

Self-Organizing Techniques for Knowledge Diffusion in Dynamic Social Networks

Luca Allodi, Luca Chiodi, and Marco Cremonini

Abstract In this paper, we model a knowledge diffusion process in a dynamic social network and study two different techniques for self-organization aimed at improving the average knowledge owned by agents and the overall knowledge diffusion within the network. One is a weak self-organization technique requiring a system-level central control, while the other is a strong self-organization technique that each agent exploits based on local information only. The two techniques are aimed at increasing the knowledge diffusion by mitigating the hype effect and the network congestion that the system dynamics shows systematically. Results of simulations are analyzed for different configurations, discussing how the improvements in knowledge diffusion are influenced by the emergent network topology and the dynamics produced by interacting agents. Our theoretical results, while preliminary, may have practical implications in contexts where the polarization of interests in a community is critical.

1 Introduction

Dynamic social networks represent a multidisciplinary research strand that has been studied with increasing interest since it emerged from sociology and complex systems research. Social networks exhibit peculiar characteristics with respect to non-social networks, in particular regarding the degree correlation of adjacent nodes and the clustering [1–3]. With respect to node degree, social networks are typically

Luca Allodi
DISI - University of Trento, Via Sommarive 5 I-38123 Povo (Tn), e-mail: luca.allodi@disi.unitn.it

Luca Chiodi
DI - University of Milano, Via Bramante 65 I-26013 Crema (Cr), e-mail: luca.c.chiodi@gmail.com

Marco Cremonini
DI - University of Milano, Via Bramante 65 I-26013 Crema (Cr), e-mail: marco.cremonini@unimi.it

assortative, meaning that the degree correlation of adjacent nodes is positive, i.e., nodes of high degree tend, on average, to be connected with other nodes of high degree. The second peculiar characteristic, *clustering*, has been defined in term of network transitivity, that is, given an edge between a pair of nodes A and B and another edge between nodes A and C , a network is said to be highly transitive if it is likely that there will also be a connection between nodes B and C [2]. For social networks, it has been observed that the clustering coefficient is typically greater, possibly orders of magnitude greater, than in typical random graphs [6, 7].

In the literature, many works have applied models of social networks to real case studies. The email exchange in a community of people, for instance, represents a relevant case study [4], as well as the dynamics showed by individuals joining and leaving groups of interests, which may stem from leisure (e.g. the case of online games) to scientific research or corporate projects [5].

In this paper, we model a social network with a fixed number of nodes (i.e. in the following we will refer to *agents* interacting according to some rules, rather than nodes and edges of the network) and mechanisms for the dynamic creation of connections among them. Our model shows the emergent characteristics of typical social networks, such as *assortativeness* [1, 2], *high clustering coefficients* [6, 7], and transition from several clustered communities of agents to a giant interconnected component [8]. The process of knowledge diffusion consists of agents that know a variable set of *topics*, each one characterized by an *interest* and a *quality*. Agents interact by selecting, first, a topic based on their own interests and then the agent in their neighborhood that owns the best quality associated to the chosen topic. In this sense, knowledge spreads from agents knowing more about a topic to those that know less. At network level, each time an interaction between two agents succeeds a directed link, from the requestor to the respondent agent, is created, if it were the first interaction, or the link's weight is increased with the number of exchanges. This interaction model is similar to important models of network influence based on belief and interest, such as the seminal one by DeGroot [9] and of knowledge diffusion, such as the one discussed by Cowan *et al.* [10]. The aim of this work is to study two self-organization strategies, *weak* and *strong self-organization*, according to the definition introduced in [11], aimed at improving knowledge diffusion, both based on the communication efficiency as the heuristic observed—at network level, in one case, at agent level in the other—to adjust the behavior. The effects of the two strategies are compared to the natural network behavior under different configurations. Our results exhibit interesting similarities with research in different settings and scenarios. In particular, works on organizational learning and parallel problem solving [12, 13] have showed how agents that learn too fast can reduce total system knowledge.

