

Effects of Evolution on the Emergence of Scale Free Networks

Bijan Ranjbar-Sahraei¹, Daan Bloembergen^{1,3}, Haitham Bou Ammar², Karl Tuyls³ and Gerhard Weiss¹

¹Department of Knowledge Engineering, Maastricht University, The Netherlands

²Computer and Information Science Dept., Grasp Lab, University of Pennsylvania

³Department of Computer Science, University of Liverpool, United Kingdom

Abstract

The evolution of cooperation in social networks, and the emergence of these networks using simple rules of attachment, have both been studied extensively although mostly in separation. In real-world scenarios, however, these two fields are typically intertwined, where individuals' behavior affect the structural emergence of the network and vice versa. Although much progress has been made in understanding each of the aforementioned fields, many joint characteristics are still unrevealed. In this paper we propose the Simultaneous Emergence and Evolution (SEE) model, aiming at unifying the study of these two fields. The SEE model combines the continuous action prisoner's dilemma (modeling the evolution of cooperation) with preferential attachment (used to model network emergence), enabling the simultaneous study of both structural emergence and behavioral evolution of social networks. A set of empirical experiments show that the SEE model is capable of generating realistic complex networks, while at the same time allowing for the study of the impact of initial conditions on the evolution of cooperation.

Introduction

Understanding the dynamics of networked interactions is of vital importance to a wide range of research areas. For example, these dynamics play a central role in biological systems such as the human brain (Bullmore and Sporns, 2009) or molecular interaction networks within cells (Barabási and Oltvai, 2004); in large technological systems such as the word wide web (Easley and Kleinberg, 2010); in social networks such as Facebook (Backstrom et al., 2011; Ghanem et al., 2012; Ugander et al., 2011); and in economic or financial institutions such as the stock market (Chapman et al., 2012; Jackson, 2008). Recently, researchers have focused on studying the evolution of cooperation in networks of self-interested individuals, aiming to understand how cooperative behavior can be sustained in the face of individual selfishness (Hofmann et al., 2011; Nowak and May, 1992; Santos and Pacheco, 2005; Ranjbar-Sahraei et al., 2014).

Many studies have targeted the discovery of structural properties of networks that promote cooperation. For instance, Santos and Pacheco (2005) show that cooperation has a higher chance of survival in scale-free networks; Oht-

suki et al. (2006) find a relation between the cost-benefit ratio of cooperation and the average node degree of a network that determines whether cooperation can be sustained; and Van Segbroeck et al. (2010) look at heterogeneity and clustering to find that these structural properties influence behavior on the individual rather than network-wide level. Others have focused on the role of the particular interaction model between neighboring nodes in determining the success of cooperation. For example, Hofmann et al. (2011) simulate various update rules in different network topologies and find that the evolution of cooperation is highly dependent on the combination of update mechanism and network topology. Ranjbar-Sahraei et al. (2014) propose a mathematical model, based on control theory, that allows individuals to choose their actions from a continuous range between pure defection and pure cooperation, and show that this model leads to a higher degree of cooperation than the traditional binary choice models. Control theory is also used by Bloembergen et al. (2014) aiming at ways of influencing the behaviors in social networks.

These studies have assumed the network to be fixed, looking only at the evolution of cooperation over time. In contrast, real-world social networks are not fixed, but continuously change as individuals make and break their ties (Kossinets and Watts, 2006). To this end, Zimmermann and Eguíluz (2005) and Santos et al. (2006) allow individuals to choose with whom to interact, e.g. by giving them the possibility to break ties with 'bad' neighbors and replacing them with a random new connection, and show that such a mechanism may promote cooperation. However, these works still assume a network to be in place, only modifying the connections between nodes over time.

Here, we investigate what happens when nodes are added to the network during interaction. Specifically, we start with an empty network, and add a new node at each time step. Simultaneously, the existing nodes in the network interact following the Continuous Action Iterated Prisoner's Dilemma (CAIPD) model of Ranjbar-Sahraei et al. (2014). New nodes are attached following preferential attachment. Usually, preferential attachment is assumed to follow the

Barabási-Albert model (Barabási and Albert, 1999) where links are formed to existing ones proportional to their degree. However, in many social scenarios it intuitively makes sense to look at other individuals' performance rather than their degree when determining with whom to interact - connecting with high performing individuals may give you an edge. We empirically compare both methods of preferential attachment, looking at the structure of the networks formed in detail.

This paper proceeds as follows. First, relevant background is provided on networks and game theory, and an overview of the continuous action iterated prisoner's dilemma (CAIPD) model and preferential attachment is given. These lay the foundation for the proposed Simultaneous Emergence and Evolution (SEE) model that is detailed thereafter. Finally, empirical evaluations highlight the properties of the proposed model.

