level decay parameter, usually 0.5):




Figure 1: Example of a small ontology with abstract concepts (spelled-out words) and concrete ones (drawings).

$$A_i = \log \sum_{k=1}^{pres} t_k^{-d} + \sum_j^{cues} w_j S_{ji} + \epsilon$$

Retrieval is contextualized by cues available through spreading activation. It is proportional to the strengths of association (S_{ji}) of all of the cues with the target chunk. While the base-level term (first term of the sum) can be seen as a prior, spreading activation models the conditional probability of retrieval given the available cues. Finally, ϵ is sampled from a logistic distribution shaped by canonical parameters. A_i must surpass a minimum *retrieval threshold*.

The model is implemented using the ACT-UP toolbox, which makes the components of the ACT-R theory are directly accessible. The cognitive model does not specify other model components (perceptual, manual, procedural), as they are neither subject to evaluation nor considered to make a significant contribution to learning or convergence effects.

3.2 Communication model

We assume that the communication system, or *language*, is a *system of signs*. Concretely, it is a set of tuples (*signs*), each associating a *meaning* with a set of up to three *symbols* (a simplifying assumption). If the communication system uses natural language, symbols consist of spoken or written words. The communication system established by the participants of Garrod’s and Fay’s

experiments uses drawings as symbols—the principle stays the same. Agents start out with a knowledge base containing signs for concrete concepts that are immediately representable as drawings or nouns; the target concepts to be conveyed by the participants, however, are more abstract and require the combination of such concrete concepts. A concept such as *hospital*, for instance, could involve the drawings for *house*, *ambulance*, and a *sad face*. A participant could choose among many ways to express *hospital*.

The goal of our cognitive models is to communicate meaning from one agent to another one. Put in natural language-oriented terminology, the *director* role is the *speaker*, a role that involves selecting the right concrete concepts that can express a given target concepts; the *matcher* role (*listener*) involves decoding the concrete drawings (or words) to retrieve the target.

A single ACT-R model implements the *director* and *matcher* roles. As a director, the model establishes new combinations of drawings for given target concepts. As a matcher, the model makes guesses. In each role, the model revises its internal mappings between drawings and target concepts. The model is copied to instantiate a community of agents, one for each node in the network.

The simplest form of representing a communication system in ACT-R memory *chunks* is as a set of signs. Each sign pairs a concept with a set of drawings. Competing signs can be used to assign multiple drawings for one concept. To reflect semantic relationships, we need to introduce a sub-symbolic notion of relatedness. We use ACT-R's spreading activation mechanism and weights between concepts to reflect relatedness. Spreading activation facilitates retrieval of a chunk if the current context offers cues related to the chunk. Relatedness is expressed as a value in log-odds space (S_{ji} values).

When the model is faced with the task to draw a given concept such as *Russell Crowe* (one of the concepts in the experiment) or *Hospital* (as in Figure 1) that has no canonical form as a drawing, a related but concrete concept is retrieved from declarative memory (such as *Syringe* in the example). In drawing-based communication, this would be a concept that can be drawn, while in natural-language based communication, this is an existing drawing expressing a similar, partial or otherwise related concept. We request two other such con-

cepts, reflecting the desire of the communicator to come up with a distinctive rather than just fitting depiction of the target concept. The case of a model recognizing a novel combination of drawings is similar; we retrieve the concept using the drawings as cues that spread activation, making the target concept the one that is the most related one to the drawings.

After drawings have been produced or recognized and mapped to a target, the target or guessed concept, along with the component drawings, is stored symbolically in memory as a chunk for later reuse (*domain sign*). These signs differ from the pre-existing concepts in the network, although they also allow for the retrieval of suitable drawings given a concept, and for a concept given some drawings. When drawing or recognizing at a later stage, the memorized domain signs are strictly preferred as a strategy over the retrieval of related concepts. The system of domain signs encodes what is agreed upon as a language system between two communicators; they will be reused readily during drawing when interacting with a new partner, but they will be of only limited use when attempting to recognize a drawing combination that adheres to somebody else's independently developed communication system.

Thus, the model has two avenues to express and recognize an abstract concept: by associative retrieval and by idiomatic domain concept. A message constructed by domain concept retrieval is often decoded by the matcher by association, and vice versa.