2 Related work

Works more closely related to ours are those that have investigated the knowledge diffusion in networked environments and that have described self-organizing systems. In particular, Cowan and Jonard, two economists, modelled and analyzed the dynamics of knowledge diffusion in a multi-agent scenario [14]. They showed how network structure affects the dynamics of knowledge diffusion, and they demonstrated that the average knowledge is maximal when the structure is a small-world. Their approach and, qualitatively, many of their results are closely related to our work. Different from them, ours is a dynamic model and we studied the emergent behavior of the system, while in their case a number of static network configurations have been considered. From a different field, but equally important for our research, the work by Brun *et al.* [15] has analyzed the importance of feedback loops in designing self-adaptive systems. Their findings, although not explicitly addressed at modelling dynamic social networks, are relevant in our context. Similarly to them, our self-organizing strategies exploit a control loop, either at system level or at agent level, to adjust the behavior. Walter *et al.* presented a model closely related to ours, although based on a static network and different in the research goal [16]. They considered a model of trust-based recommendation system on a social network, which assumed the transitivity of trust along a chain of relationships connecting agents. Differently from them, we admit only a limited degree of trust transitivity (which is restricted to the best friend-of-friends). Important for the analysis of mixing patterns and community structures in networks is the work by Girvan and Newman [8]. This research analyzed most of the characteristics that our model of social network presents and that we have tested and discussed in this work, from the assortative mixing to the formation of communities, from the relevance of friend-of-friend relationships to the dynamics of the growing network.

3 Model Description

We consider a set of N agents, n_1, n_2, \dots, n_N , each one characterized by a *Personal state* PS_{n_i} (what n_i knows) and a *Friend state* FS_{n_i} (who n_i knows). The *Personal state* has the form $PS_{n_i} = (\bigcup_{j \in T_i} (\text{topic}_j, \text{quality}_{i,j}, \text{interest}_{i,j}))$, where T is the set of topics that the population knows; each agent n_i knows a variable subset of them, $T_i \subseteq T$. The *Friend state* has the form $FS_{n_i} = (\bigcup_{j \in N_i} (n_j, \text{answers}_{i,j}))$, where n_j are the identifiers of agents connected with n_i and $\text{answers}_{i,j}$ is a counter to keep track of the number of interactions with each peer. The setup has been defined to be the most neutral, with topics T_i assigned to each agent and associated qualities selected randomly, interests distributed uniformly and no friends. More specifically:

Topics: A random set T_i of *topics* is defined for each agent. The maximum number of topics assigned to the agents can be limited by setting the maximum rate $\lambda_T \in (0, 1]$, so that $|T_i| \leq \lambda_T \cdot |T|$.

Quality and Interest: The *quality* associated to each topic of an agent’s Personal state is set to a random value in $[1, 100]$. For the *interest*, the initial value is equally distributed among all topics, and is calculated as $100/|T_i|$.

Friends: Agents start with no friends at setup, therefore, in the early stage of network formation, the local search fails and the selection of peers turns to the random choice.

Topic 0: To permit the selection of new topics, a dummy topic called *topic*₀ is always held by each agent. If *topic*₀ is selected, as for all ordinary topics, the system choose, randomly, another topic *not belonging* to the agent’s Personal state. Next, the agent looks for a peer based on the new topic. The *quality* associated to *topic*₀ is always zero, while the *interest* is calculated as for the other topics during a simulation.

3.1 Network Construction

The network is dynamically formed according to the following steps:

1. Given agent $n_{i'}$, select a topic (*topic* _{j_*}) in the Personal state. The choice of the topic is a weighted random selection with values of the associated interests (*interest* _{i,j_*}) as weights, this way topics with higher interest are more likely to be selected;
2. Among $n_{i'}$ friend agents and their "best friend" holding topic (*topic* _{j_*}), select agent $n_{i''}$ with maximum value of topic’s quality (*quality* _{i,j_*});
3. If *quality* _{i'',j_*} > *quality* _{i',j_*} then the communication takes places and agent $n_{i'}$ increases *quality* _{i',j_*} of *topic* _{j_*} ;
4. Otherwise, if none holds *topic* _{j_*} or exhibits a topic’s quality greater than that of agent $n_{i'}$, then select an agent $n_{i'''}$ randomly among the whole population;
5. if $n_{i'''}$ holds *topic* _{j_*} and *quality* _{i''',j_*} > *quality* _{i',j_*} , then the communication takes places and *quality* _{i',j_*} increases, otherwise the communication fails.