Background

This section provides background knowledge needed for the remainder of the paper. Firstly, preliminaries on the theory of networks and games are given, constituting the foundation for the model of the evolution of cooperation used in this work. Hereafter, the continuous action iterated prisoners dilemma (CAIPD) is introduced. Finally, preferential attachment, used for the generation of Scale Free (SF) networks, is detailed.

Networks

Networks describe collections of entities (nodes) and the relation between them (edges). Formally, a network can be represented by a graph $\mathbb{G} = (\mathcal{V}, \mathcal{W})$ consisting of a non-empty set of nodes (or vertices) $\mathcal{V} = \{v_1, \dots, v_N\}$ and an $N \times N$ adjacency matrix $\mathcal{W} = [w_{ij}]$ where non-zero entries w_{ij} indicate the (possibly weighted) connection from v_i to v_j . If \mathcal{W} is symmetrical, such that $w_{ij} = w_{ji}$ for all i, j , the graph is said to be undirected, meaning that the connection from node v_i to v_j is equal to the connection from node v_j to v_i . In social networks, for example, one might argue that friendship is usually mutual and hence undirected. This is the approach followed in this work. In general however this need not be the case, in which case the graph is said to be directed, and \mathcal{W} asymmetrical. The neighborhood, \mathbb{N} , of a node v_i is defined as the set of nodes it is directly connected to, i.e. $\mathbb{N}(v_i) = \cup_j v_j : w_{ij} > 0$. The node's degree $\text{deg}[v_i]$ is given by the cardinality of its neighborhood.

Several types of networks have been proposed that capture the structural properties found in large social, technological or biological networks, two well-known examples being the small-world and scale-free models. The *small-world* model exhibits short average path lengths between nodes and high clustering, two features often found in real-world networks (Watts and Strogatz, 1998). Another model is the *scale-free* network, characterised by a heavy-tailed degree distri-

bution following a power law (Barabási and Albert, 1999). In such networks the majority of nodes will have a small degree while simultaneously there will be relatively many nodes with very large degree, the latter being the hubs or connectors of the network. For a detailed description of networks and their properties, the interested reader is referred to Jackson (2008).

Game Theory

Game theory models strategic interactions in the form of games (Gibbons, 1992). Each player has a set of actions, and a preference over the joint action space that is captured by the received payoffs. The goal for each player is to come up with a strategy (a probability distribution over its actions) that maximizes his expected payoff in the game. A strategy that maximizes the payoff given fixed strategies for all opponents is called a best response to those strategies. The players are thought of as individually rational, in the sense that each player purely tries to maximize his own payoff, and assumes the others are doing likewise. However, this reasoning might not always lead to a beneficial outcome for everyone, and might even be detrimental to all players in the game. Often, there is tension between individual rationality on the one hand, and social welfare on the other.

This archetypal dilemma is aptly captured by the Prisoner's Dilemma (Axelrod and Hamilton, 1981). In this one-shot interaction, players simultaneously choose between either cooperation or defection, after which payoffs are distributed based on their joint action. Cooperation is costly, however cooperators distribute benefits among the other players. Defectors do not pay a cost, but do receive benefits from cooperators as well. In this game, defection (free-riding) is a best response against any opponent strategy, and therefore individually rational players can be expected to defect. However, if all players would cooperate their distributed benefits would outweigh the cost of cooperation, and hence all players would be strictly better off. Herein lies the dilemma.

In this work the players are nodes in the network, repeatedly playing a game with their neighbors. The players have no knowledge of the underlying game, however this repeated interaction allows for adaptation, i.e. to learn a better strategy over time based on the payoff received. The game used in this paper is a generalization of the classical Prisoner's Dilemma, in that the players can have a continuous strategy defining their level of cooperation rather than a binary choice, and payoffs are calculated accordingly. The dilemma, however, remains.

Continuous Action Iterated Prisoner's Dilemma

The continuous action iterated prisoner's dilemma (CAIPD) is a mathematical model of the evolution of cooperation on complex social networks, proposed in Ranjbar-Sahraei et al. (2014). The model describes how the individuals in the net-

work, placed on the nodes, interact with their neighbors according to the aforementioned prisoner's dilemma. Cooperators incur a cost c for each interaction, while their neighbor receives a benefit b with $b > c$. Defectors free-ride, in the sense that they do not pay costs while still receiving benefits from cooperative neighbors.

