

Emergence of Social Networks in Systems with Attraction

Elpida Tzafestas

Institute of Communication and Computer Systems
National Technical University of Athens
Zographou Campus, Athens 15773, GREECE
brensham@softlab.ece.ntua.gr

Abstract. In this paper we study the effect of a reactive attraction mechanism in a population of agents that can develop partner preferences. We use as base task for agent pair interaction the benchmark IPD game in its noisy version. Attraction makes an agent unconditionally cooperative toward an attractive opponent, hence introducing a parallel relation between agents, independent from the social interaction context itself, i.e. from the game. Partner selection can exploit such a given psychological condition and discover attracted partners, despite the fact that attraction itself is nowhere represented in the agent reasoning mechanism but is modeled as an uncontrollable process. The mechanics of partner selection are studied systematically in simulation.

Keywords: Cooperation, attraction, partner selection, social network.

PAPER SHORT VERSION !!

1 Introduction

Our motivation behind the introduction of an attraction mechanism is the general observation that in human societies, and especially in economic contexts, the agents' behavior can be heavily influenced by external psychological and social factors and also many times it can be driven to behaviors outside their normal scope. By "external" we mean a factor or process that is not influenced itself by the primary agent task and does not normally participate in it. We are using the benchmark iterated prisoner's dilemma (IPD) [1][2] in its noisy version [3][4] as a study vehicle with a stronger bias toward defection, where we feel it could make sense to introduce such an external attraction factor. More specifically, we believe that biological evolution or, equivalently, social experience would spontaneously exploit any external factor that would induce better agent scores. This is particularly true for noisy environments where agent scores may degrade abruptly, and especially when interactions are lengthier.

The attraction mechanism relies on our everyday experience that people tend to be good and cooperative with other people that attract them and tend to be "regular" with the rest. This translates in our model as:

*If (attracted by the opponent) then play ALLC (always cooperate),
Else play as usually (for example, TFT)*

We should note that noise is applied to the outcome of this behavior as well. We performed experiments with populations of agents playing a noisy IPD. The agents are interconnected via a “web of attraction” where each agent is connected to (attracted by) a number of others. The normal behavior of an agent is usually one of ALLC, ALLD (always defect), TFT and Adaptive TFT [5], but we have also experimented occasionally with STFT (Suspicious TFT) or other strategies. We experimented with both uniform or mixed populations, whose agents have the same or diverse normal strategies. The reason we use mostly ALLC, ALLD, TFT and Adaptive TFT is that we want to make sure we explore the limits of our attraction mechanism by studying its effect on the extreme behaviors (ALLC and ALLD that act without feedback) as well as on the most intelligent ones (TFT that retaliates immediately and Adaptive TFT that tries to make sense of a situation).

In previous work [5], we have classified usual IPD strategies in two categories: “retaliating” (or “rational”) and “irrational”. Retaliating strategies are those mostly TFT-derived strategies that basically seek cooperation in the long run, but may start by exploring the opponent’s reaction to a few initial D moves and will certainly retaliate the opponent’s defections in some possibly intricate way (apparently, this general behavioral organization is the best choice so as to achieve maximum scores in the long run). For example, the suspicious tit-for-tat (or STFT) strategy starts by defecting, and then plays usual tit-for-tat. On the contrary, irrational strategies are those that do not employ any feedback from the game and play blindly using some innate law (although in some cases this can work, in the general case such strategies lose in the long run). For example, periodic strategies repeat patterns of C’s and D’s, such as CDD, CCD, CDCD etc. In this sense, ALLC and ALLD are irrational strategies whereas TFT and Adaptive TFT are retaliating ones. As in real life, we would expect irrational strategies to profit more from our attraction mechanism or other similar mechanisms and rational ones to be less dependent on such add-ons. In sum, we expect cooperation to be able to emerge in social interactions even in the absence of rationality and good reason.

Before proceeding to describe the experimental setup and the results obtained, we should stress the fact that the attraction mechanism described is in our own terms irrational in that it does not depend on any real feedback of the agent. Our results then suggest that the coupling of reasoning mechanisms with reactive ones (such as attraction, be it physical, emotional, social or other) may be advantageous to social behavior and this is in line with current trends in cognitive and social science.

2 Basic interaction and partner selection

In previous work [6] we have studied how the above attraction mechanism affects tournaments of agents interconnected via a “web of attraction”. It is concluded there that the “social fitness”, in terms of total score, profits from the introduction of the attraction mechanism, and especially for the irrational agents that are more vulnerable

to order in the absence of noise. Because each agent seeks to maximize its personal score and because this score depends crucially on its attraction relations with the other agents, it is reasonable to try to select partners. However, selection will be based solely on obtained scores between pairs of agents and not on explicitly perceived relations outside the interaction, i.e. the noisy IPD game. So, attraction affects interactions but cannot be directly perceived by the agents; hence, it is modeled as an uncontrollable (unconscious, at present) process.