Best friend-of-friends. Given agent $n_{i'}$, and a selected *topic* _{j_*} , for each of its friends, the "best friend" agent is the one owning *topic* _{j_*} and the higher value of the *answer* attribute. The reason for this setup is that we consider unrealistic in a social context to scan all agents with a distance of 2 from the one selected. The selection based on the *answer* attribute represents a basic form of transitive trust. It is worth noting that the inclusion of "best friends" fosters network transitivity and the formation of triads, two key characteristics of social networks.

Start up. At start up, agents have no connection with others (i.e., Friend state is empty). When, for an agent, the 5-steps algorithm is executed, a topic is selected in *Step1*, then *Step2* and *Step3* fail and in *Step4* a random agent is selected. If *Step5* succeeds, then the connection is established. This mechanism triggers the network formation at start up.

3.2 State Update

After a successful interaction, the agent that started the communication is updated. The *quality* and the *interest* of the topic for which the communication took place increase and the other interests, associated to the other topics owned by the agent, are decreased. We decided that the topic's quality increases with increasing marginal increments, according to the assumption that an agent distrusts another one when they interact for the first time and this distrust progressively diminishes as interactions occur. The discount starts at a given value (i.e. ρ) and goes to zero exponentially. Motivations for this assumption could be found in the literature about information aggregation [17] and collective behavior [18] and refers both to the prevalence of egocentrism in assimilating new information and to trust dynamics. The *quality gain* obtained by agent n_i is:

$$\delta quality_{j^*} = \frac{quality_{i',j^*} - quality_{i',j^*}}{\gamma + \rho e^{-\frac{x}{\theta}}} \quad (1)$$

with: $\gamma \geq 1$ setting the nominal fraction of $\delta quality$ that n_i could learn from another agent; x the value of the attribute *answers* representing the number of past interactions that agent n_i had with agent $n_{i'}$; ρ the initial discount; θ the factor that controls the rate at which agents increase their trust towards the others.

The dynamics we have assumed for the *interest* associated to the topic for which the interaction took place is similar to that of the quality, but with two important differences: It only depends on the $\delta quality$ value and, accordingly, all other interests on topics owned by the agent decreases (studies in cognitive science have showed the tendency of people to shift their attention and interest, rather than behave incrementally [18]). The function is:

$$\delta interest_{i',j^*} = \alpha \left(1 - e^{-\frac{\delta quality_{i',j^*}}{\beta}} \right) \quad (2)$$

with $\alpha > 1$ and $\beta > 1$ the two parameters that control, respectively, the scale and slope of the interest increase.

Parameter β is key for our following analysis of the knowledge diffusion and the self-organization strategies. Self-organizing mechanisms could tune β to reduce ($\beta \uparrow$) or increase ($\beta \downarrow$) the speed at which the agent's interests change.

Finally, all interests associated to topics different from $topic_{j^*}$ are reduced by $\delta interest_{i',j \neq j^*}(t_k, t_{k-1}) = \delta interest_{i',j^*}(t_k, t_{k-1}) / (|T_i| - 1)$, that is the value of the interest gain for $topic_{j^*}$ at t_k divided by the number of topics $|T_i|$ minus one.¹

¹ The interest reduction applies to $topic_0$ as well, which is included in the total number $|T_i|$ of topics known by agent n_i .

3.3 Metrics.

Three metrics have been defined: *Communication Efficiency (CE)*, *Average Knowledge (AK)* and *Knowledge Diffusion (KD)*. The first one is a heuristic observed either at system-level by a central control process or locally by each agent evaluating its own behavior. *CE* measures how often agents are able to successfully interact with others with respect to the number of requests they made during a simulation. Γ is the number of requests made by all agents:

$$CE = \frac{\text{Total No. of Answers}}{\text{Total No. of Requests}} = \frac{\sum_{i=1}^N \sum_{j \in N_i} \text{answers}_{i,j}}{\Gamma} \quad (3)$$

The meaning is that if $CE = 0$, then there has been no communication since every interaction failed; if $CE = 1$ every interaction succeeded.

