Title

Social Learning in Complex Networks: The Role of Building Blocks and Environmental Change

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN
1069-7977

Authors
Barkoczi, Daniel
Galesic, Mirta

Publication Date
2013

Peer reviewed

Social Learning in Complex Networks: The Role of Building Blocks and Environmental Change

Daniel Barkoczi (barkoczi@mpib-berlin.mpg.de)
Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany

Mirta Galesic (galesic@mpib-berlin.mpg.de)
Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany

Abstract

We explore the interaction between information sampling and the structure of the social environment in the case of two prominent social learning strategies: *imitate-the-best* and *imitate-the-majority*. In a series of simulations a group of agents made repeated choices between options. We varied the building blocks of the strategies used by agents, the structure of the social network and characteristics of the task environment. A key factor influencing strategies' success is the speed with which they are able to respond to environmental change. In general, *imitate-the-best* provides a faster response compared to *imitate-the-majority* and larger samples help the former but hurt the latter. Less efficient networks decrease the performance of both, but are more detrimental for *imitate-the-majority*. Our findings highlight the role of sampling and social structure in the study of social learning, an area not sufficiently explored before.

Keywords: Social learning; information sampling; social networks; simple heuristics; simulation; decision-making

Introduction

Humans and other animals obtain information via social learning. This is an efficient way to save the time and effort involved in individual trial-and-error learning and is known to underlie our capacity for culture. Despite the diverse list of empirical evidence for its use in the wild (Laland, 2004; McElreath, et al. 2008), theoretical models exploring the adaptive nature of social learning strategies lack sufficient detail to explain when we should expect to observe them. Most models study unstructured groups and focus only on the decision phase of implementing a strategy (e.g. *imitate-the-majority*), leaving open an important dimension affecting strategy performance: the interaction between information sampling and the structure of the social environment. The present study is an attempt towards filling this gap in the literature.

Social learning is often based on limited samples of the social environment. Most communities consist of sizable groups where an individual cannot survey all other group members within reasonable time before making a decision. Consider migrating animals deciding between multiple directions, individuals in an organization trying to jointly solve a problem or stock traders trying to predict the best investment option (Couzin, Krause, Franks & Levin, 2005; March, 1991). In such situations the way information about options is sampled from the social environment is likely to be an important aspect of any strategy. The structure of the social network in which social learning takes place can then in turn affect the options available for sampling. Previous work has shown that different network structures and their efficiency can affect the diversity of options in the population and the time it takes groups to converge on a solution (Lazer & Friedman, 2007; Mason & Watts, 2012). How does the performance of different strategies depend on the way they sample information and on the social environment in which they are embedded?

To address this question we study two representative social learning strategies: *imitate-the-best* and *imitate-the-majority* (Boyd & Richerson, 1985; Laland, 2004) and model them as decision heuristics that consist of different building blocks: search, stop and decision rules (Gigerenzer, Todd & the ABC Research Group, 1999). By explicitly modeling these three phases we are able to test their relative contribution to strategy success in different social environmental structures.

Overall, a general characteristic shared by many social learning strategies, including those we study here, is that they alter the structure of the social environment by increasing the frequency of the correct option (i.e. the one with the highest payoff) and simultaneously decreasing the diversity of options in the group. This is a result of their bias towards specific sources (best member, majority) and their selectiveness (e.g. copy only if payoff better)[1]. This property has been extensively studied in the context of biased cultural transmission (Boyd & Richerson, 1985) and suggests a key factor influencing strategy success in a changing environment: the speed with which they increase the frequency of the correct option in the group and, therefore, their ability to respond to environmental change. Our goal here is to show how this speed can be influenced by the strategy's building blocks (their sampling and decision rule) and by the structure and efficiency of the social network.

In what follows we derive specific expectations, based on previous literature and preliminary analytic calculations,

[1] One can relax this assumption if other selective forces (e.g. natural selection) are at work.

about the effects of different building blocks and network structures on *imitate-the-best* and *imitate-the-majority*.

We consider a hypothetical situation where a group of agents make repeated choices between two options (one correct, the other incorrect). Whenever the environment changes, the previously correct option becomes incorrect and vice versa.