Formally, in CAIPD the individuals in a network are represented by N vertices $v_i \in \mathcal{V}$ for $i = \{1, \dots, N\}$ on a weighted graph $\mathbb{G} = (\mathcal{V}, \mathcal{W})$. The characteristics of the graph \mathbb{G} are described by the symmetrically weighted $N \times N$ adjacency matrix $\mathcal{W} = [w_{ij}]$. Namely, the connections between the i^{th} and j^{th} individual are denoted by $w_{ij} \in \{0, 1\}$ with all $w_{ii} = 0$. The latter prevent individuals from self-interaction. In contrast to other existing models, CAIPD allows individuals to choose their level of cooperation from a continuous range between pure cooperation and pure defection, rather than a binary choice between those. This choice is captured by the individuals' state variable $x_i \in [0, 1]$. A value of $x_i = 0$ corresponds to pure defection while $x_i = 1$ represents pure cooperation; however x can take on any arbitrary value between those extremes. Extending the prisoner's dilemma to account for this continuous nature of cooperation, a player pays a cost cx_i while the opponent receives a benefit bx_i , again with $b > c$. This way a defector (i.e., $x_i = 0$) pays no cost and distributes no benefits. Then, the fitness of player i can be calculated as:

$$f_i = -\text{deg}[v_i]cx_i + b \sum_{j=1}^N w_{ij}x_j \quad (1)$$

where $\text{deg}[v_i]$ denotes the number of neighbors of the individual v_i . Each player observes how well their neighbors are doing, and will then take over one of their neighbors' strategies with some probability based on their respective fitness difference. In particular, player i adopts the strategy of neighbor j with probability

$$p_{ij} = w_{ij} \cdot \text{sigmoid}(\beta(f_j - f_i)) \quad (2)$$

where $\text{sigmoid}(x) = \frac{1}{1+\exp(-x)}$, and β is a parameter varying the opponents' influence on individual i .

A network with a state \mathbf{x} and topology \mathbb{G} is defined as $\mathbb{G}_{\mathbf{x}} = (\mathbb{G}, \mathbf{x})$ with $\mathbf{x} = [x_1, x_2, \dots, x_N]^T$ representing the state of each node in the network. Such a network $\mathbb{G}_{\mathbf{x}}$ can then be regarded as a dynamical system, where \mathbf{x} evolves with respect to time. This evolution depends on a nonlinear mapping

$$\dot{\mathbf{x}} = [h_1(\mathbf{x}), \dots, h_N(\mathbf{x})]^T. \quad (3)$$

Specifically, the dynamics of the i^{th} player are described by

$$h_i(\mathbf{x}) = \frac{1}{\text{deg}[v_i]} \left[\sum_{j=1}^N p_{ij} (x_j(t) - x_i(t)) \right] \quad (4)$$

where $\text{deg}[v_i]$ denotes the number of neighbors of player i (i.e., the degree of node v_i); x_i and x_j denote the current cooperation level of players i and j , respectively; and p_{ij} represents the probability with which player i adopts the strategy of player j , as defined in Equation 2. For a detailed description of the CAIPD model, the interested reader is referred to Ranjbar-Sahraei et al. (2014).

The CAIPD model is more general than other existing models, in that it can be used to model the evolution of cooperation on arbitrary complex social networks. Moreover, the continuous nature of the model, allowing for a degree of cooperation rather than a binary choice, is better suited to model many real-world problems (Killingback and Doebeli, 2002). Finally, the deterministic nature of the mathematical model makes it computationally less complex, allowing to study a wider range of scenarios. For these reasons, we adopt this model in the current paper as well.

Preferential Attachment

The preferential attachment model was proposed by Barabási and Albert to explain the power-law degree distribution that is present in many complex real-world networks (Barabási and Albert, 1999). This model is based on the assumption that, in many social settings, the chance of making new connections grows proportionally with the number of connections that you already have (also known as the rich-get-richer phenomenon). The Barabási-Albert model simulates this by growing the network over time, adding one new node at a time, and linking it to a fixed number of existing nodes, these being chosen proportionally to their current degree. Specifically, starting from an initial network of m_0 nodes, at every time step one new node is added to the network. The new node forms $m < |m_0|$ connections to existing nodes, where the probability p_i that the new node connects to existing node v_i is proportional to its degree:

$$p_i = \frac{\text{deg}[v_i]}{\sum_j \text{deg}[v_j]}. \quad (5)$$

Preferential attachment generates a long-tailed degree distribution following a *power-law*:

$$P(k) \sim k^{-\alpha}$$

with k denoting the degree of the nodes. The power-law exponent for the Barabási-Albert model is $\alpha = 3$; in comparison, many real-worlds complex networks have been shown to lie in the range $2 \leq \alpha \leq 4$ (Barabási and Albert, 1999; Newman, 2005).