AK and *KD* metrics are used to evaluate two different characteristics of the knowledge diffusion dynamics and to analyze its efficiency and the benefits of the proposed self-organizing techniques. *AK* is calculated as the average quality with respect to the topics actually owned by agents. *KD*, instead, is the average quality with respect to the case of perfect diffusion of knowledge (i.e., all agents holding all topics). While *AK* is maximised by increasing only the average quality of each agent, regardless of the number of known topics, *KD*, instead, depends from the diffusion of topics among agents.

$$AK = \frac{\sum_{i=1}^N \sum_{j=1}^{|PS_i|} \text{quality}_{i,j}}{\sum_{i=1}^N (|PS_i| - 1)} ; KD = \frac{\sum_{i=1}^N \sum_{j=1}^{|PS_i|} \text{quality}_{i,j}}{|N| \times (|T| - 1)} \quad (4)$$

4 Network Simulations

During simulations, the number of agents ($N = 100$) and duration ($\Gamma = 50000$) have been kept constant. Key parameters that were varied are: the *number of topics in the network* $|T|$, the *maximum rate of topics assigned to agents at setup* λ_T (e.g., $\lambda_T = 0.1$ means that at setup agents know *at most* 10% of topics T), and β defined in Equation 2. Combined, they deeply influence the emergent network structure and the aggregation of agents in communities. In our simulations, the parameter λ_T affects the mean network degree. The typical transition [8] from small communities to the giant one happens for $\lambda_T \approx 0.55$.

Figure 1 shows the results for four typical base configurations. In all cases, *CE* has an initial spike (i.e. the *hype effect*) as a result of polarization of agent interests, which, due to the positive feedbacks of successful interactions and the quality increase, tend to exhibit bursts of interaction with the same peer for the same topic. While this effect greatly increases the performance of the network in the early stages, it also quickly dissipates and *CE* abruptly drops until agents start rebalancing their interests by choosing other topics. The dynamics of the communication has

visible effects on KD : An inefficient communication, in general, implies a slower diffusion of knowledge within the network, as showed by the values of KD (see case A with C and case B with D). Similar results, although on a different model (based on broadcasting in a static network), have been presented in [10]. With respect to AK , instead, the four cases do not exhibit relevant differences. In all cases, at the end of the simulation, the average knowledge owned by agents has reached a level close to 80%, meaning that, limited to the topics they individually own, they have reached good quality. As a final remark to this set of results, we stress the fact that where AK and KD strongly differ (configurations A and B are those with the largest differences), the network behavior is inefficient with respect to the goal of diffusing knowledge both in quality *and* in quantity. In those cases agents, on average, have improved on topics they owned at setup, but did not acquire much knowledge about new ones.

5 Weak and Strong Self-organization

The analysis of these configurations has showed that reducing the tendency of agents to polarize, caused by interests on just a few topics that grow too fast, communication and knowledge diffusion generally improve. Self-organization mechanisms could modify at run-time some critical parameters. In particular, we consider the interest function and parameter β . The reason for choosing to change only β is because it has both a intuitive meaning in real-world contexts and a direct reference with previous research. Increasing β means, essentially, to reduce the speed of learning avoiding to concentrate on few issues only [13], or be driven more by our own belief than by information received from others [17].

To decide when β should be adjusted, we adopted a simple heuristic based on the dynamics of the CE : when its trend changes from increasing to decreasing, β is modified. With this rationale for our approach to self-organization, two strategies have been designed:

Weak Self-Organization is a system-level strategy that assumes the presence of a central control process able to observe the system dynamics of CE and to adjust β , a global parameter, at run-time when communication starts dropping.

Strong Self-Organization is a local strategy where each agent observes its own communication dynamics and acts on its own interest function. To this end, the model has been modified to introduce a *local communication efficiency* CE_{n_i} , for each agent, and a *local interest function*, with the same form of Equation 2, except for β_{n_i} , now specified for each agent, rather than as a global parameter. Another modification has been to set a threshold ω that automatically triggers the change of β_{n_i} for each agent.

