Promoting Cooperation through External Interference

The Anh Han[1], Long Tran-Thanh[2], Simon Lynch[1], Theodor Cimpeanu[1], Francisco C. Santos[3]

[1]School of Computing, Media and the Arts, Teesside University
[2]School of Electronics and Computer Science, University of Southampton
[3]INESC-ID and Instituto Superior Tecnico, Universidade de Lisboa
Email: t.han@tees.ac.uk

The problem of promoting the evolution of cooperative behaviour within populations of self-regarding individuals has been intensively investigated across diverse fields of behavioural, social and computational sciences (Nowak, 2006; Perc et al., 2017). In most studies, cooperation is assumed to emerge from the combined actions of participating individuals within the populations, without taking into account the possibility of external interference and how it can be performed in a cost-efficient way. However, in many scenarios, cooperative behaviours are advocated and promoted by an exogenous decision maker, who is not part of the system (e.g. the United Nation interferes in political systems for conflict resolution or the World Wildlife Fund organisation interferes in ecosystems to maintain biodiversity). Thus, a new set of heuristics capable of engineering a desired collective behaviour in a self-organised multiagent system is required. Here we summarize our recent works to bridge this gap, in which we employ theoretical analysis and computer simulations based on evolutionary game theory (Nowak, 2006), to study cost-efficient interference strategies for enhancing cooperation in the context of cooperation dilemma games, for both well-mixed (Han and Tran-Thanh, 2018) and square-lattice structured populations (Han et al., 2018).

We consider finite populations (of size N) of individuals who interact with each other through the one-shot Prisoner's Dilemma game (PD) (Nowak, 2006), where in each interaction two players simultaneously choose either to cooperate (C) or defect (D). Mutual cooperation (mutual defection) yields the reward R (penalty P) and unilateral cooperation gives the cooperator the sucker's payoff S while the defector the temptation T. A PD is characterized by the ordering $T > R > P > S$. In a well-mixed population, each player interacts with all others in the population while in a square lattice the player interacts with its four immediate neighbors. A player's fitness is its averaged payoff over all its interactions, which is then used for strategy update. Namely, a player A with fitness f_A chooses to copy the strategy of a randomly selected player in the population (well-mixed) or randomly selected neighbor (structured) with a probability

given by the Fermi function, $(1 + e^{\beta(f_A - f_B)})^{-1}$, where β represents the intensity of selection (Traulsen et al., 2006). When $\beta = 0$ corresponds to neutral drift while $\beta \to \infty$ leads to increasingly deterministic selection. Weak or even close to neutral selections (small β) are abundant in nature, while the strong selection regime has been reported as predominant in social settings. As an alternative to this stochastic update rule, one can also consider a *deterministic update* in which agents copy, if advantageous, the most successful player in their neighbourhood.

An interference strategy or scheme can be generally defined as a sequence of decisions about which C players in the population to invest in (i.e. reward the player an amount, denoted by θ), in order to achieve the highest level of cooperation while minimising the total cost of investment. These decisions can be made by considering different aspects of the population such as its global statistics and/or its structural properties. In the context of a well-mixed population, an interference scheme solely depends on its composition (i.e. how many C and D players there are at the time of decision making). In this case, we have derived analytical conditions for which a general interference scheme can guarantee a given level of cooperation while at the same time minimising the total cost of investment (for rewarding cooperative behaviours), and show that the results are highly sensitive to the intensity of selection by interference. Moreover, we have studied a specific class of interference strategies that make investments whenever the number of C players reaches a certain threshold, denoted by t ($\forall t \in \{1, \ldots, N-1\}$), showing that there is a wide range of t that it outperforms standard institutional incentive strategies—which unconditionally interfere into the system regardless of its composition, corresponding to $t = N - 1$ (Chen et al., 2015). Figure 1a shows the optimal threshold of t for varying the intensity of selection, β, where the minimal expected cost of interference is obtained while guaranteeing at least ω-frequency ($0 \leq \omega \leq 1$) of cooperation. We can observe that, when β is sufficiently small (weak selection), an intermediate value of t would lead to most cost-efficient interference strategies, while for a sufficiently strong selection, it is best to