It is worth noting that although the Barabási-Albert model is a well known model for generating scale-free graphs, individual properties of the nodes other than their degree are typically not taken into account. In real world scenarios, however, both structural as well as behavioral properties affect the preferential attachment process. For instance, in a

co-authorship network, new authors may indeed have a tendency to team up with existing authors that already worked together with many others (i.e., high degree nodes, a structural property), but will also consider paper quality, number of citations, etcetera (i.e., individual behavioral properties).

In this work, we adapt the preferential attachment model of network growth to take behavioral properties into account when forming new links, where the behavior itself follows the dynamics of the CAIPD model. This approach of Simultaneous Emergence and Evolution (SEE) is detailed next.

Simultaneous Emergence and Evolution

Aiming at a unification, the Simultaneous Emergence and Evolution (SEE) model incorporates two evolutionary procedures. The first is concerned with the evolution of behaviors in the network, which follows from the CAIPD model. The second deals with the construction of the network itself. Here, preferential attachment is used. Contrary to previous works, however, the links that each new individual forms with existing nodes depend on the current fitness of those nodes under the CAIPD dynamics, rather than on their degree. Next, an in-depth description of the SEE algorithm is presented.

The SEE Algorithm

Starting from m initially connected individuals, new nodes are added one at a time. The initial state of these m nodes, as well as of each new node, are set randomly to either pure defection or pure cooperation with equal probability. Each new node is connected to m existing ones with a probability proportional to the fitness of the existing nodes, computed according to Equation 1. The connection probability p_i (i.e., the probability that a new node is connected to i) is defined as

$$p_i = \frac{f_i}{\sum_j f_j} \quad (6)$$

where f_i is the fitness of node i and the sum runs over all N pre-existing nodes $j = 1, 2, \dots, N$. Therefore, nodes with high fitness tend to quickly accumulate more neighbors, while nodes with low fitness are unlikely to be chosen as the connector for a new node. An upper limit size of N_{\max} is defined. This ensures that the network halts its expansion after reaching size N_{\max} .

In parallel to the structural emergence of the network, the CAIPD model is used to evolve the individual behaviors of the existing nodes. At each iteration, the adjacency matrix \mathcal{W} is updated. Fitnesses are then computed according to Equation 1. These new fitness values are then used to update the state of each node, and therefore of the network as a whole, using the dynamical model of Equations 3 and 4. This is in practice performed with an adequately small step size. The SEE model allows to vary the update rates of both evolutionary processes independently. For example, the behavior of the individual nodes might evolve faster or slower

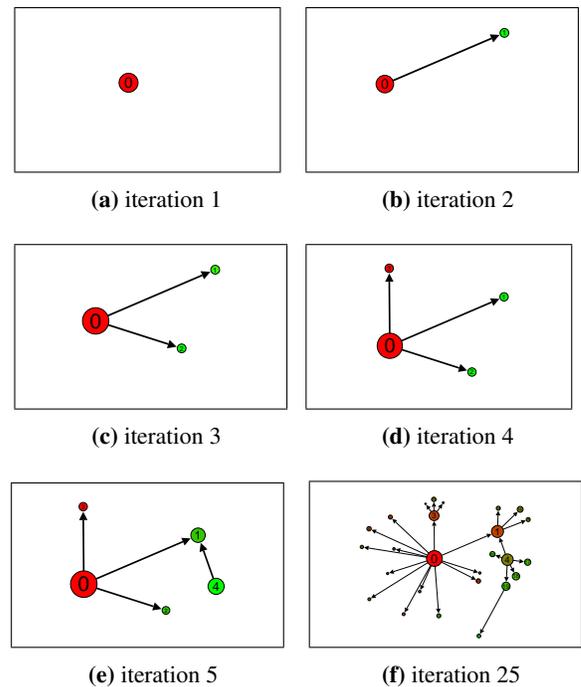


Figure 1: Emergence and evolution of a network according to the SEE model. Node size reflects individual’s fitness and its color denotes the state (red for defection to green for cooperation). The direction of arrows shows how individuals influence each other.

than the rate at which new nodes are added. This ratio between the update rate of the behavior and the update rate of the network is defined by R_{Evo} , such that when at each time step k a new node is added, the CAIPD model progresses R_{Evo} steps.

Illustration of SEE for a Sample Network

In this section an illustration of the SEE model on a sample network is presented. Initially, there is just one node with pure defection state $x_0 = 0$, as depicted in Figure 1(a). At the second iteration, Figure 1(b), a cooperating individual is entering the environment (i.e. individual 1), and gets attached to the defector (i.e., individual 0). At this stage, the defector acquires some benefits from the cooperator, while imposing a cost on the cooperator. This results in a higher fitness for the defector than the cooperator (depicted using the node size in the figure).