Effects of decision rules. In general, as long as the correct option is used by the majority of agents in a group and the environment is stable, both *imitate-the-best* and *imitate-the-majority* will converge to the correct option. However, under the assumption that the best member can be reliably identified within the sample, the *imitate-the-best* will always converge faster because it requires only a single agent with the correct solution to reach a decision, whereas *imitate-the-majority* requires at least two out of three. As soon as the environment changes, the correct option will be in minority. In this case, *imitate-the-best* will still be able to find it, however, as predicted by the Condorcet Jury Theorem (CJT), *imitate-the-majority* will never find the correct option because it requires that the proportion of agents with the correct option be higher than 0.5 (e.g. Grofman, Owen & Feld, 1983).

Effects of information sampling and sample size. The CJT prediction may no longer hold when sampling is involved. Even if the correct option is in minority, *imitate-the-majority* may still be able to find it. Sampling as opposed to group-level aggregation can create situations where the correct option is more frequent in one's sample than overall in the group. When agents with such samples choose the correct option, this further increases the correct option's frequency in the group as a result of the environment altering feature of social learning discussed earlier (Boyd & Richerson, 1985). Smaller samples are more likely to produce such situations, both because they are more likely to be biased and because they require fewer agents with the correct option in order to reach a decision. This suggests two situations where smaller as opposed to larger samples should benefit *imitate-the-majority*. First, whenever the group is converging towards the incorrect option, smaller samples will delay this process and keep the payoffs of the group higher for the longer time. Second, when the correct option is in minority, smaller samples will make it more likely to accidentally have a majority of agents with the correct option. In contrast, for *imitate-the-best* larger samples are always more advantageous, because they increase the chance of finding at least one agent with the correct solution.

Effects of network structure. Previous studies have demonstrated that higher network efficiency increases the speed with which information spreads and consequently decreases the diversity of information in the group. More efficient networks should, therefore, favor all strategies. Network efficiency depends on a variety of factors (Mason & Watts, 2012); here we focus on clustering and average path length. As networks become more clustered and average path lengths increase, their efficiency decreases, and they maintain diversity for a longer time (Lazer & Friedman, 2007). We hypothesize that in such networks, the speed with which different strategies can find the correct option will become more important. As a result, the difference in speed between *imitate-the-best* and *imitate-the-majority* should become even larger. More clustered networks could have an additional effect by enabling the occurrence of relatively homogeneous clusters using the same option. If this option is incorrect, *imitate-the-majority* using a sample within that cluster will not be able to find the correct option. In contrast, *imitate-the-best* should be less affected by diversity of information as it only requires a single agent with the correct option.

Method

Overview

We simulated a situation where multiple agents (N=100) had to make repeated choices between different number of options by acquiring information from their contacts. The choices they made directly affected their payoffs.

We created three social networks differing in their efficiency (as measured by clustering and average path length). Each agent had the same number of contacts in the network (d=10) and was assigned one of four decision strategies. Each strategy sampled randomly among one's contacts but differed in its stopping and decision rule. The agents' task was to make repeated choices between different number of options (2 or 10) at each time-step using their decision strategy. The environment could change on each time-step (t_i) with some probability (p_c) affecting the payoff of options at the next time-step (t_{i+1}). The simulation was run for t=1000 time-steps and each condition was replicated 30 times[2]. To evaluate the performance of different strategies we tested them both in isolation and in an evolutionary competition where better performing strategies could replace worse performing ones. More specifically the simulation consisted of the following steps:

1) at $t=0$ agents were placed in the networks and randomly assigned a decision strategy and an initial option
2) from $t=1$ onwards, agents sampled the options and corresponding payoffs at t_{i-1} of their contacts
3) made a choice between sampled options based on their decision rules
4) only in the evolutionary competition: switched strategies with a small probability (introduced from t=50)
5) the environment changed with a certain probability, leading to a different option with the highest payoff
6) payoffs for the choice from step 3) were determined

Note that there is a lag between the information acquired from contacts and the realization of the agent's payoff in the sense that information is collected before environmental change occurs, thus allowing for the possibility of acquiring

[2] Sensitivity analyses revealed that running the simulation for 2000 time-steps and 60 replications produced identical results.

outdated information when the environment changes to a new state.

