

Selection for 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 evolutionary potential of a reactive attraction mechanism in a population of agents that interact via cooperation games, as exemplified by the noisy iterated prisoner’s dilemma. 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. An additional partner selection mechanism 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. We show how evolution can select attraction at various levels of increasing complexity and how each level triggers the selection of the next one, provided the corresponding mechanisms emerge behaviorally. We also discuss the implications of these experiments for the study of the evolution of complex cognitive capacities and functionalities.

Keywords: Selection, cooperation, attraction, partner selection, social network.

1 Introduction

Our motivation behind the introduction of an attraction mechanism is the general observation that in biological societies, and especially in human ones, 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.

In section 2, we summarize previous results [6] that 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.

In section 3, we study the impact of the 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. 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. The same applies to the bottom-up study of biological organizations, for example to the study of colonization in microbial populations, where specific sets of agents group together to form coherent wholes that interact only a little with external agents.

The focus of the paper is to show, in section 4, how the various attraction levels of increasing complexity and sophistication can be selected one after the other and how each level triggers the selection of the next one, provided it emerges behaviorally. Our results using artificial selection suggest that evolution in this case can proceed by “discovering” behaviors that build on previous behaviors by directly observing the results of the latter and adding regulatory and monitoring components. Finally, in section 5, we discuss how such components could emerge within a biological organism and what it would take from an environmental perspective in order to stabilize these behaviors.

2 Attraction in society

We adopt the noisy version of IPD in which there is a nonzero probability that an agent’s action will be switched to the opposite, i.e. from *COOPERATE* to *DEFECT* or vice versa. It has been shown that retaliating strategies such as TFT can score quite badly in the presence of noise, despite their superiority in the non-noisy domain [3][4]. This happens because even accidental defections may lead to a persistent series of mutual defections by both players, thus breaking cooperation. The usual approach is to introduce some degree of explicit generosity to account for opponent’s misbehaviors or to attempt opponent modeling.

We use instead an attraction mechanism which 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 tournament 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.

Our results sum up as follows [6]:

- Counter-intuitively, a random pairing scheme, where each agent is attracted by M others randomly and not necessarily reciprocally, is clearly superior to an exact pairing scheme where each agent is mutually attracted by exactly M other agents. This is because random pairing induces persistent cooperative behavior against non-reciprocally attracted agents, thus inviting them to cooperation by retaliating. As a consequence, score improvement in the case of TFT is much more pronounced than in the case of Adaptive TFT which is by default more robust to noise.
- Attraction makes sense in societies that are not uniform ALLC. Especially the ALLD behavior profits most from the introduction of attraction because it is given the chance to exploit others. On average an ALLD strategy exploits and is being exploited in a web of attraction so that its average score rises beyond its theoretical equilibrium. This is why the scores with noise improve for the ALLD, unlike what happens with retaliating agents.
- The bigger the attraction factor, the bigger the score improvement for any strategy. Again, the improvement is inversely proportional to the degree of “rationality” of the strategy: an ALLD gains more than a TFT that in turn gains more than Adaptive TFT.
- Finally, for mixed populations, the impact of the attraction web can be significant on the resulting score ranking, so that an agent can drop from top to bottom or climb from bottom to top on reinitialization of the attraction relations. Note also that, in mixed populations with attraction the ranking of an agent is not a function of its degree of “rationality”: the cleverest agents do not score best. Rather, an “irrational” agent may be the most performant agent in an appropriate context, whereas an otherwise intelligent agent can sink to the end of the score queue. So, attraction is a mechanism that perturbs the regular social/economic relations and calls for elaborate partner selection to both exploit its opportunities and smoothen its drawbacks.
- Attraction can also yield impressive results in spatial contexts [7]. Agents on a grid can stabilize to a robust configuration of arbitrary composition, and especially consisting of high proportions of “irrational” agents. Composition after various perturbation events (such agent injection) can change arbitrarily, and especially “irrational” agents may proliferate and even take over the population, unlike what happens without attraction.