For further illustration, Figures 1(c)-1(e) show the attachment of three more individuals with defecting or cooperating states (chosen randomly) after joining the network. In parallel to this network emergence, individuals influence each other as described by the CAIPD model, resulting in a simultaneous evolution of their behavior. Figure 1(f) shows the structure and behavioral state of the network after the 25th iteration.

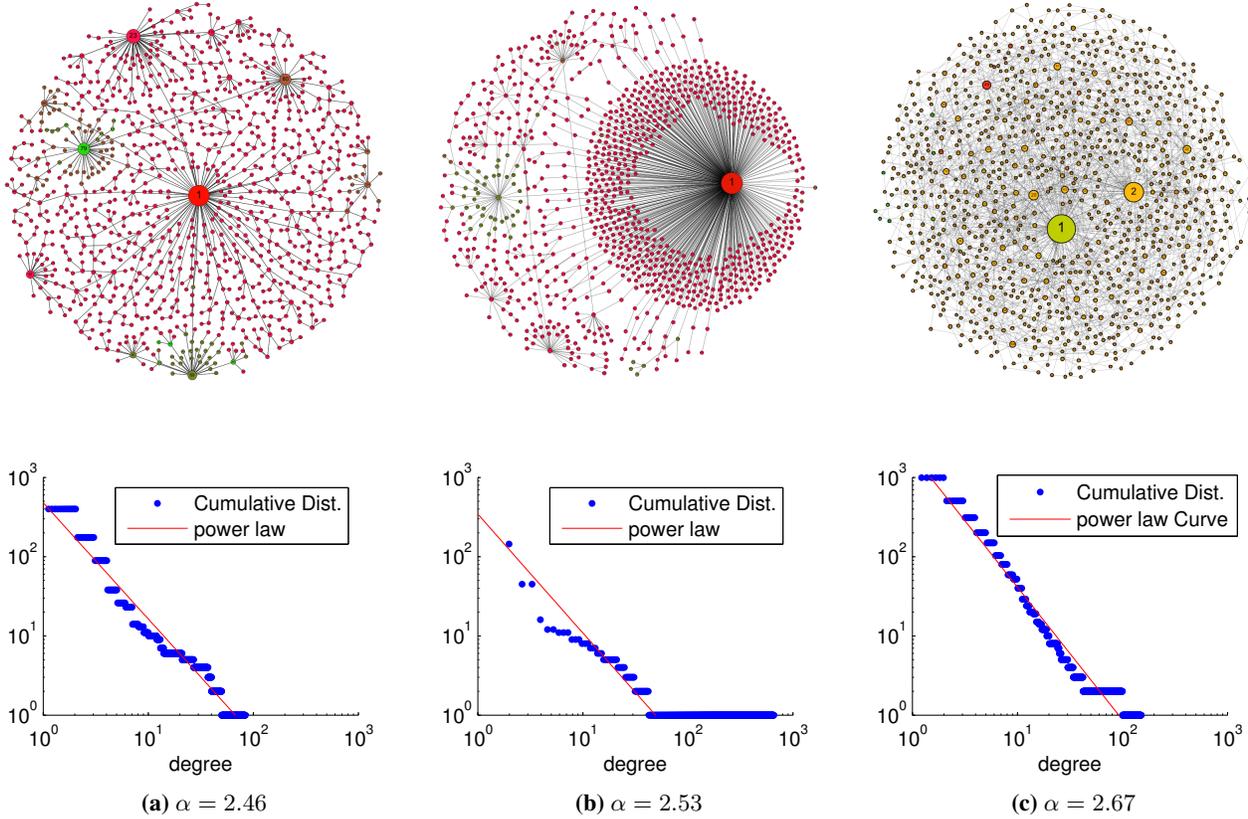


Figure 2: Sample network topologies generated by the SEE model (a) and (b) for $m = 1$ and (c) for $m = 2$, after 1000 iterations. The state and degree of the individuals are denoted by the color (red for pure defection to green for pure cooperation) and size of the nodes. The cumulative degree distribution of each network, shown as blue dots, shows how close this network follows a power law curve, shown as red line, with exponent α .

Experiments and Results

In this section we first illustrate sample networks generated using the proposed SEE model, and show the scale-free characteristics that emerge. Hereafter, the cumulative degree distribution of 8000 different networks generated for 8 different settings of the SEE model will be studied in detail by computing the power law exponent in these networks. Finally, the evolution of cooperation resulting from the proposed SEE model is compared to the standard Barabási-Albert model of preferential attachment. In all experiments, the upper limit for network size is set to $N_{max} = 1000$. The number of links added for each new individual, m , is set to either 1 or 2 (indicated where applicable). In the CAIPD model the step size is 0.1; $b = 4$, $c = 1$ and $\beta = 1$.

