

# Infection-Based Norm Emergence in Multi-Agent Complex Networks

Norman Salazar and Juan A. Rodriguez-Aguilar and Josep Ll. Arcos<sup>1</sup>

**Abstract.** We propose a computational model that facilitates agents in a MAS to collaboratively evolve their norms to reach the best norm conventions. Our approach borrows from the social contagion phenomenon to exploit the notion of *positive infection*: agents with *good behaviors* become infectious to spread their norms in the agent society. By combining infection and innovation, our computational model helps a MAS establish better norm conventions even when a sub-optimal one has fully settled in the population.

## 1 Introduction

Norms have become a common mechanism to regulate the behavior of agents in multi-agent systems (MAS). They exist to balance agents' interests with respect to the society's in such a way that each agent can pursue its individual goals without preventing other agents'. However, learning and establishing an adequate set of norms is not trivial. This process, usually referred to as either self-organization or emergence.

In societies, conventions result when members agree upon a specific behavior. Thus, a norm convention refers to a set of norms that has been established among the members of a society. One of the trends of thought in social studies is that norm conventions emerge by propagation or contagion, where social facilitation and imitation are key factors [2, 1]. From a MAS perspective, the studies in [8] [7] show that norm emergence is possible. However, these works limit to analyze norm propagation, leaving out norm innovation (discovery of new norms), a key factor for the evolution of societies. When the aim is to help a MAS establish conventions in dynamic environments, propagating norms may not be enough since propagation assumes that at least some agent in the society knows the correct set of norms, which is not always the case. Additionally, the problem can become even more difficult when the aim is not only to establish (any) convention(s), but the best convention(s).

We propose an evolutionary computational model that facilitates agents in a MAS to collaboratively evolve their norms to reach the best norm conventions for a wide range of interaction topologies. At this aim, we take inspiration on the argument in the social sciences literature that behavior conventions arise from a social contagion [1]. Although further evolutionary approaches appear in the literature [4], they are usually applied either: (i) as a centralized process; or (ii) as an individual self-contained process for each agent. Both approaches can be potentially slow and tend to be off-line processes, and thus unsuitable to our purpose of dynamically adapting norms.

## 2 An Evolutionary Infection-Based Model

We propose a computational model that helps agents in a MAS reach *norm conventions* that maximize the social welfare. At this aim, we assume, in line with the distributed nature of the problem, that we can achieve our goal by maximizing agents' *individual welfares*.

The social sciences literature argues that conventions in societies are reached through social contagion [1]: behaviors spread between individuals akin to an infectious disease. Hence, we chose to model the social contagion into a MAS framework. However, we target beneficial conventions that if possible tend to maximize the social welfare. Considering the social welfare as a composition of individual welfares, it makes sense to let the individual behaviors that impact positively on it, here named *good behaviors*, be more infectious. Nevertheless, positive infection at most achieves a total replication of the best-known behavior among agents. Therefore, we also require a norm innovation mechanism. Hence we expect that a MAS can reach norms that are dominant in the society so that no better ones can be found and no worst ones can upstage them. However, if some unaccounted factor(s) alter(s) the MAS so that the current norms become obsolete (the social welfare deteriorates), the infectious process will re-configure the norms toward a better social welfare.

We propose an evolutionary algorithm (EA) approach that helps agents in a MAS reach the best norm conventions. In our infection-based EA, each agent has genes that encodes its behavior. Agents can infect other agents with their genes following the *survival of the fittest* concept: the highest its individual welfare, the more infectious. Furthermore, it realizes innovation (exploration) by letting agents mutate their genes. This process runs distributedly: each agent decides whether to infect or mutate based on local knowledge. Thus, each agent is endowed with: i) an **evaluation function** to assess its individual welfare; ii) a **selection process** to choose a peer to infect, out of its local neighborhood, based on its fitness; iii) an **infection operator** to inject some of its genes into the selected agent; and iv) an **innovation operator** to mutate its genes to create new behaviors.