3 Partner selection

It was concluded before that the “social fitness”, in terms of total score, profits from the introduction of the attraction mechanism, and especially for the non-retaliating

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 stick with the right partners. In our previous work [6], we performed experiments with ad hoc groupings of agents that are attracted only with one another, for example in a “two-sexes-like” manner, where each agent of one group (females) can only be attracted by agents of the other group (males) and vice versa. Our results with groups of equal cardinality within a population implied that the bigger the pool an agent can find attracted partners in, the higher the obtained scores.

The next logical step is to discover that it pays 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 performed experiments [8] characterized 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 have experimented with a population of TFT agents ($N=30$, $M=5$, $K=5$) and we have verified that the average population score converges to the theoretical maximum (3 per round, which is the reward value in an IPD game where both agents play cooperatively). The actual attraction factor, i.e. the average of the proportion of attracted partners per agent, also converges to the theoretical maximum of 1 (100%).

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.

We performed an additional perturbation study [8], equivalent to a regular invasion study performed in usual IPD games, where we reinitialized the attraction web after the social system had stabilized. For example, we run a system for 100 cycles until stabilization and then for another 100 rounds with a new attraction web. It was found 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 demonstrates a lower performance (inferior

average score and attraction factor after stabilization), apparently due to the non-equiprobable initial partner preferences, unlike what happens when the system runs from the beginning with this same attraction web. 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.

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 and they thus demonstrate the same qualitative long-term performance despite any re-initialization. The transitive results can however be more variable and the convergence speed can be 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 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 Rounds of selection

The previous mechanisms can be put in order of increasing complexity as follows:

- Level 0: Retaliating behaviors (TFT, Adaptive TFT, etc.)
- Level 1: Retaliating behaviors with attraction
- Level 2: Partner selection based on attraction
- Level 3: Meta-regulation of attraction-driven partner selection

We have performed artificial selection experiments from level to level by creating populations of agents containing both agents of one level and of the next one. As

before, 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, from level 2 onwards) and the normal strategy of the agents. The length of each noisy IPD game has been set to 100 cycles and the degree of noise to 10%. Selection proceeds as follows: for every generation we run a tournament between all agents and we rank the results. The behavioral “profiles” of half of the agents, the lowest ranking ones, are replaced by the behavioral profiles of the highest scoring agents, to produce the next generation of agents and evolution continues.

The results of an example experiment from level 0 to level 1 are given in figure 1, where $N=30$ (out of which 15 play TFT and 15 play Adaptive TFT) and $M=15$. The behavioral profile of the agents in this case consists of 1 bit, the attraction bit: when it is on, the agent takes its attraction relations into account as described in section 1, otherwise it uses its regular strategy always. We observe in the experiment that the average percentage of an active attraction bit rise consistently across generations to finally stabilise at a high value not far from the theoretical maximum. The same thing applies to the average population score as well.




Fig. 1. (Left) Evolution of the average score per agent in the artificial selection experiment relating levels 0 and 1. (Right) Evolution of the percentage of an active attraction bit in the population.




Fig. 2. Evolution of the percentage of an active attraction bit in the artificial selection experiment relating levels 0 and 1 (diverse population with high M , see text).

The conditions for evolution of attraction (level 1) are a high proportion of retaliatory agents in the population and a big attraction factor M , i.e. a critical mass of retaliatory agents densely connected in a web of attraction should already be present

for attraction to spread. This is not unrealistic, because such a critical mass can be imagined in many natural systems either because most agents of a society are descendants of the same parent (such as multiple bacterial “clones”) or because of imitation. As a demonstration of the condition of evolution, figure 2 gives the results of an example experiment from level 0 to level 1 with $N=30$ (5 TFT, 5 Adaptive TFT, 5 ALLC, 5 STFT, 5 ALLD and 5 CDD) and $M=15$. The population diversity combined with a high attraction factor M leads to rapid attraction extinction.