Sample Networks Generated by the SEE Model

Consider an evolution ratio R_{Evo} of 1 in a network that initiates from a single individual which is set initially to either pure defection or pure cooperation. When applying the SEE model, various different network structures can be expected to emerge, as there is stochasticity involved in both initial-

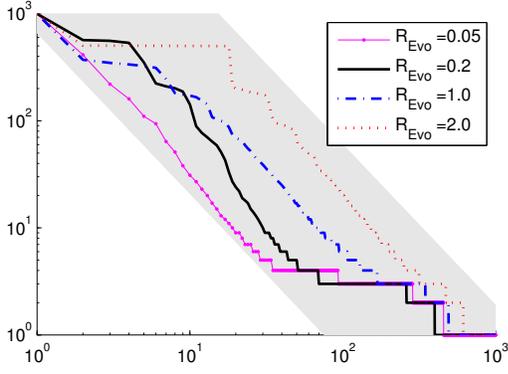
ization of the nodes' states and their attachment. Three samples of such networks are illustrated in the top portion of Figures 2(a)-(c).

In order to study whether the networks generated by the SEE model follow a power law degree distribution, the cumulative degree distribution, i.e., the number of nodes with degree greater than or equal to k , of the sample networks in Figures 2(a)-(c) are shown on a log-log scale¹. The results indeed show a power-law degree distribution with exponent close to $\alpha = 2.5$ for the sample networks in Figure 2. Next, we study the average cumulative degree distribution of networks generated by the SEE model in more detail.

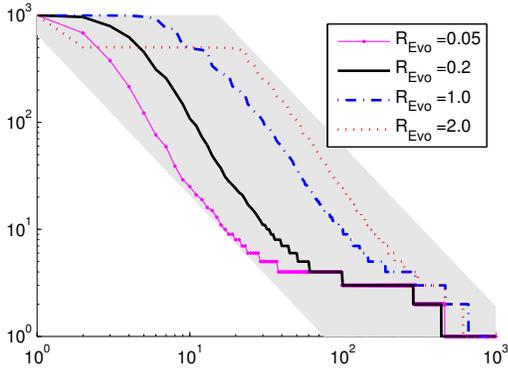
Degree Distribution in the SEE Model

In this section, we provide an empirical study on a large number of different networks generated using the proposed SEE model. We analyse different settings for R_{Evo} , ranging from 0.05 (slow evolution) to 2 (fast evolution). Figures 3(a)

¹The cumulative distribution of a power law distribution is also power law but with an exponent $\alpha - 1$, i.e., $\int cx^{-\alpha} = \frac{c}{1-\alpha}x^{-(\alpha-1)}$, where c is the power law coefficient



(a) $m = 1$



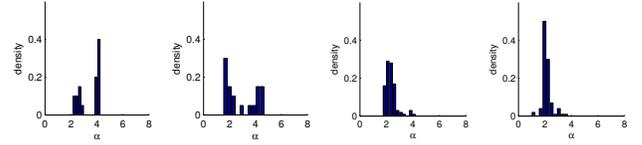
(b) $m = 2$

Figure 3: Cumulative Degree distribution of nodes for different evolution rates. The gray region contains the power law distributions corresponding to the scale free networks with exponent $\alpha = 2.5$ and different coefficients of the power law.

and (b) show the average cumulative degree distribution of these networks for $m = 1$ and $m = 2$, respectively. For each combination of settings, the SEE model is run 1000 times, with initial nodes randomly set to either cooperation or defection, and the results are averaged.

It can be observed from Figures 3(a) and (b) that the networks that emerged using the SEE model, on average, largely follow a power law degree distribution with exponent close to $\alpha = 2.5$. When evolution is slow (i.e., $R_{Evo} \rightarrow 0$) the power law is less clearly present, in particular towards the high end of the degree distribution. A possible explanation is that, as the CAIPD evolution slows down, the fitness of the nodes gets updates less frequently as there are fewer interactions. Hence, having more neighbors does not immediately translate to a potential higher fitness.