The results of the system behavior with *weak self-organizations* are shown in Figure 2 (A-D). Qualitatively, we can see that with this simple technique the initial polarization of the network, with its negative effect, is not reduced. The benefit of the weak self-organization strategy becomes evident when the network has absorbed the

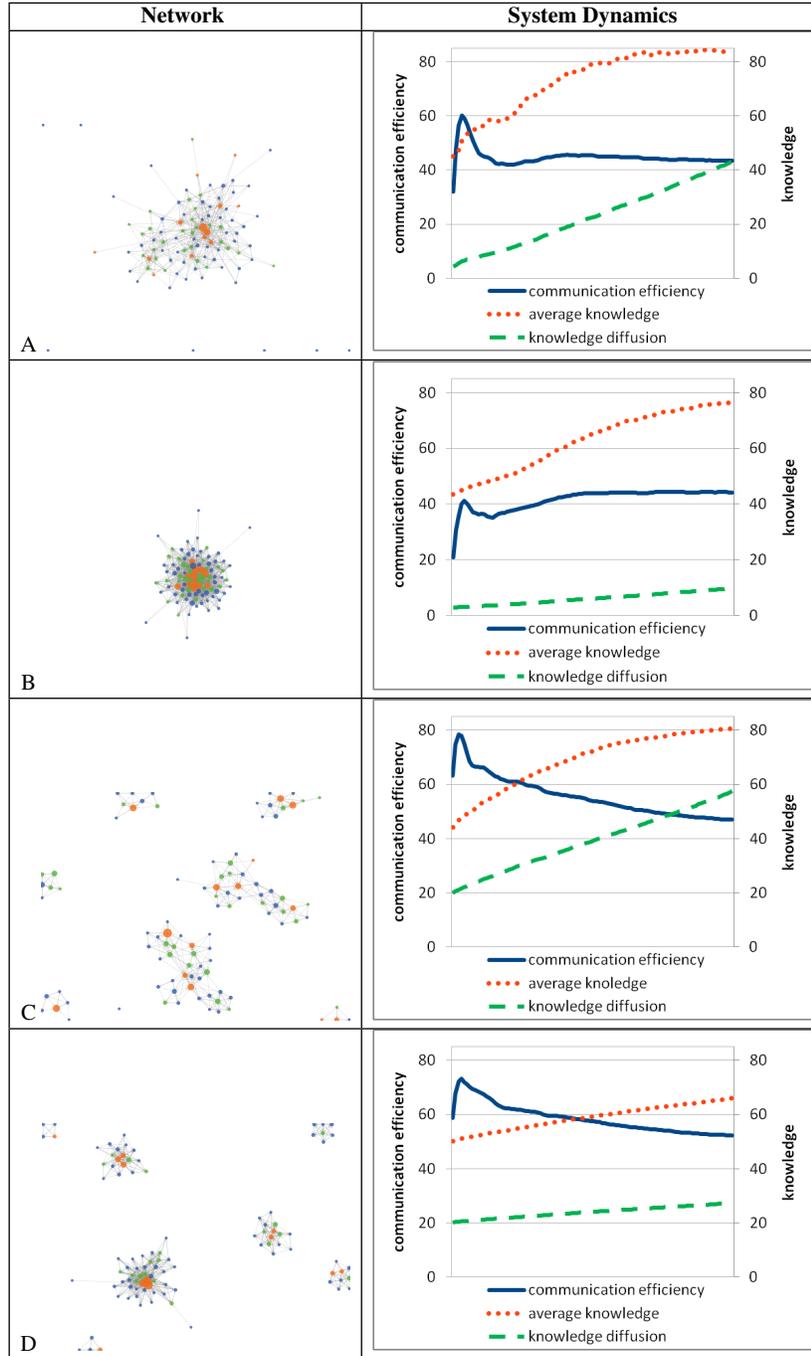


Fig. 1 System dynamics with different configurations. The x-axis represents simulation time (number of ticks T); the y-axis is a scale 0-100. Parameters: A) $\lambda_T = 0.1$, $|T| = 20$; B) $\lambda_T = 0.1$, $|T| = 100$; C) $\lambda_T = 1.0$, $|T| = 20$; D) $\lambda_T = 1.0$, $|T| = 100$.

excessive polarization following the spike and regains a higher level of efficiency. From that point on, the CE remarkably improves in all four cases, with respect to the corresponding ones of Figure 1.

The worsening of AK is a consequence of the lower clustering and the more frequent random choice of peers (i.e. compare A with C and B with D). This effect is amplified by adjusting β : Agents increment the number of topics, but do not sufficiently interact to improve their AK .