The identification accuracy of the model shows characteristics observed in empirical work (Fay et al. 2008). See Reitter and Lebiere (subm) for a detailed description of the model and its evaluation.

3.3 Knowledge

Agents start out with shared world knowledge. This is expressed as a network of concepts, connected by weighted links (S_{ji}). The distribution of link strengths is important in this context, as it determines how easily we can find drawing combinations that reliably express target concepts. Thus, the S_{ji} were sampled randomly from an empirical distribution: log-odds derived from the frequencies of collocations found in text corpus data. From the *Wall Street Journal* corpus we extracted and counted pairs of nouns that co-occurred in the same sentence (e.g., “market”, “plunge”). As ex-




Figure 2: Identification accuracy for isolated pairs and communities: (a) human data as provided by Fay (p.c.), (b) simulation. One-tailed standard-error based 95% confidence intervals (upper bounds for communities, lower bounds for pairs) for human data; two-tailed 95% via bootstrapping for simulations. As in the human data, both community pairs and isolated pairs converge most in the early rounds, but community pairs lose much accuracy when switching partners.

pected, the frequencies of such collocations are distributed according to a power law.

Such knowledge is, however, not fully shared between agents. Each agent has their own knowledge network resulting from life experience. This difference is essential to the difficulty of the task: if all agents came to the same conclusions about the strongest representation of target concepts, there would be little need to establish the domain language. We control the noise applied to the link strengths between concepts j and i for agent M (S_{ji}^M) by combining the common ground S_{ji} (shared between all agents) with a random sample N_{ji}^M in a mixture model: $S_{ji}^M = (1 - n)S_{ji} + nN_{ji}^M$; sign identification accuracy was found to be stable for n up to about 0.4; we set it to 0.3 for Simulation 1.

4 Simulation 1

Networks of individual cognitive agents were created to differentiate performance between four different network structures. **Random networks** contain N nodes with randomly assigned links

between them, on average d links for each node (Erdős and Rényi, 1959). **n -dimensional Grids** contain N nodes with a constant number of links d per node, with links between neighbors along each dimension. The width w is kept the same along each dimension, i.e. there are w nodes per row. We use 6-dimensional lattices. **Trees** are directed acyclic graphs with 1 link leading up, and $d - 1$ links (branching factor) leading down the hierarchy of a total of N nodes. **Scale-free networks** are constructed using the *preferential attachment* method as follows (Barabasi and Albert, 1999). N nodes are created and each is connected to one randomly selected other node. Then, two links $\langle a, b \rangle$ and $\langle a', b' \rangle$ are chosen randomly out of the existing set of links, and a new link $\langle a, b' \rangle$ is added, until the mean degree d (links per node) is reached. Preferential attachment ensures that nodes with a high number of links acquire further links more quickly than other nodes (*the rich get richer*). This yields a power-law distribution of degrees. Our scale-free networks display small world properties.

For the first Simulation, we control N at 85 and d at 5¹. 35 iterations were simulated in each trial; 20 trials were run. During each round, each agent (network node) plays one game (16 concepts) with one of its neighbors. The order of neighbors is shuffled initially, but constant across the rounds. A variable *Round* coded iterations from 1 to 35.

Results Figure 3 shows the learning curve for agent pairs in the four networks. Agents in all networks converge. Confidence intervals obtained via bootstrapping indicated no apparent differences at any specific iteration. A linear model was fitted estimating the effects of network type overall (as a baseline) for each of the four types. It also fitted interactions of iteration (1–35) with the network types, which indicate significant learning effects as follows. For each network type, we found a significant learning effect (effect of *Round*) (β 0.002, $p < 0.001$).

Planned comparisons of the learning rate in Small World networks revealed no difference with either of the other three network types ($p > 0.3$).

¹We found that networks need to be sufficiently large to display meaningful differences in community structure. The sizes were chosen to be computationally feasible (4h/CPU core per network).




Figure 3: Identification accuracy between connected agents for communities of different network structures.

5 Simulation 2

The success of a community is not only determined by how successfully individuals communicate in their local environment, that is, with their network neighbors. Communities require communicative success outside of well-acquainted agents. Agents' languages would ideally converge on a global scale. One way to test this is to have randomly paired agents play the Pictionary game at regular intervals throughout the game and thus measure identification accuracy outside of the network that defines the social structure.