To get a more detailed insight, the distribution of the exponent of power law distribution that is fit to the constructed networks is illustrated in Figures 4(a)-(d) and 4(e)-(h) for

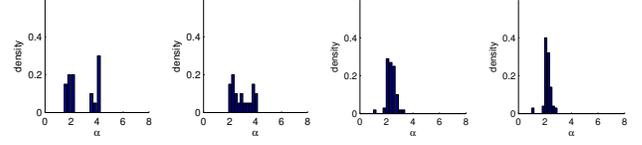


(a) $m = 1$
 $R_{Evo} = 0.05$

(b) $m = 1$
 $R_{Evo} = 0.2$

(c) $m = 1$
 $R_{Evo} = 1.0$

(d) $m = 1$
 $R_{Evo} = 2.0$



(e) $m = 2$
 $R_{Evo} = 0.05$

(f) $m = 2$
 $R_{Evo} = 0.2$

(g) $m = 2$
 $R_{Evo} = 1.0$

(h) $m = 2$
 $R_{Evo} = 2.0$

Figure 4: Distribution of the power law exponent for various evolution ratios in the SEE model.

$m = 1$ and $m = 2$, respectively. These figures show that the SEE networks with slow evolution rates exhibit power law degree distribution with exponents $1 < \alpha < 5$ (i.e., some of the networks fall outside the range of typical real-world complex networks). Increasing the evolution ratio shrinks the range of α values that are observed and centers their distribution around $\alpha = 2.5$, yielding realistic scale-free networks. Moreover, it is interesting to note that a bifurcation seems to take place when R_{Evo} decreases (in Figures 4(a), (b) and (e)): the distribution of power law exponents splits into two parts with their mass centered around 2 and 4. This phenomenon warrants a closer inspection in future work.

SEE model vs. Barabási-Albert model

In the previous section the scale-free characteristic of the SEE model was studied and it was shown that the degree distribution of these networks follows a power law degree distribution with $\alpha \approx 2.5$. In this section we study the influence of the SEE model on the evolution of behavior in the network. We compare the proposed SEE model, which uses preferential attachment based on fitness (see Equation 6), with the standard Barabási-Albert model that uses the degree (see Equation 5). For all experiments, $N_{MAX} = 1000$, and the evolution ratio R_{Evo} is set to 1.

Figures 5(a) and (b) show the evolution of cooperation under the SEE model, specifically the figures show the final cooperation level in the network depending on whether the initial nodes were either defectors or cooperators. Similarly, Figures 5(c) and (d) show the same results when the Barabási-Albert (B-A) model is used for the preferential attachment. It is clear from these figures that the final cooperation level in the network greatly depends on the initial state of the first individuals. When the initial nodes are cooperators, the network tends to cooperate to a large (> 0.5) degree, whereas the situation reverses when the initial nodes

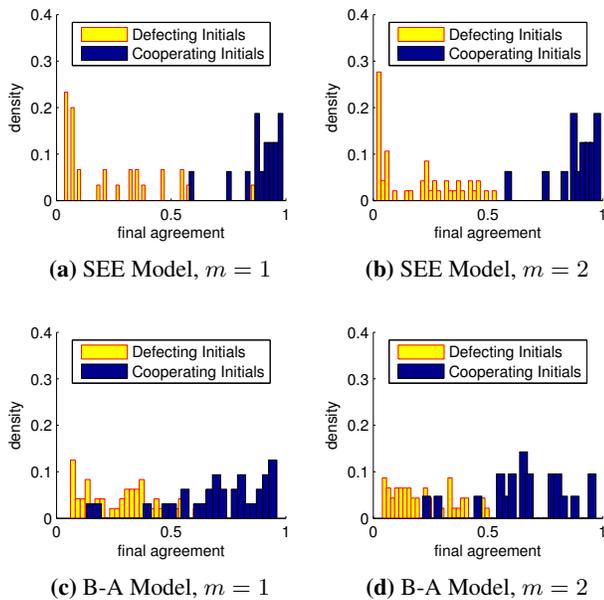


Figure 5: Final degree of cooperation as a function of the initial members’ state in the SEE model in (a) and (b) and the Barabasi-Albert (B-A) model in (c) and (d). The colors indicate the state of the initial members: blue for cooperators, and yellow for defectors.

are defectors. This effect is most clear under the SEE model, where a large fraction of the networks eventually reaches either a high (≈ 1) or low (≈ 0) degree of cooperation. When preferential attachment according to Barabasi-Albert is used, this effect is less strong. Here we observe a broad mix of final cooperation levels; moreover the divide between initial cooperators and initial defectors is less clear.

Clearly, the above results demonstrate that final agreements depend on the initial state of the first individuals in the network. This phenomenon is manifested in both the SEE and the Barabasi-Albert model. Having proposed a generalized and formal framework for analysing evolution of cooperation and network emergence, this aspect constitutes a major direction in our future work, where SEE can be used to acquire analytical conclusions describing the effects of such a dependence.