Decision strategies

We studied four decision strategies that differed in their building blocks (see Table 1). For each strategy we assumed that agents sample among their contacts randomly, and stop after collecting either a small (n=4) or a large sample (n=10)[3]. They then decide to try an option that is either endorsed by the majority of the sample contacts or by the agent that had the best payoff in the last time-step. In all cases agents only switch to a new option if that option's payoff was higher at the previous time-step than the option they are currently using. In situations where these two payoffs are equal or when the majority rule results in ties, agents chose randomly.

Table 1: Decision strategies

Sampling rule	Stopping rule	Decision rule
random sample of contacts	n=4	*imitate-the-majority*
	n=10	*imitate-the-majority*
	n=4	*imitate-the-best*
	n=10	*imitate-the-best*

In order to keep track of a changing environment any social learning strategy requires that there is some form of individual learning generating novel options, therefore, we allowed new information to enter the population through copying error, a parameter we fixed at p_e=0.01. That is, on each step there was a 0.01 chance that the agent does not consider the option used by its contacts, but a randomly selected option, however, agents only switched to this option if it had a higher payoff at the previous time-step. This lies in contrast with other studies which allowed new information to enter the group by assuming that whenever other agents' payoffs are lower or equal, the agent does not stick with its own option but explores other options randomly (Lazer & Friedman, 2007; Mason & Watts, 2012). These studies, therefore, allowed for a higher amount of innovation than our model. In this way we explore the performance of social learning strategies when aided with only a minimum amount of individual learning.

Decision environment

Two factors affecting the decision environment were varied in different simulations: a) the number of options available and b) the rate of environmental change. To manipulate the first factor we assumed that agents choose either between 2 or 10 options with payoffs ranging from 1 to 2 and from 1 to 10 respectively. At any given time, only one option had the highest payoff. On the first time-step agents were assigned options randomly. In conditions with 2 options, we varied the initial proportion of the correct option in the group (p_{init}=0.2, 0.5 or 0.7). For 10 options each option had the same initial proportion. For the second factor we assumed that the payoffs of options can change on each time-step with probability p_c=0.001, 0.01, 0.1 or 0.4 reflecting a discreet scale between slow and fast rates of change. We ran all possible combinations of environmental change on all 3 network structures described below.

Network structure

Three different networks were created, ranging from most efficient to least efficient as measured by two standard indicators in the network science literature (Mason & Watts, 2012): clustering coefficient and average path length. The clustering coefficient measures the extent to which the network is dominated by isolated cliques, which from a communication perspective decreases the efficiency of a network by making it harder for information to spread the higher the clustering. Consider an example where small groups of tightly connected agents exchange information but because groups are isolated from other groups information spreads much slower between these small units.

Another measure of efficiency is average path length, the average number of steps it takes to get from any agent to any other agent in the network. The shorter the path length the more easily information can spread. The efficiency of a network is known to affect how quickly information spreads from one part to another, however, it can also enable maladaptive information to spread more rapidly as in the case of panics following flu pandemics or stock bubbles. Many real-world networks are known to have both high clustering and low average path lengths thus representing an intermediate level of efficiency. These small-world networks (Watts & Strogatz, 1998) can be mimicked by performing random re-wirings on edges of a *lattice*. In line with previous studies (e.g. Schwenk & Reimer, 2007), we started by first generating a random directed lattice and then rewired it with a 0.1 probability to obtain a small-world network[4]. In addition we created a fully-connected network absent of any structural properties to be able to compare to previous studies that focused on unstructured groups (see Table 2). All three networks had a fixed degree of 10 and a total of 100 nodes (d=10, n=100).

[3] Sensitivity analyses with sample sizes n=3 and n=9 produced similar results and we do not report them here.

[4] Other networks with lower values of rewiring produce similar results, therefore, we omit them.