We characterize our experiments by the number of agents N participating in a population, the attraction factor M (number of agents that an agent is attracted to), the partner set size K (number of agents that an agent interacts with) and the normal strategy of the agents. In each round, an agent selects K partners to interact with and receives a total score. The length of each noisy IPD game has been set to 100 cycles and the degree of noise to 10%.

Partner selection is done based on a simple probabilistic preference scheme where each agent maintains a set of probabilities of interaction with each other agent of the society. All partners are equiprobable in the beginning of an experiment and the corresponding preferences develop according to the following reinforcement algorithm:

Let s be the score against opponent i and avg be the average score with all opponents in this round.

If $(s > avg)$ then increase preference for partner i ,

Else decrease preference for partner i .

Re-normalize all preferences so that they express probabilities

We give below the results for the example case of a population of TFT agents ($N=30$, $M=5$, $K=5$). Figure 1 left shows the evolution of the average score per agent, and figure 1 right shows the evolution of the actual attraction factor, i.e. the average of the proportion of attracted partners per agent.

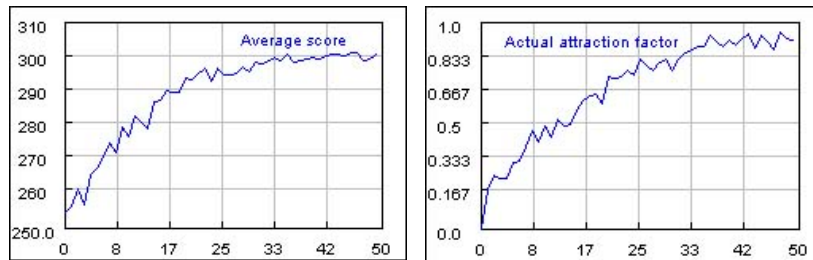


Fig. 1. (Left) Evolution of the average score per agent: it approaches 300, which is the theoretical best for a pair of TFT agents playing a 100-cycle IPD game. (Right) Evolution of the actual attraction factor, i.e. the average of the proportion of attracted partners per agent: it approaches 1, which is the theoretical maximum.

Qualitatively similar results are obtained with uniform populations of other behavior and with mixed populations. The precise quantitative gains in average score remain however highly variable across different cases and depend on the actual behavior mix and on the attraction web.

3 More social regulation

We have performed an additional perturbation study, equivalent to a regular invasion study performed in usual IPD games, where we reinitialize the attraction web after the social system has stabilized. For example, we run a system for 100 cycles until stabilization and then for another 100 rounds with a new attraction web. In the following figure we present comparative results for all three cases: the initial system (Case 1), the same system after attraction reinitialization (Case 2), and the second (reinitialized) system replayed with initial equiprobable preferences per agent (Case 3). It can be clearly seen that the (re)stabilization potential of such a social system depends on the initial conditions, that are the initial preferences of each agent. More specifically, the system when reinitialized after stabilization (Case 2) demonstrates a lower performance (inferior average score and attraction factor after stabilization), apparently due to the non-equiprobable initial partner preferences, unlike cases 1 and 3. This is both undesirable from a pure modeling perspective and biologically unrealistic: we would rather expect to see the system re-stabilize to the same final situation after any such perturbation.

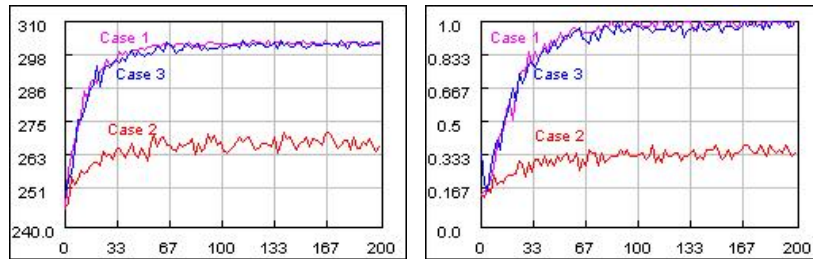


Fig. 2. (Left) Evolution of the average score per agent in the three cases (see text). (Right) Evolution of the actual attraction factor, i.e. the average of the proportion of attracted partners per agent in the three cases (see text). Apparently, case 2 cannot restabilize well enough, that is to the same near-optimal situations as the other two cases.