Combined, these results suggest that when knowledge diffusion is adjusted at system level by manipulating the interest dynamics, the number of information in the network $|T|$ with respect to the number of agents N is a critical characteristics, that not only strongly influences the original dynamics, but it also affects the efficacy of the self-organization solution. On the contrary, it could be worthless to force structural modifications, such as inducing the creation of larger communities or splitting giant ones.

Finally, it is remarkable to note how CE always increases of about the same rate in all four cases. This suggests that when self-organization strategies are applied and benefits should be evaluated, an increase in the communication efficiency should not be considered a reliable sign of better knowledge diffusion or a parameter suitable for comparative analyses.

In case of *strong self-organisation*, agents adjust their local β_{n_i} parameter by evaluating their local CE_{n_i} . For the simulations, the threshold, shared by all agents, is set to $\omega = 0.8$, meaning that when an agent sees its own CE_{n_i} dropping below 80%, then it adjusts β_{n_i} to 500; when, instead, it sees CE_{n_i} raising over 80%, it switches back to $\beta_{n_i} = 5$.

The results are showed in Figure 2 (E-H). The more complex dynamics of the strong self-organizing technique has clearly deeper effects on system behavior than the weak self-organization case, which lead to further possibilities to manipulate the evolution of a dynamic social network. By considering the numerical results of Table 1, we observe that reducing the initial spike produces variable benefits. With respect to AK , the network structure is clearly the dominant factor influencing the performances, with highly clustered configurations C and D increasing AK , while lower clustered ones A and B reducing AK . Similar results have been found in [14]. With respect to KD , the solution always produces a gain, but in this case the number of topics is the dominant factor. Configurations A and C with few topics increase KD of more than 16%, while configurations B and D with more topics increase KD of about 4-5%. In more detail, from Table 1, we note that the differences observed for the weak self-organization case between configurations with few topics A and C and configurations with more topics B and D are confirmed in case of strong self-organisation. The same holds for the differences between network structures, A and B having one main component with respect to C and D with small clusters. With respect to KD , there is a tendency of the strong self-organization technique to outperform the weak one, in all configurations except D , meaning that there is a better knowledge diffusion. The opposite holds with respect to AK , due to the reduced tendency of agents to polarize on just few topics. The reasons is that in the weak technique, the global parameter β is adjusted when the *average* communication ef-