Conclusions

The recent interest to study social networks and their behavior has led to many studies, which can roughly be divided in two streams. The first stream has studied the emergence of these networks, and has in particular tried to find generative models that can explain certain structural properties of real-world complex networks, such as a scale-free degree distribution. The second stream of research has focussed on the evolution of behavior on such networks, when the nodes represent individuals that interact according to some rules. Most notably, interest has been in the evolution of coopera-

tion in social networks, aiming to identify properties of both the network and the interactions that sustain cooperation.

This paper aims to unify these two streams, by studying the simultaneous evolution of behavior on a social network, and the structural emergence of the network itself. The Simultaneous Emergence and Evolution (SEE) model proposed in this paper combines a modified version of preferential attachment, used to generate scale-free networks, with the continuous action iterated prisoner’s dilemma (CAIPD) model, describing the evolution of cooperation. Using the proposed model, a number of different networks, emerging from different initial conditions, have been studied. It has been shown that the SEE model yields realistic scale-free networks, despite the fact that the preferential attachment is based on individual’s fitness rather than degree.

Moreover, results show that that both structural emergence and behavioral evolution are intertwined, mutually influencing each other, and should therefore be studied in tandem. Aiming at a better understanding of such phenomena, the SEE model provides a fundamental and general framework that allows the analysis of these processes as they co-evolve.

An interesting direction for future work is to include the possibility of rewiring as well in the SEE model, whereby existing nodes may break or create links at any time.

References

- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–6.
- Backstrom, L., Boldi, P., and Rosa, M. (2011). Four degrees of separation. *Arxiv preprint arXiv:1111.4570*.
- Barabasi, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–12.
- Barabasi, A.-L. and Oltvai, Z. N. (2004). Network biology: understanding the cell’s functional organization. *Nature reviews. Genetics*, 5(2):101–13.
- Bloembergen, D., Ranjbar-Sahraei, B., Ammar, H. B., Tuyls, K., and Weiss, G. (2014). Influencing social networks: An optimal control study. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI2014)*.
- Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews. Neuroscience*, 10(3):186–98.
- Chapman, M., Tyson, G., Atkinson, K., Luck, M., and McBurney, P. (2012). Social networking and information diffusion in automated markets. In *Joint International Workshop on Trading Agent Design and Analysis (TADA) and Agent-Mediated Electronic Commerce (AMEC)*, pages 1–14.
- Easley, D. and Kleinberg, J. (2010). *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press.

- Ghanem, A. G., Vedanarayanan, S., and Minai, A. A. (2012). Agents of influence in social networks. In *Proc. of 11th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2012)*.
- Gibbons, R. (1992). *A Primer in Game Theory*. Pearson Education.
- Hofmann, L.-M., Chakraborty, N., and Sycara, K. (2011). The Evolution of Cooperation in Self-Interested Agent Societies: A Critical Study. *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 685–692.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Killingback, T. and Doebeli, M. (2002). The continuous prisoner’s dilemma and the evolution of cooperation through reciprocal altruism with variable investment. *The American Naturalist*, 160(4):421–438.
- Kossinets, G. and Watts, D. J. (2006). Empirical analysis of an evolving social network. *Science*, 311(5757):88–90.
- Newman, M. E. (2005). Power laws, pareto distributions and zipf’s law. *Contemporary physics*, 46(5):323–351.
- Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829.
- Ohtsuki, H., Hauert, C., Lieberman, E., and Nowak, M. A. (2006). A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505.
- Ranjbar-Sahraei, B., Ammar, H. B., Bloembergen, D., Tuyls, K., and Weiss, G. (2014). Theory of cooperation in complex social networks. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI-14)*.
- Ranjbar-Sahraei, B., Bou-Ammar, H., Bloembergen, D., Tuyls, K., and Weiss, G. (2014). Evolution of Cooperation in Arbitrary Complex Networks. In *13th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 677–684.
- Santos, F. and Pacheco, J. (2005). Scale-Free Networks Provide a Unifying Framework for the Emergence of Cooperation. *Physical Review Letters*, 95(9):1–4.
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. (2006). Cooperation prevails when individuals adjust their social ties. *PLoS computational biology*, 2(10):e140.
- Ugander, J., Karrer, B., Backstrom, L., and Marlow, C. (2011). The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*, pages 1–17.
- Van Segbroeck, S., de Jong, S., Nowe, A., Santos, F. C., and Lenaerts, T. (2010). Learning to coordinate in complex networks. *Adaptive Behavior*, 18(5):416–427.
- Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442.
- Zimmermann, M. G. and Eguíluz, V. M. (2005). Cooperation, social networks, and the emergence of leadership in a prisoner’s dilemma with adaptive local interactions. *Physical Review E*, 72(5).