INTELIGENCIA ARTIFICIAL

Efficiency of the Emergence of Consensus in Complex Networks – assessing force influence

Paulo Urbano, João Balsa, Luis Antunes and Luis Moniz
Faculty of Sciences of Lisbon
{pub,jbalsa,xarax,hal}@di.fc.ul.pt

Abstract Interactions between agents in multi-agent systems have an inherent potential for conflict and must be coordinated. The adhesion to uniform behaviour is one way of coordinating actions: social conventions and lexicons are good examples of coordinating systems, where uniformity promotes shared expectations of behaviour and shared meanings. This paper deals with the emergence of a uniform collective choice, a consensus, inside a population of artificial agents. The efficiency of the formation of a uniform shared decision or choice is an important issue. The nature of interactions and also the nature of society configurations may promote or inhibit consensual emergence. We study the efficiency of uniform decision formation along two dimensions: rules of interaction and network topologies. We compare two different interaction behaviours: the well know and used Highest Cumulative Reward, which is a kind of reinforcement learning behaviour, against a recent behaviour named Recruitment based on Force. We also compare those interaction behaviours along five types of social link networks: fully connected, regular, random, scale-free and small-world.

Keywords: coordination, multiagent systems, convention emergence.

1 Introduction

The achievement of global conventions in multi-agent systems was first addressed by Shoham and Tennenholtz [12] in the beginning of the 90s but it is the subject of recent research [5, 6, 11]. Conventions specify a choice common to all agents, and are a straightforward means for achieving coordination in a multi-agent system. "In multi-agent systems such as multi-planners systems, it is crucial that the agents agree on certain rules, in order to decrease conflicts among them and promote cooperative behaviour"[12].

The issue at stake here relates to collective choice and coordination mechanisms: a homogeneous group is in presence of several strategies and has to select one of them. The strategies are considered equally good: what is important is that the choice is consensual, which particular strategy is chosen is irrelevant. An example of such a collective choice is the lane of traffic on a given country. It is irrelevant whether right lane or left lane is chosen, as long as everyone uses the same.

Legislating conventions before hand (offline) or developing a central control mechanism for generating them can be a difficult and intractable task in dynamic situations [14]. This explains the interest on the design of emergent conventions through a co-learning process [13]: individual agents occasionally meet and observe each other, and they may update their strategies based on this local information. The guarantee of achieving a global consensus and the efficiency of convergence towards it have been natural concerns in the design of local interaction behaviours. Shoham and Tennenholtz [12, 13] compared several strategy update rules and one of them has been selected as a benchmark: the Highest Cumulative Reward (HCR) [13].

In the initial research on convention emergence [12, 13, 14], there were no restrictions on interactions; any agent could interact by chance with any other individual. Kittock [8] introduced interaction graphs in order to specify restrictions on interactions and made experiences with the HCR update rule along different interaction graphs. He also made experiments with regular graphs and his experiments showed that interaction topology is important in the efficiency of convention emergence—he conjectured that efficiency depends on the diameter of the graph. In what concerns the number of interactions needed to accomplish consensus, Kittock observed a variation with the number of agents of $O(N^3)$ for regular graphs and $O(NlogN)$ for fully connected ones. Other graph properties like clustering coefficient and average path length are also important.

However, regular graphs are not very realistic. "If we pay attention on real networks, we find out that most of them have a very particular topology, they are complex networks. (…) Complex networks are well characterized by some special properties, such as the connectivity distribution (either exponential or power-law) or the small-world property" [5, 6]. Delgado et al have made experiments with HCR update rule for fully connected, regular, scale-free and small-world graphs [4], and their results were consistent with Kittock's, confirming the relation between efficiency and graph diameter. Delgado observed also that scale-free and small-world networks were as efficient as fully connected ones, but small-world networks were slower to converge to a unique choice.

In [16] a new rule for strategy update was introduced, named Recruitment based on Force with Reinforcement (RFR). This rule showed faster convergence than HCR in the case of fully connected networks. In RFR, agents have different power to influence others, but their force is not defined a priori in a hierarchy network as in [9], rather it evolves dynamically along the interactions. Agents submit to stronger agents, copying their strategies, but also inheriting their force, in a double mimetic process.