## 3 Empirical Results

Agents in a MAS interact with each other by engaging in iterative games with multiple rounds. During a round each agent randomly selects a neighbor agent to play with (an opponent). A play consists in both agents doing an action, either A or B (actions constrained by their current norms). Plays are rewarded with a payoff, which is accumulated after each game round. The payoff for as round can be: -1, 1 or  $\alpha$  based on the agent's current action and the action of the selected by the neighbor (different ones, both B and both A respectively). This payoff can help capture pure coordination games [8][7]

<sup>1</sup> IIIA, Artificial Intelligence Research Institute, CSIC, Spanish National Research Council, Spain, email:{norman, jar, arcos}@iiia.csic.es

( $\alpha = 1$ ) and coordination games with equilibrium differing in social efficiency [6] ( $\alpha > 1$ ).

Each agent,  $ag_i$  has two parameterized norms: one to help it decide what action to take based on the last opponent's past action; and another one to decide the action to take when no past action is known. To this end, the agent keeps on its memory the action performed by its last opponent without distinguishing who the opponent was. Thus, our model has the task of finding for each agent the norm parameters that maximize the social-welfare  $u$ .

It is well known that the behavior of infections is affected by the type of topology on which a population interacts [9, 5]. Therefore, in order to empirically analyze such effects in our infection-based model we chose the following interaction topologies: *small-world*,  $W_{1000}^{10,0.1}$ ; *scale-free*,  $S_{1000}^{10,-3}$ ; and *random graphs*,  $R_{1000}^{10}$ .

We know beforehand that four cooperative-only norms exist (norms that always try to cooperate), and also that they are the strongest attractors. Two of them always make agents do A (*A-conventions*) and the other two always do B (*B-conventions*). A-conventions give higher payoffs when  $\alpha > 1$ .

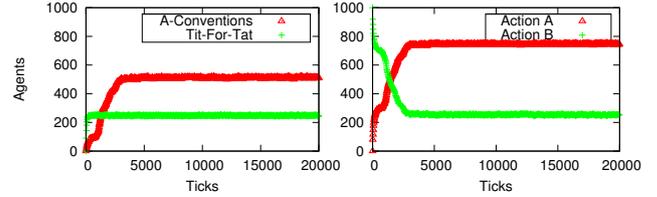
Our experiments aimed at showing that our model can help establish the best norm convention(s), maximize the social welfare, for a wide range of initial agent settings (norm configurations) and under the most common interaction topologies. Therefore, each experiment is composed of: i) an interaction topology model; ii) a payoff:  $\alpha \in \{1, 1.5, 2\}$ ; and iii) an initial norm distribution, consisting in initializing the norms of every agent using five distributions: a) **random** (norms are randomly set); b) **attractor-free** (norms set from the non-cooperative-only norms); c) **low sub-optimal** (norms of 25% of the agents set from the B-Conventions); d) **high sub-optimal** (75% of agents with norms from the B-Convention); and e) **fully sub-optimal** (norms of all agents were set from the B-Conventions). We run 50 simulations of each experiment. In a simulation agents interact and infect each other, as described above, during 20000 ticks. To measure if a convention is established, we counted the agents with the same norms per tick, and the agents doing A or B per tick. The counts of each simulation in the experiment were then aggregated using the inter-quartile mean.

**Pure coordination game** [ $\alpha = 1$ ]. The experiments show that the population converges an A-convention if initially more than 50% of the agents doing action A; otherwise, a B-convention settles down. Importantly, a MAS establishes the cooperative-only norms even though for this game other conventions can achieve the same result. Since the A and B-conventions are equally valuable, in this case the MAS establishes one of the best conventions regardless of the initial norm distribution and independently of the interaction topology.

**Different social efficiencies** [ $\alpha > 1$ ]. When using **random** initial distribution, a MAS readily establishes in an A-convention for  $\alpha > 1.0$  independently of the interaction topology. The same occurs for the **attractor-free** and the **low sub-optimal** initializations, even though in the former, at startup no agent knew the best norms.