Table 2: Types of networks used in the simulation (n=100, d=10)

Network	Clustering coefficient	Average path length	Rewiring probability
Lattice	0.67	5.55	p=0
Small world	0.31	2.35	p=0.1
Fully connected	1	1	p=1

Evolutionary competition

In order to properly evaluate each strategy we look at their performance both in isolation (in homogeneous groups using the same strategy) and by directly testing different strategies against each other (heterogeneous groups) in an evolutionary competition. In the former we are interested in isolating the factors contributing to the success of different strategies, whereas in the latter we wish to evaluate them in a competitive setting where the performance of a strategy can depend on the strategies used by other agents in the group. Evolutionary competitions are a popular method in the study of social learning (e.g. Rendell, et al. 2010) where the strategy accumulating the highest payoff has the best chance of reproducing and spreading in the population, while the worst performing strategies die out. The prevalence of a strategy is, therefore, a clear-cut measure of its success in a given environment.

There are many ways to implement an evolutionary dynamic. Here we chose the ‘imitation process’ (Nowak, 2006) in order to reflect a plausible real-life scenario. We assumed that on each time-step, randomly selected agents change their strategies to one of their contacts’ strategy with a probability proportional to the cumulative payoff of that contact. If none of the contacts has a higher payoff, the agent keeps its strategy, and in situations of equal payoff random choice is implemented. We fixed the parameter specifying the probability of strategy change to p_s=0.02 thus expecting 2 agents switching strategies on each time-step. Evolutionary dynamics were introduced from the t=50 time-step to allow for a burn-in period.

Simulation results

Figure 1A shows the overall performance of the four different decision strategies observed in isolation, measured by their rate of environmental tracking (percentage of agents using the correct option on each time-step). We show the results for 2 options, probability of environmental change p_c=0.001, and initial probability of correct option p_{init}=0.5, averaged across networks[5]. To make the main results easier to view, we focus on the time-steps before and after environmental change occurring at t=100. Figure 1B shows the frequency of different strategies in the evolutionary competition averaged across networks for the same environmental condition. Overall, *imitate-the-best* consistently outperforms *imitate-the-majority* both in homogeneous and heterogeneous groups. This result holds in all network structures and in environmental conditions.

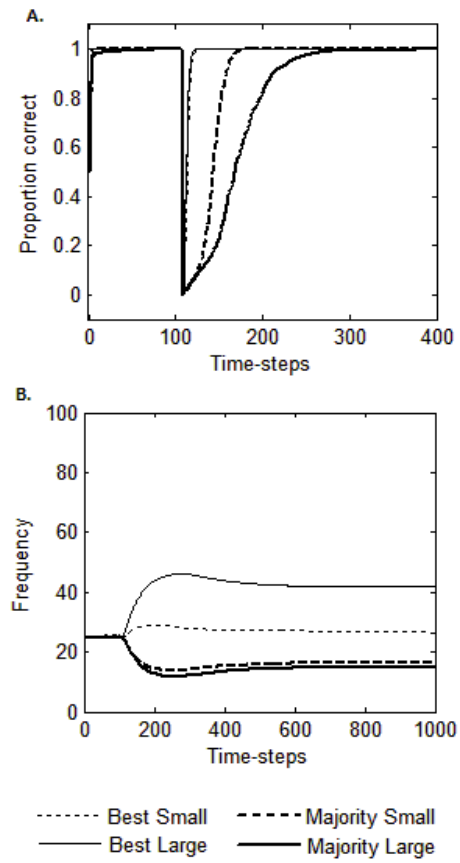

Figure 1. Panel A. Performance of strategies observed in isolation. Panel B. Frequency of strategies in the evolutionary competition. Results are shown for environmental conditions p_c=0.001 and p_{init}=0.5, averaged across networks.

Effects of information sampling

From Figure 1A we can see the number of time-steps it takes groups using each of the strategies to converge on the correct solution after the environment has changed. As expected, *imitate-the-best* benefits somewhat from larger samples, however, even its small sample version outperforms both versions of *imitate-the-majority*. The opposite is the case for *imitate-the-majority,* which is hurt by larger samples and actually performs better when it samples fewer people. This result highlights that speed with which different strategies can recover after environmental

[5] Results for 10 options and other rates of environmental change and initial probability of correct option do not change the main conclusions and we do not present them here.

change is crucial to their success and demonstrates that different sampling regimes should be adopted depending on the decision rule used.