Recently, we have proposed some experiments to investigate how permeability between contexts in a set of heterogeneous networks can have a decisive influence in convergence for a global consensus [1, 2]. Our findings indicate that the co-existence of several concomitant relations with context permeability connecting them qualitatively changes the progression towards convergence [1], while allowing convergence to be achieved more often and faster in hard cases, such as scale-free networks, especially if connected to at least one regular network with degree greater than 2 [2].

In the present paper, we advance research on the emergent collective adoption of a common strategy. We foster the study on the performance of RFR rule for strategy update. We apply RFR to random, regular and complex networks, comparing it with the HCR strategy update rule, following the above mentioned works of Kittock and Delgado.

2   Interaction Graph topologies

The interaction graph topology is a general way of modelling restrictions on interactions. Restrictions could be due to spatial barriers, communicating links, different castes, social groups, etc. We have experimented with five network topologies: fully connected, regular, scale-free, small-world and random. The average path length is calculated by finding the shortest path between all pairs of nodes, adding them up, and then dividing by the total number of pairs. It indicates us, on average, the number of steps it takes to get from one member of the network to another. The diameter of a graph is the longest-shortest-path between nodes. The clustering coefficient is a measure of "all-my-friends-know-each-other" property, which is sometimes described as "the friends of my friends are my friends." The clustering coefficient of a node is the ratio of existing links connecting a node's neighbours to each other to the maximum possible number of such links. The clustering coefficient for the entire network is the average of the clustering coefficients of all the nodes.

2.1   Regular Graphs

By definition, a graph is considered regular when every node has the same number of neighbours. We are going to use a special kind of regular graph, explored in [8] and named *Contract Net with Communication Radius K* in [15]. $C_{N,K}$ is the graph (regular ring lattice) on N nodes such that node i is adjacent to nodes *(i+j) mod N* and *(i−j) mod N* for *1≤j≤K*. In a $C_{N,K}$ graph, every node has connectivity *2*K*. These are highly clustered graphs but have very long path lengths (average path length and diameter grow linearly with the number of nodes).

2.2   Fully Connected

In this type of graph topology, named K_N, there are no restrictions on the pattern of interactions: each agent is connected to every other agent in the society. This means that an agent can potentially interact with any other agent. K_N is a special case of a regular graph where each agent has *N-1* neighbours, in a group of N agents.

2.3   Random Graphs

$R_{N,K}$ are random graphs with N nodes and average connectivity of K. Every node has precisely k neighbours chosen randomly. The clustering coefficient of $R_{N,K}$ tends to *0* and the average path length is small and grows logarithmically with *N*.

2.4   Scale-free

This network type has a large number of nodes connected only to a few nodes and a small number of well-connected nodes called hubs. The power law distribution highly influences the network topology. It turns out that major hubs are closely followed by smaller ones. These ones, in turn, are followed by other nodes with an even smaller degree, and so on. As the network changes in size, the ratio of hubs to the number of nodes in the rest of network remains constant—this is why it is named scale-free. The connectivity of a scale-free network follows a power law $P(k) \sim k^{-\gamma}$. Such networks can be found in a surprisingly large range of real world situations, ranging from the connections between websites to the collaborations between actors.

To generate the scale-free graphs we have used the Albert and Barabási extended model [3], since Delgado argues that it allows some control over the exponent (γ) of the graph [5]. The inspiration of this algorithm is that of "preferential attachment," meaning that the most "popular" nodes get most of the links. The construction algorithm relies on four parameters: m_0 (initial number of nodes), m (number of links added and/or rewired at every step), p (probability of adding links), and q (probability of edge rewiring). The algorithm starts with m_0 isolated nodes and at each step performs one of these three actions until the desired number N of nodes is obtained:

(1) with probability p, add m ($\leq m_0$) new links. Pick two nodes randomly. The starting point of the link is chosen uniformly and the end point of the link is chosen according to the probability distribution: $\Pi_i = (k_i + 1)/\sum_j (k_j + 1)$, where Π_i is the probability of selecting the i^{th} node and k_i is the number of edges of node i. This process is repeated m times.

(2) with probability q, m edges are rewired. That is, repeat m times: choose uniformly at random one node i and one link l_{ij}. Delete this link and choose a different node k with probability $\{\Pi_l\}_{l=1,...,N}$ and add the new link l_{ik}.