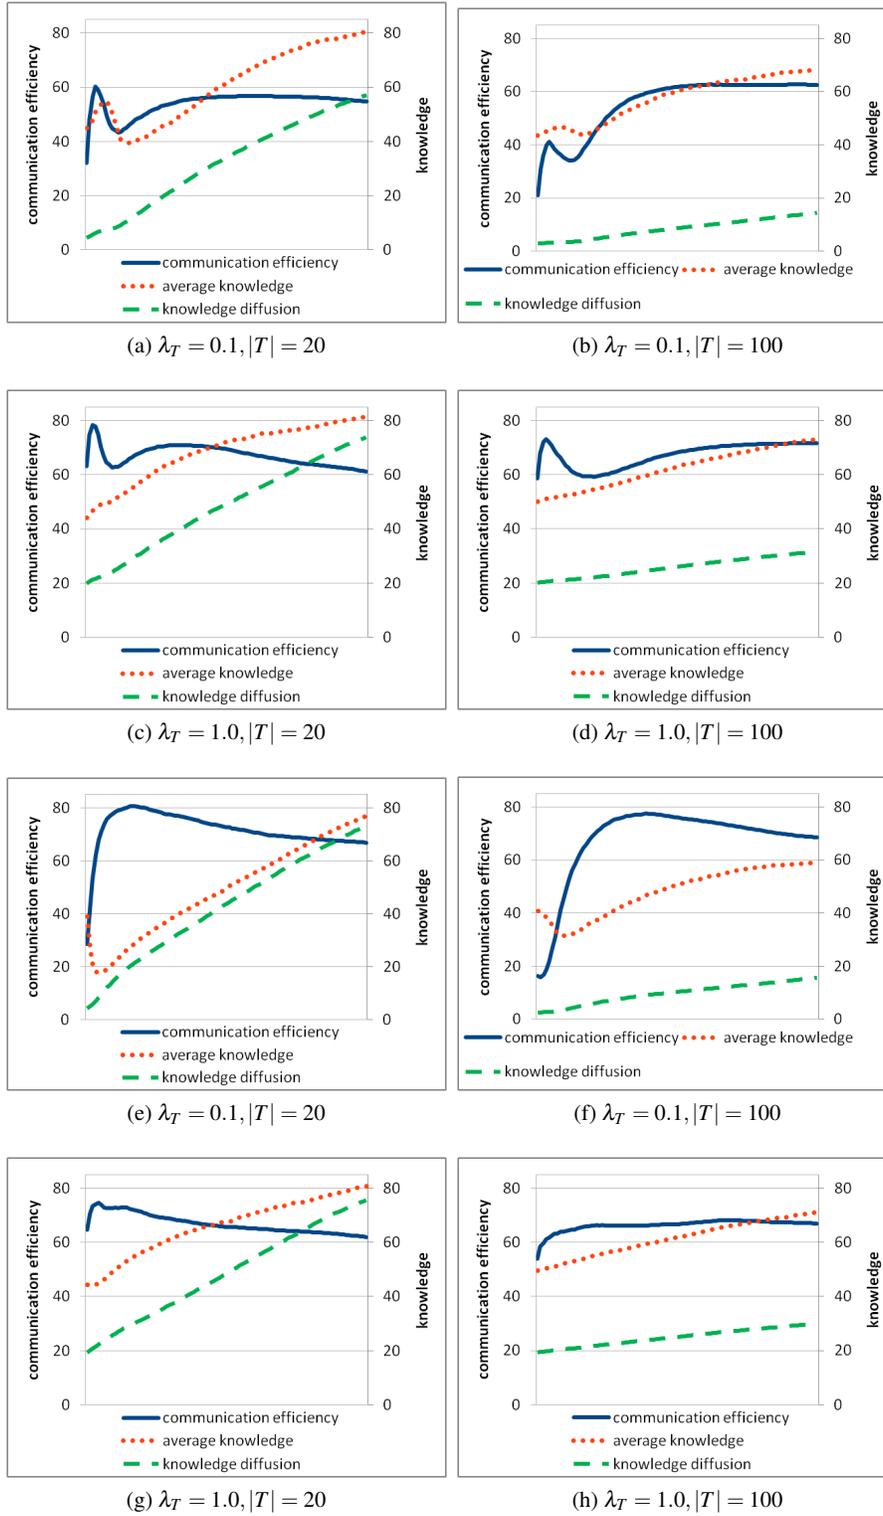


Fig. 2 (A-D): Weak self-organization; (E-H): Strong self-organisation.

efficiency of agents drops. This means that some agents could be already extremely polarized and almost unable to interact, while others could be still efficiently interacting. Heterogeneity of agents behavior affects the effectiveness of the weak self-organization technique. Differently, the strong self-organization techniques operates locally, therefore each agent adjusts its own behavior when needed. This has the effect of preventing extreme polarization.

6 Conclusions

In this paper we presented a model of dynamic social network based on knowledge exchange among agents. The results, although still preliminary, are promising and some strikingly similitudes with previous studies based on different assumptions and different network structures have been found.

In particular, it appears that social network analysis applied to different self-organization strategies could provide important insights for relating one strategy with another. In particular, our focus is on those phenomena that typically exhibit network congestion due to excessive agent polarization, which, as a consequence, exhibit an initial exceptional communication efficiency, followed by a steep decrease. Some well-known examples are: the *hype effect* typically present in the adoption cycle of new technologies [19], the formation of blockbusters in cultural markets [20] or the choice of news published by media [21]. As a final example, in school education, either in traditional classes or through e-learning systems, there is anecdotal evidence that polarization may emerge and could be detrimental for the overall level of knowledge acquired by students. This could be the case of excessive interest of students on few topics only with respect to a more balanced distribution of efforts and time for learning. In all these examples, intuitively, our model and self-organization techniques could be well-suited for describing the effects of manipulating the speed of interest grow.

	Base System			Weak Self-Organization			Strong Self-Organization		
	CE	AK	KD	ΔCE	ΔAK	ΔKD	ΔCE	ΔAK	ΔKD
A	39.50	77.44	37.60	+15.18%	-1.07%	+16.74%	+28.10%	-0.32%	+35.17%
B	44.12	76.14	9.68	+18.45%	-8.04%	+4.73%	+24.41%	-17.35%	+5.98%
C	46.96	80.41	57.41	+14.26%	+1.13%	+16.42%	+14.97%	+0.41%	+18.24%
D	52.39	66.11	27.36	+19.18%	+7.08%	+4.30%	+14.48%	+4.98%	+2.60%

Table 1 Results of weak and strong self-organization techniques.