The results of the experiment from level 1 to level 2 are given in figure 3, where $N=30$ (9 TFT, 9 Adaptive TFT, 3 ALLC, 3 STFT, 3 ALLD and 3 CDD), $M=15$ and $K=5$. The behavioral profile of the agents in this case consists of another bit, the partner selection bit: when it is on, the agent selects partners as described in section 3, otherwise it interacts with randomly selected partners. As before, we observe that the average percentage of an active partner selection bit rise consistently across generations to finally stabilize at the theoretical maximum (100%). The same thing applies to the average population score as well that stabilizes at a value above the theoretical maximum of 3 per agent (this is due to the presence of the non-retaliating agents ALLD and CDD).




Fig. 3. (Left) Evolution of the average score per agent in the artificial selection experiment relating levels 1 and 2. (Right) Evolution of the percentage of an active partner selection bit in the population.

The conditions for evolution of partner selection (level 2) are a high behavioral diversity in the population and a big attraction factor M , bigger than the interaction factor K . Again this is not unrealistic, because diversity is given in a natural system and interactions usually have a cost, thus only a few can happen, whereas attraction comes for free. However, as before, for systems that fail to meet these conditions, partner selection may not evolve.

The results of the experiment from level 2 to level 3 are given in figure 4, where $N=30$ (15 TFT and 15 Adaptive TFT), $M=15$ and $K=5$. The behavioral profile of the agents in this case consists of a third bit, the meta-regulation bit: when it is on, the agent meta-regulates its partner selection behavior as described in section 3, otherwise it uses regular partner selection. This experiment is trickier than the previous ones, because in order to test the need for meta-regulation we should artificially design the corresponding environment with an appropriate perturbation factor. To achieve this effect, we run all simulations for 50 cycles before reinitializing the attraction web and running the system for another 100 cycles at the end of which we measure and record average population scores. This method introduces the artificial bias toward selection

for perturbation-resistant behaviors, like the meta-regulated partner selection behavior. We observe that the meta-regulatory behavior rises steadily in the evolving population, albeit slowly.




Fig. 4. Evolution of the percentage of an active meta-regulation bit in the population in the artificial selection experiment relating levels 1 and 2. The actual curve (curve with local peaks) is visualized together with its moving average with a window of 10 (steadily increasing curve).

The condition for evolution of regulatory partner selection behavior (level 3) is the existence of persistent perturbations without behavioral changes and adaptation, i.e. a high-frequency perturbation scheme where things change very rapidly compared to the agent life span. Once more, we know from experience that this frequently happens in social systems, especially in higher vertebrates, and it is not surprising that the more difficult the social environment, the more intelligent one has to be. As always, easier and more stable environments will not allow higher cognitive functions to evolve.

5 Discussion

The above artificial evolution experiments can serve as a vehicle for understanding the emergence and evolution of complex cognitive functionalities out of more primitive ones. More specifically, we observe that each subsequent level of behavior builds on the possibilities of the previous level with the aid of a monitoring and regulatory mechanism:

Level 1: With retaliating behavior present, react differentially toward attractive agents.

Level 2: With attraction behavior present, identify friend agents (attractive or attracted) and keep them while discarding the rest of the agents.

Level 3: With partner selection behavior present and an additional exploration factor, regulate the exploration factor so as to quickly identify friend agents under all social conditions.

These may be presented schematically in the table found in the next page.

Each subsequent level can therefore emerge behaviorally by superposition of novel perceptual capacities that are coupled with novel regulatory actions. Both perception and action innovation may be the result of a selective behavioral generation process within the agent which produces a wealth of capacities and actions that are tried out in

the real environment. The perception-action couplings that happen to be beneficial for the agent can therefore be selected and stabilize in the social environment, while irrelevant capacities and/or couplings will die out. Thus, each subsequent evolution level can evolve by using the previous level as a building block and further “reasoning” on its results. In our experiments, we have accordingly shown how an initially reactive and uncontrollable mechanism, such as attraction, may be exploited and integrated by the agents together with accompanying reasoning capacities so as to allow much more complex behaviors of much greater value.