(3) with probability 1-p-q add a new node with m links. These new links will connect the new node to m other nodes chosen according to $\{\Pi_l\}_{l=1,...,N}$.

Using this algorithm, the parameter γ is a function of m and p ($\gamma = (2m(1-p)+1)/m+1$).

2.5   Small World

The Small World graphs are highly clustered graphs (like regular graphs) with small average path lengths (like random graphs, described above). To generate small world graphs we use the Watts-Strogatz model [17, 18]. It depends on two parameters, connectivity (*K*) and randomness (*P*), given the size of the graph (*N*).

This model starts with a $C_{N,K}$ graph and then every link is rewired at random with probability *P*, that is, for every link l_{ij} we decide whether we change the "destination" node with probability *P*; if this is the case, we choose a new node k uniformly at random (no self-links allowed) and add the link l_{ik} while erasing link l_{ij}. In fact, for *P = 0* we have $W_N = C_{N,K}$ and for *P = 1* we have a completely random graph (but not scale-free). For intermediate values of *P* there is the "small-world" region, where the graph is highly clustered (which means it is *not* random) but with a small characteristic path length (a property shared with random graphs).
Albert-Barabasi model graphs have not the small-world property and reciprocally the Watts-Strogatz model does not generate scale-free graphs (it generates an exponential connectivity distribution, not a power law).

3   Strategy update rules

Agent societies consist of *N* agents on a graph, where every agent is located on a node of the graph. Its adjacent nodes are its neighbours. In order to make experiments and simulations we have adopted a simple agent model where they have at their disposal a finite repertoire of strategies. We only deal with the two strategies case. We use here the concept strategy in a very abstract way: it can be a social norm, like driving on the left or on the right

lane, the meaning of a word, an orientation for flocking, etc. In order to focus on the essential features of agent interactions, the agent environment consists solely of other agents, which in turn depend on the network topology. So, each agent has to adopt one of the strategies from the repertoire and through mutual interactions they can change their adopted strategies along time. A consensus, or collective choice, exists when all the agents are using the same particular strategy.

From the point of view of each agent, there is an interaction scenario of a sequence of pairwise asymmetric encounters, where it meets randomly one of its neighbours. After an encounter, each agent updates its strategy, i.e., it selects the strategy it will use in the next interaction—the result need not necessarily be a change in strategy adoption. Therefore, agents need strategy update rules (behaviours). We assume that each agent updates its strategy at each encounter. Shoham and Tennenholtz [12, 14] have studied the effects of updating less frequently on the efficiency of global choice emergence. We only consider asymmetric encounters, where only one of the agents applies its strategy update rule, based on the strategies used by all the individuals involved in the interaction. Thus, interactions are always considered from the point of view of some particular agent. We now describe the two strategy update rules whose performance we subsequently compare. In the first scenario the agent and its selected partner strategies are crucial for the update, and in the second case, it is the simultaneous strategies of its neighbours and its own strategy that influence the update rule.

3.1 Highest Cumulative Reward/External Majority/Feedback Positive with Score

The most referred strategy update rule is the *Highest Cumulative Reward* update rule (HCR), which was developed in the context of game theory by Shoham and Tennenholtz [13]. Intuitively, a game involves a number of players each of which has available to it a number of strategies. Depending on the strategies selected by each agent, they each receive a certain payoff. The payoffs are captured in a payoff matrix.

Thus, returning to the context of this paper, when two agents meet they play a pure coordination game, which is an instance of the class of coordination games introduced by Lewis [10]. The pure coordination game is defined by the following symmetric payoff matrix:

	A	B
A	+1,+1	-1,-1
B	-1,-1	+1,+1

Suppose that every player has two available strategies, say A and B. If both players play A, both players receive a payoff of 1. If they play B they receive a payoff of 1. When the players do not agree, for example, player 1 plays A and player 2 plays B, they will both receive a payoff a -1; the remaining situation is symmetric. The condition on the entries of the payoff matrix makes it clear that the best action consists in playing the same strategy, i.e., coordinating.

According to the HCR update rule, an agent switches to a new strategy if and only if the total payoff obtained from that strategy in the latest m interactions is greater than the payoff obtained from the current strategy in the same last m interactions. The m parameter may not have limit, implying that the full history of pairwise meetings will play a role in the strategy selection process, or we can implement a forgetting mechanism by limiting m. The agents' memories register the payoffs that each strategy has received during the last m encounters. When an agent receives new feedback it discards its old memory to maintain the memory at a fixed size.