As mentioned before, without sampling, *imitate-the-majority* will converge on an incorrect option whenever the proportion of agents using the correct option is smaller than 0.5. As expected, these results do not hold when decisions are based on sampled information as opposed to overall group aggregation. As visible in Figure 1A, *imitate-the-majority* is able to find the correct option even when the proportion of agents using it falls under 0.5. As a sensitivity check, we reran our simulations with p_{init}=0.2 and copying error of p_e=0, allowing no new information to enter the population. Even then, *imitate-the-majority* can still converge on the correct option, in particular when it uses small samples.

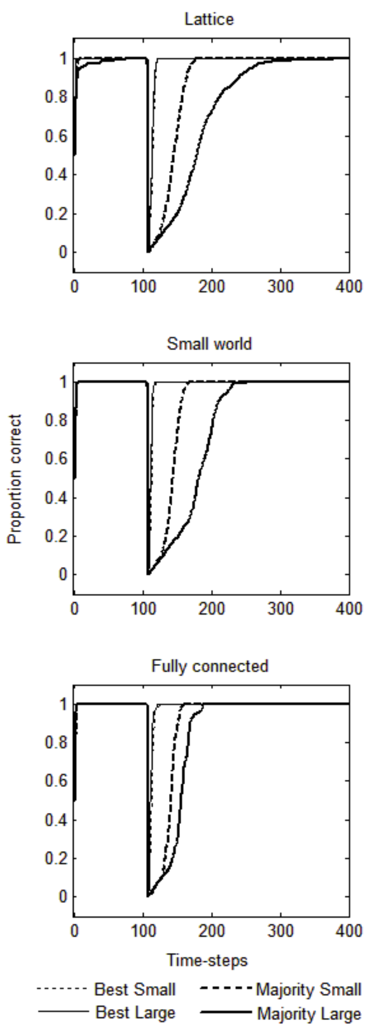

Figure 2. Performance of different strategies in the three network structures. Results are shown for environmental conditions p_c=0.001 and p_{init}=0.5

Effects of network structure

Overall, we find that regardless of strategy, more efficient networks are faster at spreading information and that this helps groups in all conditions. However, we observe an effect for network structure on the relative difference between strategies. Figure 2 shows that the difference between strategies is least pronounced in the fully connected network absent of any structural properties, however, as networks become more structured (thereby decreasing the efficiency and speed with which information flows), the difference between *imitate-the-best* and *imitate-the-majority* becomes more pronounced.

The effect of network structure is especially visible immediately after environmental change. In networks with high clustering and long path lengths such as lattice, relatively isolated agents may form homogeneous groups possessing the same information. In these situations, *imitate-the-majority* has problems finding the correct option. The larger the sample, the more prone is this strategy to get stuck. As expected, the performance of *imitate-the-best* is less affected by network structure.

Discussion

Our goal was to study how information sampling and the structure of the social environment affect the performance of two representative social learning strategies: *imitate-the-best* and *imitate-the-majority*. We modeled social learning strategies as heuristics consisting of different building blocks and embedded them in three social networks in a task involving repeated choices between multiple options.

Overall, we find that *imitate-the-best* consistently outperforms *imitate-the-majority* and our results suggest that the reason underlying this finding is the speed with which different strategies are able to respond to environmental change. This speed is affected both by different building blocks and the structure of the social environment. *Imitate-the-best* is always faster at finding the good option because its decision rule requires fewer correct instances in the sample and larger samples are always beneficial. In contrast, sample size has a counterintuitive effect on *imitate-the-majority* with smaller samples increasing the likelihood and thereby the speed of finding the correct option. The relative difference between *imitate-the-best* and *imitate-the-majority*, however, is moderated by network structure. More efficient networks (those with lower clustering and shorter path lengths) benefit all strategies and decrease the difference between them while less efficient networks (with more clusters and longer path lengths) increase the difference by having a worse impact on *imitate-the-majority*.