Departing from a **high sub-optimal** distribution, a MAS establishes in a B-convention when  $\alpha = 1.5$  for all interaction topologies. However, by setting  $\alpha$  to 2.0, the small-world networks manage to establish an A-convention. Thus, agents will not consider a new convention unless its benefit is significant enough. As to the scale-free case, a greater benefit is needed.

The **fully sub-optimal** distribution represents the worst case scenario (figure 1). In this case, innovation becomes a key factor. When the innovation probability is low, ( $p_{mutation} = 0.003$ ), the MAS is unable to converge to the best convention, because innovating agents are not able to overcome the high peer pressure. Even more, in-



**Figure 1.** Results for scale-free with fully sub-optimal initialization. Left) agents per norm; Right) agents per action

fect scale-free networks are hard to overcome [3, 5]. Hence, we increased the mutation probability ( $p_{mutation} = 0.055$ ) so that scale-free ( $\alpha = 2.0$ ) and small-world ( $\alpha > 1$ ) converged to an A-Convention. This occurred because a small group of agents playing *tit-for-tat* kind of norms starts to appear. Agents with this strategy can coexist with B-Convention agents with a small or non-negative effect to their accumulated payoffs. Therefore, when agents with an A-convention norms appear, they have a higher chance of having neighbors that will cooperate with them. However, a high mutation presents the disadvantage that a small part of the population will be constantly trying to innovate (for our case around 20%).

We conclude that *highly-clustered* agent communities (e.g. small-world) are more open to positive infections, whereas the *low-clustered* ones (e.g. scale-free) are harder to infect if some infection has settled. This is similar to some results shown in [6]. However, our evolutionary model can overcome the difficulty of re-infecting *low-clustered* networks (by using a high innovation through mutation rate) whenever we are ready to pay the following cost: a small subgroup of agents unable to settle on a set of norms.

Finally, we claim that i) a convention is always reached, and ii) under certain conditions this convention is the best one for all topologies. Moreover, when these conditions are not met, e.g. a suboptimal convention is fully established, our model can still reach the best convention through innovation.

## ACKNOWLEDGEMENTS

The first author thanks the CONACyT scholarship. The work was funded by IEA (TIN2006-15662-C02-01), AT (CSD2007-0022), S-sourcing, and the Generalitat of Catalunya grant 2005-SGR-00093.

## REFERENCES

- [1] R Burt, 'Social contagion and innovation: Cohesion versus structural equivalence', *American J. of Sociology*, **92**, 1287–1335, (1987).
- [2] R Conte and Mario P, 'Intelligent social learning', *Artificial Society and Social Simulation*, **4**(1), 1–23, (2001).
- [3] Z Dezso and A.L. Barabási, 'Halting viruses in scale-free networks', *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, **65**(5), 055103, (2002).
- [4] D Moriarty, A Schultz, and J Grefenstette, 'Evolutionary algorithms for reinforcement learning', *Artificial Intelligence Research*, **11**(1-1), 241–276, (1999).
- [5] R Pastor-Satorras and A Vespignani, 'Epidemic dynamics and endemic states in complex networks', *Physical Review E*, **63**, 066117, (2001).
- [6] J.M. Pujol, J Delgado, R Sangüesa, and A Flache, 'The role of clustering on the emergence of efficient social conventions.', in *IJCAI 2005*, pp. 965–970, (2005).
- [7] Y Shoham and M Tennenholtz, 'On the emergence of social conventions: Modeling, analysis, and simulations', *Artificial Intelligence*, **94**(1-2), 139–166, (1997).
- [8] A Walker and M Wooldridge, 'Understanding the emergence of conventions in multi-agent systems', in *ICMAS 1995*, pp. 384–389, (1995).
- [9] D.J. Watts and S. H. Strogatz, 'Collective dynamics of 'small-world' networks', *Nature*, **393**(6684), 440–442, (June 1998).