The *External Majority* strategy update rule (EM) was introduced by Shoham and Tennenholtz [12] and is the following: adopt the strategy that was observed more often in other agents in the last m interactions, but remain with your current strategy in case of equality. Shoham and Tenneholtz [14] showed that EM and HCR are equivalent strategy update rules in the case there is a repertoire of two strategies to select. In EM, memory is used to register the strategies observed during the last interactions. An agent updates its memory after observing its partner strategy and then decides to change to a new strategy only in case it was more frequently observed than the current one.

In the context of lexical emergence, Kaplan [7] introduced a strategy update rule called *Positive Feedback with Score*, which is pretty much the same as EM. The only difference is that, in case of equality, the agent does not necessarily remain with its current strategy but chooses randomly one of the previously most seen strategies. Kaplan considered the full history of encounters for strategy update.

3.2   Recruitment based on Force with Reinforcement (RFR)

In this strategy update rule there is a new attribute, besides the strategy, called *force*. Thus, agents are characterised by two attributes: *strategy* and *force*. These attributes can be observed during encounters. During a dialogue (asymmetric), involving two agents, one is the observing agent and the other is the observed one. The observing agent "fights" metaphorically with its partner, comparing its force with the partner's force. If the observing agent is the stronger one, or if they have identical force, it will loose the fight; otherwise it will be the winner. The winner's behaviour is: (1) if they have the same current strategy their forces are reinforced by 1 unit, otherwise (2) it does nothing. The loser's behaviour is: (1) it imitates both strategy and force in case they have different current strategies, otherwise imitates the force of the winner agent and increments its force by 1 unit. In summary, stronger agents recruit weaker agents for their parties, enlarging the influence of their options. As the recruited agents will be at least as strong as the winners, they will be better recruiters.

At the beginning of an experiment, every agent has the same value of force (0) and their forces evolve along with interactions. Therefore, there is no a priori (off line) power hierarchy. This shows a clear contrast with the work of Kittock on authority [9], where agents have fixed different influences on one another's behaviour, modelled by the probability of receiving feedback during encounters: the more influential agents receive feedback with some probability, while the less influential agents always receive feedback.

The force attribute can be interpreted not as the strength of an individual because, being imitated, it is diffused along agents, and does not belong to any agent, but as the force of the strategy the player is adopting. The more the strategy is diffused the more it will have stronger representatives. So when a player observes a stronger agent it is recruited, inheriting his force, i.e., updating the information about the strategy it is now adopting. There is a positive reinforcement when an agent faces another with the same strategy during a meeting, which is the amplification mechanism for strategy diffusion.

RFR does not try to simulate any natural behaviour and it was introduced in [16] as the best outcome of experimenting with several strategies involving the idea of emergence of hierarchies in consensus emergence. We think it is simpler than the HCR as agents do not need to maintain the recent history of encounters in spite of using force as an extra attribute. It maintains the essential properties of HCR, which is the capacity to adapt, locality (an agent relies only is the information gathered in interactions), and no more cognitive skills than the capability to imitate.

4   Experiments and Results

Experiments were conducted using the most recent version of the Netlogo platform (version 4.0.2 [19], released December 2007).

The system starts with half of the agents adopting randomly one of the strategies (50% possibilities for each). In each step, every agent, in an asynchronous way, is selected and chosen for asymmetric strategy updating. The order of selected agents is completely random and changes in each iteration. We use the same measure of performance as in [8, 5]: average number of interactions to a fixed convergence, where convergence means the fraction of agents using the majority strategy. We made one hundred runs for each parameter setting and in each run we have measured the number of encounters until 90% convergence and calculated the average performance of the different runs.

Our main goal was to compare the performance of the two strategy update rules (behaviours) described in section 3: Highest Cumulative Reward (HCR) and Recrutiment based on Force with Reinforcement (RFR). In order to choose the size of memory of HCR update rule we have made a lot of experiments with different types of networks and the size = 3 achieved the best performance. Both Kittock [8, 9] and Delgado [5, 6] chose size = 2 in their convention experiments, but size 3 HCRs outperformed size 2 HCRs in our own tests. So we will only present the comparison between RFR and the best HCR (HCR-3). The comparison between these two behaviours was made along the different kinds of networks described in section 2, using different parameter settings. Again, we only present here the most representative experiments.