References

1. E. M. Jin, M. Girvan, and M. E. J. Newman, "Structure of growing social networks," *Physical Review E*, vol. 64, no. 4, pp. 046132+, Sep. 2001.
2. M. E. J. Newman and J. Park, "Why social networks are different from other types of networks," *Physical Review E*, vol. 68, no. 3, 2003.
3. B. Skyrms and R. Pemantle, "A dynamic model of social network formation," *Proceedings of the National Academy of Sciences*, vol. 97, no. 16, pp. 9340–9346, 2000.
4. J. R. Tyler, D. M. Wilkinson, and B. A. Huberman, *Email as spectroscopy: automated discovery of community structure within organizations*. The Netherlands: Kluwer, B.V., 2003, pp. 81–96.
5. M. E. J. Newman, "Coauthorship networks and patterns of scientific collaboration," *Proceedings of the National Academy of Sciences*, pp. 5200–5205, 2004.
6. M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, no. 2, pp. 167–256, 2003.
7. D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, Feb 1998.
8. M. E. J. Newman and M. Girvan, "Mixing patterns and community structure in networks," in *Statistical Mechanics of Complex Networks*, ser. Lecture Notes in Physics, R. Pastor-Satorras, M. Rubi, and A. Diaz-Guilera, Eds. Springer Berlin / Heidelberg, 2003, vol. 625, pp. 66–87.
9. M. H. DeGroot, "Reaching a Consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
10. R. Cowan and N. Jonard, "Knowledge creation, knowledge diffusion and network structure," in *Economics with Heterogeneous Interacting Agents*, A. Kirman and J.-B. Zimmermann, Eds. Springer, 2001, vol. 503, pp. 327–343.
11. G. Di Marzo Serugendo, M. P. Gleizes, and A. Karageorgos, "Self-organization in multi-agent systems," in *The Knowledge Engineering Review*. United Kingdom: Cambridge University Press, 2005, vol. 20:2, pp. 165–189.
12. K. D. Miller, M. Zhao, and R. J. Calantone, "Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model," *Academy of Management Journal*, vol. 49, no. 4, pp. 709–722, 2006.
13. D. Lazer and A. Friedman, "The Network Structure of Exploration and Exploitation." *Administrative Science Quarterly*, vol. 52, no. 4, pp. 667–694, 2007.
14. R. Cowan and N. Jonard, "Network structure and the diffusion of knowledge," *Journal of Economic Dynamics and Control*, vol. 28, no. 8, pp. 1557–1575, June 2004.
15. Y. Brun, G. Di Marzo Serugendo, C. Gacek, H. Giese, H. Kienle, M. Litoiu, H. Müller, M. Pezzè, and M. Shaw, "Engineering self-adaptive systems through feedback loops," in *Software Engineering for Self-Adaptive Systems*, B. H. Cheng, R. Lemos, H. Giese, P. Inverardi, and J. Magee, Eds. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 48–70.
16. F. E. Walter, S. Battiston, and F. Schweitzer, "A model of a trust-based recommendation system on a social network," *Auton Agent Multi-Agent Syst*, vol. 16, no. 1, pp. 57–74, Feb 2008.
17. L. M. A. Bettencourt, "The rules of information aggregation and emergence of collective intelligent behavior," *Topics in Cognitive Science*, vol. 1, no. 4, pp. 598–620, 2009.
18. R. L. Goldstone and T. M. Gureckis, "Collective behavior," *Topics in Cognitive Science*, vol. 1, no. 3, pp. 412–438, 2009.
19. Gartner, "Gartner hype cycle." [Online]. Available: <http://www.gartner.com/technology/research/methodologies/hype-cycle.jsp>
20. D. J. Watts and S. Hasker, "Marketing in an unpredictable world," *Harvard Business Review*, September 2006.
21. Media Standards Trust, "Shrinking world: The decline of international reporting in the british press." [Online]. Available: <http://mediastandardstrust.org/publications/shrinking-world-the-decline-of-international-reporting-in-the-british-press/>