This simulation was identical to Simulation 1, except that we scaled up the simulation to examine whether the lack of effect was possibly due to size or density of the nodes ($N = 512, d = 6$, noise level: 0.2, repetitions: 20). In this simulation, we measured ID accuracy between pairs of randomly chosen agents after each round. For three network types, *Grid*, *Small World* and *Random* we found significant interactions with *round*, i.e. significant convergence, (all $\beta > 0.016, z > 2.1, p < 0.05$). For the network type *Tree* we found no significant interaction ($\beta = 0.012, z = 1.55, p = 0.12$).²

²All regressions in this simulation were (generalized) mixed-effects models, with ID accuracy as response via logit link, *Round* as predictor, and *Condition* as factor for four network types. A random intercept was fitted, grouped by repetition (1–20), to account for repeated measures. The predictor was centered; no substantial collinearity remained. The analysis of Simulation 1 was a simple linear model; ID accuracy




Figure 4: (Aggregate) Identification accuracy between random agent pairs for communities of different network structures.

To test the initial hypothesis, we re-coded the conditions with a *SmallWorld* factor, contrasting the small world networks with all other conditions. We found an effect of *Round* ($\beta = 0.017, z = 3.66, p < 0.001$), indicating convergence, but no interaction with *SmallWorld* ($\beta = -0.00027, z = -0.03, p = 0.98$).³

Results Figure 4 shows network-global convergence. Again, a linear model was fitted to estimate the learning rate in different network types (interaction of network type and iteration) (baseline intercepts were fitted for each network type). We found significant interactions with iteration for the following network types: *Grid* ($\beta = 0.004, p < 0.001$), *Small World* ($\beta = 0.003, p < 0.01$), and *Random* ($\beta = 0.003, p < 0.005$), but not for *Tree* ($p = 0.991$).

Planned comparisons revealed an interaction of network type and iteration for *Tree* compared to *Small World* ($\beta = -0.003, p < 0.05$), but not for *Grid* nor *Random* compared to *Small World* ($p > 0.35$). This indicates slower across-network convergence for trees than for small worlds. It also suggests that convergence across the network does not differ much between grids, random networks and small worlds.

was, for all levels, not near either extreme ($\mu = 0.77$).

³Further, unreported, experiments, showed a similar picture with a smaller network as in Simulation 1.

6 Discussion

We find that convergence is relatively stable across the four network types. Analyzing the differences between the networks, we find that the average degree, which was controlled for grids, random networks and small worlds, was substantially lower for trees ($d = 1.9$) due to the large number of leaves with degree 1. This (or the correlated algebraic connectivity of the network) may prove to be a deciding correlate with cross-network convergence. Other metrics, such as the clustering coefficient (Watts and Strogatz, 1998), which gives an indication of the degree of neighborhood cohesion

We see these results still as preliminary. More work needs to be done to investigate how well learning scales with network growth, and how network analytics such as clustering coefficients affect the dispersion of information.

Further work will explore range of networks and the possibly unique suitability of human learning mechanisms to succeed in such networks. We will explore the (subsymbolic) parameters governing adaptation, and to what extent the quantitative parameters we find universal to humans are substantially optimized to deal with the small-world networks and pareto degree-distributions found in human communities.

7 Conclusion

Cognition may appear to be adapted to the social structures prevalent in communities of flocks, packs and human teams. There are many reasons why such social structures themselves could have evolved; if cognitive constraints play a role, we expect it to be only a small factor among many. The present simulation results certainly do not support this view: they are much more compatible with a *humans-as-generalists* theory that proposes that humans have evolved to handle a variety of network structures well, or that their recency- and frequency-based learning mechanism is not specialized.

Learning, if adapted to social structure in any way, may go beyond the current, mechanistic and implicit mechanisms implemented in ACT-R and comparable theories: learning may rely on more explicit strategies, analyzing one's interaction partners and their current knowledge, and it needs to judge information according to its sources (*trust*). Meta-cognition could also play a role in determining when a set of signs is substantially

novel and better than the current system, and thus worth enduring the cost of switching from a settled set of language conventions.