For all settings we made the number of agents range from 100 to 1000, using a step of 100. We made 500 experiments for each scenario.

In Fig. 1, we can see a comparison between the average number of meetings needed for a 90% convergence using the two behaviours in fully connected (dashed lines) and scale-free networks (with γ=2.5, solid lines), and with RFR (diamonds) and HCR (squares) as update rules.
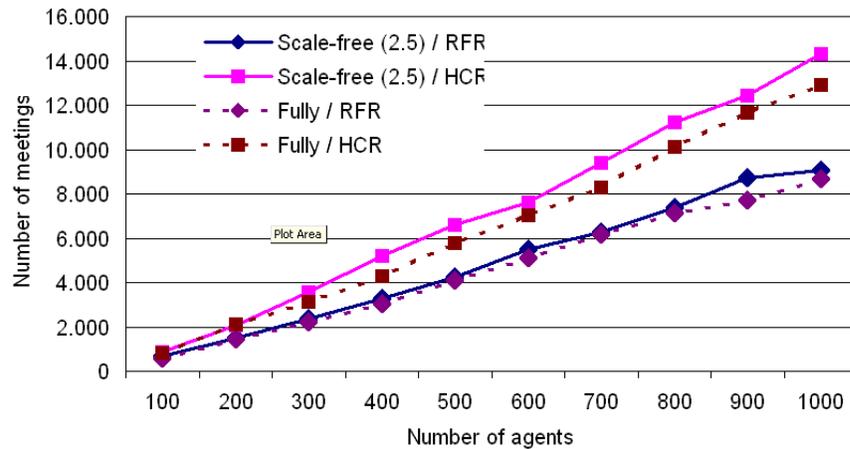
Fig. 1. RFR vs HCR in fully connected and scale-free networks.

In both cases RFR clearly outperforms HCR, since consensuses are reached in a much lower number of meetings. Besides, the difference appears to increase with N. In fact, in the above experiments, improvements vary between 25% and 37%.

Using other types of networks the difference between the two behaviours is even clearer. In Fig. 2 we compare the performance of the two behaviours in small-world networks (with P = 0.1, dashed lines) and regular networks (with K = 20, each agent with 40 neighbours, full lines). In these cases we need to use a logarithmic scale in the y axis. The average number of meetings needed to reach a consensus is again much lower when RFR (diamonds) is used. In Small-world networks improvements reach 80%. Using regular networks the difference with N greater than 500 is so huge that it is not represented.
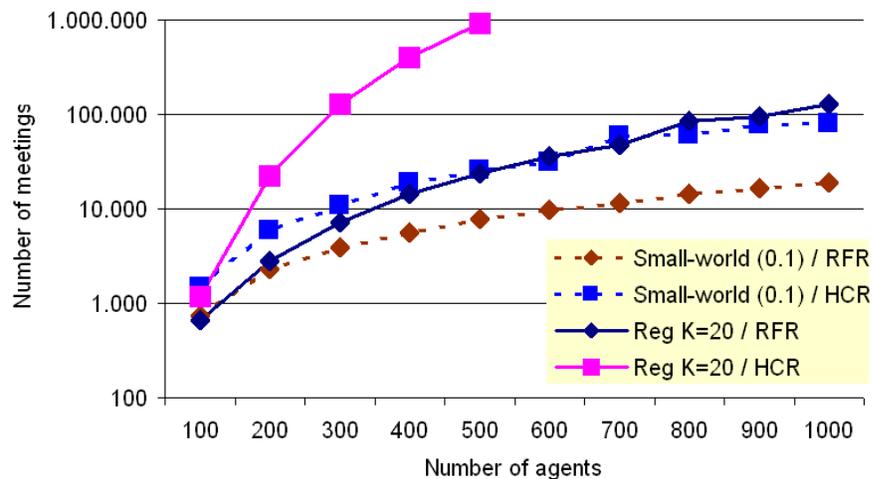

Fig. 2. RFR vs HCR in small-world and regular networks.

We also made a set of experiments with random networks for N=100 and N=200, and with different values for K, and the results also showed that RFR always performs better, with an average of 30% improvement over HCR regarding the number of meetings needed to reach consensus.