Information sampling as opposed to group-level aggregation has an additional effect on *imitate-the-majority:* it can still converge on the correct option, even if less than 50% of the group is using it. This result lies in contrast to the predictions of the Condorcet and related Theorems on full group-level aggregation of information in a single trial (Grofman, Owen & Feld, 1983).

Both *imitate-the-best* and *imitate-the-majority* have been extensively studied both theoretically and empirically (e.g. Conradt & Roper, 2003; Garcia-Retamero, Takezawa & Gigerenzer, 2006; Hastie & Kameda, 2005; Katsikopoulos & King, 2010; McElreath, Wallin & Fasolo, 2012). Much of this work has studied small and unstructured groups and focused exclusively on the decision-phase of implementing these strategies (but see Pachur, Rieskamp & Hertwig, 2005; Schwenk & Reimer, 2007 for exceptions in other contexts). We believe that this leaves many important details affecting strategy success unaddressed and can be one reason why some studies reach different conclusions. The present study is a first step towards developing a more general framework for capturing the interactions between the building blocks of social heuristics and the structure of the social and task environments that they exploit. We propose that their study can bring novel insight into our understanding of social phenomena including the evolution of different social learning rules, the diffusion of innovations in cultures or the strategy selection process in social domains.

Acknowledgments

The authors wish to thank Konstantinos Katsikopoulos and Henrik Olsson for helpful comments on an earlier draft and the Max Planck Institute for Human Development for financial support.

References

Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Conradt, & Roper, T. J. (2003). Group decision-making in animals. *Nature, 421*, 155-158.

Couzin, I. D., Krause, J., Franks, N. R. & Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature, 433*, 513-516.

Garcia-Retamero, R., Takezawa, M. & Gigerenzer, G. (2006). How to learn good cue orders: When social learning benefits simple heuristics. In R. Sun & N. Mikaye (Eds.), *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 1352-1357). Mahwah, NJ: Erlbaum

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple Heuristics that Make Us Smart.* New York: Oxford University Press.

Grofman, B., Owen, G. & Feld, S. C. (1983). Thirteen theorems in search of the truth. *Theory and Decision, 15,* 261-278.

Hastie, R. & Kameda, T. (2005). The Robust Beauty of Majority Rules in Group Decisions. *Psychological Review*, *112*, 494-508.

Katsikopoulos, K. & King, A. J. (2010). Swarm Intelligence in Animal Groups: When Can a Collective Out-Perform an Expert? *Plos ONE*, *5*, 1-5.

Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, *32*, 4-14.

Lazer, D., & Friedman, A., (2007). The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly, 52,* 667-694.

Mason, W. A., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences, 109*, 764-769.

March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organizational Science*, *2,* 71-78.

McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J. & Waring, T. (2008). Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B, 363*, 3515-3528.

McElreath, R., Wallin, A. & Fasolo, B. (2012). The Evolutionary Rationality of Social Learning. In R. Hertwig & U. Hoffrage & the ABC Research Group (Eds.), *Simple Heuristics in a Social World.* New York: OUP.

Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life.* Harvard University Press.

Pachur, T., Rieskamp, J., & Hertwig, R. (2005). The social circle heuristic: Fast and frugal decisions based on small samples. In K. Forbus, D. Gentner, & T. Reiger. (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 1077-1082). Mahwah, NJ: Erlbaum.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghiranda, S., Lillicrap, T., & Laland, K. N. (2010). Why Copy Others? Insights fomr the Social Learning Strategies Tournament. *Science, 328*, 208-213.

Schwenk, G., & Reimer, T. (2007). Social Influence and Bounded Rationality: Heuristic Decision Making in Complex Networks. In D. S. McNamara & G. Trafton. (Eds.), *Proceedings of the Twenty-Ninth Annual Conference of the Cognitive Science Society* (pp. 1497-1484). Mahwah, NJ: Lawrence Erlbaum Associates.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature, 393*, 409-410.