	Substrate	New perception	New action
Level 0		Retaliatory behavior	
Level 1	Level 0	Identify attractive agents	Always cooperate with attractive agents
Level 2	Level 1	Associate individual agents with scores	Select partners that allow high-scoring, discard the rest
Level 3	Level 2 (+ exploration factor)	Compare short and long term past scores	Regulate exploration factor

But even if the behavioral generation process is in place, this is not enough to allow subsequent behavioral levels to be favoured by the evolutionary process and become therefore selected. In our experiments, we have identified at least two factors that act on the possibility of a population to select the higher level behavior: the agent behavioral mix and its evolution in time. We have departed from the hypothesis that the initial population at some point in evolution will be mostly comprised of retaliatory agents, such as TFT and Adaptive TFT. This is only partly true, because most such populations can be vulnerable to hard defective agents and may be furthermore invaded by a mix of other non-retaliatory agents. As a demonstration of this, experiments from level 0 to level 1 may often fail in a population mix with a low attraction factor, in the sense that the active attraction bit may become extinct or it may stabilize for a long period to an intermediate value between 0 and 100%. However, if behavioral generation of the attraction mechanism (level 1) coincides with a long evolutionary period when the population consists mostly of retaliatory agents, there is high chance that the active attraction bit will rise quickly and overtake the population. Similar things apply to the experiments from level 1 to level 2 that may fail in the opposite case of retaliatory rather than mixed populations with low M or high K. Because retaliatory behaviors, with or without attraction, are still vulnerable to invasion by other behaviors, we expect level 2 to subsequently emerge and become selected at long evolutionary periods when populations are highly diverse. Finally, level 3 may be evolutionarily selected in periods when persistent perturbations exist and reliance on partner selection is crucial. As a result of all the above, a long-term artificial evolution experiment would not necessarily produce all three levels of behavior, and different subpopulations would most possibly evolve differently. In all cases, it appears that some threshold value has to be exceeded for evolution of consecutive levels to occur, or, equivalently, agent populations of a novel behavior have to acquire critical mass to allow behavioral explosion.

It is finally not unreasonable to assume that level N can start emerging even when level N-1 is not fully (100%) in place. Thus, the evolutionary history of the system is not predetermined to follow one single linear path, but rather many different branches are expected to appear depending on the internal structure and interactions of the system, as well as on external factors, mostly ecological ones. For example, local effects of imitation can create a critical mass of attracted agents that will soon start selecting partners, especially if the population includes also other agents that do not even know attraction. In this sense, we expect to see changes that will function catalytically to further diversification of the population in terms of behavioral levels of attraction. It is finally reasonable to expect that a behavioral trait may be used in a different context and role with the initial apparatus gradually disappearing, for example if the partner selection mechanism acts on new criteria influencing the social fitness, then the attraction mechanism that probed its emergence may itself disappear.

As a conclusion, we feel it will be important in the future to study the selection process itself: how basic agent strategy (TFT, ALLC etc.) co-evolves with behaviors of a different order (such as attraction) and how this combination may trigger the evolution of higher levels of functionality. To this end, we purport to integrate in the immediate future a mechanism for social imitation as well as to replace the artificial selection mechanism described in section 4 by a social selection mechanism, closer to intuitive ecological reality. We expect the first mechanism to speed up evolution, when coupled with the behavioral generation process, and the second to channel evolution steadily toward pertinent structures and away from the oscillatory nature of the artificial mechanism used so far. Also we plan to study the effect of social generations overlapping to the evolution of behavior of higher complexity.

We stress finally our belief that such an approach combining “low-level” reactive behaviors with “higher-level” reasoning behaviors is necessary at all levels of biological organization (microbial, multi-cellular, ecological etc.), because one of the main properties of all living organisms is their ability to incorporate information and experience about past behavior and corresponding outcome and to assess these in different, dynamic and noisy environments.

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, *Proceedings ESHIA 2007*, Fairfax, VA.
7. Tzafestas, E.: Attraction and cooperation in space, *Proceedings IEEE Conference on Evolutionary Computation 2007*, Singapore.
8. Tzafestas, E.: Emergence of social networks in systems with attraction, in preparation.