5   Conclusions and Future Work

In what respects our primary goal of comparing HCR with RFR, the main conclusion is that RFR always perform better than HCR. This conclusion is valid for all types of networks with different parameter settings. This is an impressive result, since HCR was considered a benchmark, and we show that it is possible to outperform it. According to our experiments RFR represents at least a 25% improvement over HCR but is much higher in many settings, most extremely in regular networks.

An interesting point is that our results for the RFR strategy update rule show that the network diameter strongly influences the performance, as noted in other settings by authors such as Kittock [8] and Delgado [5]. The smaller the network diameter is, the better. Also, the performance of fully-connected networks and scale-free ones seems to be quite similar (as can be observed in Fig. 1, comparing lines with similar marks, that correspond to the same behaviour). This is a very important result since scale-free networks are much less expensive than fully-connected ones. Nevertheless, we must perform experiments with greater values of N in order to obtain definite conclusions on this matter.

Besides performing experiments with larger values of N, two other aspects are scheduled for short-term future work. One is to consider that agents can choose between more than two strategies. The other is to explore the performance of these behaviours in networks with dynamical structure.

References

[1] Luis Antunes, João Balsa, Pauo Urbano, and Helder Coelho. The challenge of context permeability in social simulation. In Proceedings of the Fourth European Social Simulation Association, 2007.

[2] Luis Antunes, João Balsa, Paulo Urbano, and Helder Coelho. Exploring context permeability in multiple social networks. In Proceedings of the World Congress on Social Simulation, 2008.

[3] R. Albert and A. Barabási. Topology of evolving networks: local events and universality. Physical Review Letters, 85:5234-5237, 2000.

[4] A. L. Barabási. Linked: The New Science of Networks. Perseus Publishing, Cambridge, Massachusetts, 2002.

[5] Jordi Delgado. Emergence of social conventions in complex networks. Artificial Intelligence, 141(1/2):171-185, 2002. DOI: 10.1016/S0004-3702(02)00262-X

[6] Jordi Delgado, Josep M. Pujol, and Ramon Sangüesa. Emergence of coordination in scale-free networks. Web Intelligence and Agent Systems, 1(2):131-138, 2003.

[7] Frédéric Kaplan. L'Emergence D'un Lexique Dans Une Population d'Agents Autonomes. PhD thesis, Université de Paris 6, 2000.

[8] J. E. Kittock. Emergent conventions and the structure of multi-agent systems. In L. Nadel, D. Stein (Eds.), 1993 Lectures in Complex Systems, in: SFI Studies in the Sciences of Complexity. Addison-Wesley, Reading, MA, 1995.

[9] James E. Kittock. The impact of locality and authority on emergent conventions: Initial observations. In Proceedings of the Twelfth National Conference on Artificial Intelligence, pages 420-425, Seattle, WA, 1994.

[10] D. K. Lewis. Convention: A Philosophical Study. Harvard University Press, Cambridge, MA, 1969.

[11] Sandip Sen and Stéphane Airiau. Emergence of norms through social learning. In Manuela M. Veloso, editor. IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007, pages 1507-1512, 2007.

[12] Y. Shoham and M. Tennenholtz. Emergent conventions in multi-agent systems. In C. Rich, W. Swartout, and B. Nebel, editors, Proceedings of Knowledge Representation and Reasoning (KR&R-92), pages 225-231, 1992.

[13] Yoav Shoham and Moshe Tennenholtz. Co-learning and the evolution of social acitivity. Technical Report CS-TR-94-1511, Stanford University, Department of Computer Science, 1994.

[14] Yoav Shoham and Moshe Tennenholtz. On the emergence of social conventions: Modeling, analysis and simulations. Artificial Intelligence, 94(1-2):139-166, 1997. DOI: 10.1016/S0004-3702(97)00028-3

[15] Moshe Tennenholtz. Convention evolution in organizations and markets. Computational & Mathematical Organization Theory, 2(4):261-283, 1996.

[16] Publication by one of the authors of the present paper.

[17] D. J. Watts. *Six Degrees: The Science of a Connected Age*. W.W. Norton & Co., New York, (2003).

[18] D. J. Watts and S. H. Strogatz. Collective dynamics of `small-world' networks. Nature, 393:440-442, 1998.

[19] U. Wilensky, U. NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. (1999)