To face this requirement, we introduce an additional “exploration factor” used for partner selection. This expresses a probability with which an agent will choose a random partner for a particular interaction and not one based on its preference set. Because a system under stable attraction conditions does not need to re-stabilize, it is reasonable to assume that the exploration factor in this case should be 0. To account for perturbations, however, we would need a nonzero value for this factor, say 0.08 (8%). To integrate all the above requirements in one simple rule we use the following meta-regulation rule:

*Let longAvg be the average of avg over a specified number of past rounds.
If (avg > longAvg) then exploration factor = max,
Else exploration factor = 0.*

With this additional rule the agents manage to regulate how much they actively search for new partners and for potential better opportunities. Figure 3 gives the same results as figure 2, but this time with the additional meta-regulation rule. Note that the transitive results can be more variable and that the convergence speed is lower than

with the previous rule: both these features are the price to pay for the occasional use of an exploration factor in partner selection. This meta rule uses negative feedback in the sense that, if the actual behavior of the agent (expressed in the obtained score) moves in one direction (increases or decreases), then the regulated parameter (the exploration factor) will move in the other direction (it will increase or decrease respectively, which will probably lead to the score decreasing or increasing respectively). We have used such rules in previous works of ours (see for example [7]) because they allow systems to remain at the edge of behavioral change and hence to show enhanced stability in front of perturbations of various types.

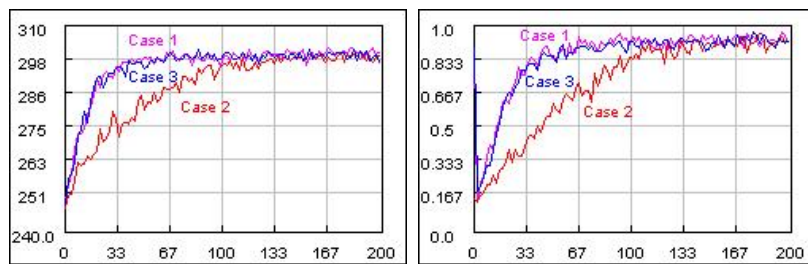


Fig. 3. (Left) Evolution of the average score per agent in the three cases (see text). (Right) Evolution of the actual attraction factor, i.e. the average of the proportion of attracted partners per agent in the three cases (see text). Apparently, all cases lead to the same final stable situation.

We should note that many variants of the meta-regulation rule can be devised that will lead to the same qualitative results. Because this rule is so simple, the non-uniqueness of its exact details is an indication that an evolutionary process can actually discover it by trial-and-error. The same setup and rules can also be used with variations of the attraction mechanism, for example with a mechanism defining different degrees of cooperation for each of the four cases: no attraction, self attracted, opponent attracted, mutually attracted.

4 Conclusion

We have studied the impact of an attraction mechanism in a social system that can develop partner preferences. It is shown how an agent in such a system can “discover” attracted agents and how this can drive the partner selection process. We have also observed that partner preferences can of course develop in the absence of attraction and even of noise. However, in all those cases, the average scores obtained do not improve, hence the social network built is practically useless, and the need for more elaborate meta-regulation vanishes as does the need for partner selection itself. The lesson from our study is that if an uncontrollable “reactive” feature such as attraction exists, that is coupled with the agent’s regular behavior, then this can trigger the development and the evolution of intricate cognitive functions, such as meta-regulation, that seek to exploit the reactive feature’s specificities (or bypass them if

they are detrimental to the agent). The impact of this general approach can be significant in the modeling of early cognitive development, where initial impulsive behaviors are known to develop to higher-level reasoning abilities, especially in the social domain.

References

1. Axelrod, R., and Hamilton, W.D.: The evolution of cooperation, *Science* 211 (1981) 1390-96
2. Axelrod, R.: The evolution of cooperation. Basic Books (1984)
3. Kraines, D., and Kraines, V: Evolution of learning among Pavlov strategies in a competitive environment with noise, *Journal of Conflict Resolution*, 39(3):439-466 (1995)
4. Molander, P.: The optimal level of generosity in a selfish, uncertain environment, *Journal of Conflict Resolution*, 31(4):692-724 (1987)
5. Tzafestas, E.: Toward adaptive cooperative behavior, *Proceedings of the Simulation of Adaptive Behavior Conference*, Paris, September (2000)
6. Tzafestas, E.: Attraction and cooperation in noisy environments, in press.
7. Tzafestas, E.: Regulation problems in explorer agents, *Cybernetics and Systems*, 36:339-377 (2005).