We have evaluated only a small, initial part of a co-evolution theory we proposed. Also, the problem we describe may be best operationalized at a higher abstraction level: *Consensus problems* and information spread have been intensively studied (e.g., Latora and Marchiori, 2001; Wu et al., 2004). Comparing community convergence in a number of differently-structured networks, so far we see little evidence supporting our hypothesis, namely that cognition (memory) has specialized to accommodate social structures as defined by contemporary network science, and that those structures accommodate cognitive properties. Instead, we find that the simulated cognitive agents converge in their communication systems quite well regardless of the network structures, at least as long as those networks are relatively small and of similar average degrees.

Acknowledgments

This work was funded by the Air Force Office of Scientific Research (MURI grant FA95500810356).

References

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Quin, Y. (2004). An integrated theory of mind. *Psychological Review*, 111:1036–1060.
- Ball, J., Heiberg, A., and Silber, R. (2007). Toward a large-scale model of language comprehension in act-r 6. In *Proceedings of the 8th International Conference on Cognitive Modeling*, Ann Arbor, MI.
- Barabasi, A. L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.
- Brighton, H., Smith, K., and Kirby, S. (2005). Language as an evolutionary system. *Physics of Life Reviews*, 2(3):177–226.
- Budiu, R. and Anderson, J. R. (2002). Comprehending anaphoric metaphors. *Memory & Cognition*, 30:158–165.
- Christiansen, M. H. and Chater, N. (2008). Language as shaped by the brain. *Behavioral and Brain Sciences*, 31(5):489–509.

- Crescentini, C. and Stocco, A. (2005). Agrammatism as a failure in the lexical activation process. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*.
- Dall'Asta, L., Baronchelli, A., Barrat, A., and Loreto, V. (2006). Agreement dynamics on small-world networks. *EPL (Europhysics Letters)*, 73(6):969.
- Erdős, P. and Rényi, A. (1959). On random graphs. I. *Publ. Math. Debrecen*, 6:290–297.
- Fay, N., Garrod, S., and Roberts, L. (2008). The fitness and functionality of culturally evolved communication systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509):3553–3561.
- Fay, N., Garrod, S., Roberts, L., and Swoboda, N. (2010). The interactive evolution of human communication systems. *Cognitive Science*, 34(3):351–386.
- Garrod, S. and Doherty, G. M. (1994). Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, 53:181–215.
- Garrod, S., Fay, N., Lee, J., Oberlander, J., and Macleod, T. (2007). Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science*, 31(6):961–987.
- Griffiths, T. L. and Kalish, M. L. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive Science*, 31(3):441–480.
- Kirby, S. and Hurford, J. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*, chapter 6, pages 121–148. Springer Verlag, London.
- Laird, J. E. and Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1):1–64.
- Latora, V. and Marchiori, M. (2001). Efficient behavior of small-world networks. *Phys. Rev. Lett.*, 87(19):198701.
- Lewis, R. L. and Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29:1–45.
- Pickering, M. J. and Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27:169–225.
- Reitter, D. (2008). *Context Effects in Language Production: Models of Syntactic Priming in Dialogue Corpora*. PhD thesis, University of Edinburgh.
- Reitter, D. and Lebiere, C. (2009). Towards explaining the evolution of domain languages with cognitive simulation. In *Proceedings of the 9th International Conference on Cognitive Modeling (ICCM)*, Manchester, UK.
- Reitter, D. and Lebiere, C. (subm.). Towards explaining the evolution of domain languages with cognitive simulation. *Cognitive Systems Research*.
- Reitter, D. and Moore, J. D. (2007). Predicting success in dialogue. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics (ACL)*, pages 808–815, Prague, Czech Republic.
- Steedman, M. (2000). *The Syntactic Process*. MIT Press, Cambridge, MA.
- Sun, R. (2001). Cognitive science meets multi-agent systems: A prolegomenon. *Philosophical Psychology*, 14(1):5–28.
- Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.
- Wu, F., Huberman, B. A., Adamic, L. A., and Tyler, J. R. (2004). Information flow in social groups. *Physica A: Statistical and Theoretical Physics*, 337(1-2):327 – 335.