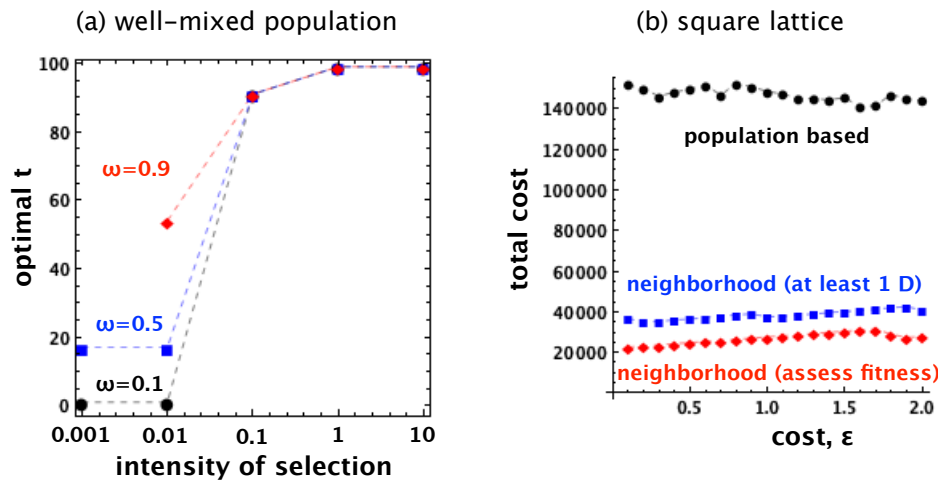**(a) well–mixed population**   **(b) square lattice**

Figure 1: **a)** In a well-mixed population (size $N = 100$), the optimal threshold of t where a minimal expected cost of interference is obtained while guaranteeing at least ω-frequency of cooperation, which increases with the intensity of selection (Other parameters: $T = 2$; $R = 1$; $P = 0$; $S = -1$; $\theta = 5$); **b)** In a square lattice population (size $N = 100^2$), neighborhood-based (local) interference strategies are more cost efficient in ensuring high cooperation ($\approx 100\%$) than population-based (population cooperation level) ones. Other parameters: $T = 1.8$; $R = 1$; $P = 0$; $S = 0$ (weak PD); deterministic update was used but results are robust for stochastic update (Han et al., 2018).

always interfere, meaning that standard institutional incentive strategies (i.e. $t = N - 1$) would be most effective.

With a structured population, individuals (even of the same strategy) might reside in different kinds of neighborhood (with different cooperativeness levels), and therein local information might be useful to enhance cost-efficiency and cooperation. To this end, we test several interference paradigms (Han et al., 2018) that make investment decisions based on a player's current cooperativeness level (the number of C players in the neighborhood), and compared their efficiency with the population-based strategies (as in the well-mixed case). Our systematic analysis reveals a simple strategy that invests when there is at least one D player in the neighborhood and does not invest otherwise, is highly cost-efficient in promoting cooperation (see Figure 1b). Furthermore, when additional information regarding the fitness levels (i.e. individual income information) of players in a neighbourhood is accessible, further improvement can be made by more accurately influencing D neighbours for behavioural change (to become cooperators).

Future works include analysis of other types of population structures such as the scale-free ones and their variants (Cimpeanu et al., 2019) and more complex interference strategies such as those vary the cost of investment over time or combine different forms of incentives (Chen et al., 2015; Han, 2016).

In short, we have studied how cooperation can be promoted in a cost-efficient way from an external decision maker's perspective. It provides new insights regarding heuristics capable of engineering a desired collective behaviour in a self-organised complex system, not only in social and biological contexts, but also Artificial Life scenarios such as swarm-based and multi-robots systems (Bonabeau et al., 1999; Han et al., 2012).

References

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford U. Press.

Chen, X., Sasaki, T., Brännström, Å., and Dieckmann, U. (2015). First carrot, then stick: how the adaptive hybridization of incentives promotes cooperation. *Journal of The Royal Society Interface*, 12(102):20140935.

Cimpeanu, T., Han, T. A., and Santos, F. C. (2019). Emergence of coordination with asymmetric benefits via prior commitment. In *ALIFE (in press)*.

Han, T. A. (2016). Emergence of social punishment and cooperation through prior commitments. In *AAAI*, pages 2494–2500.

Han, T. A., Lynch, S., Tran-Thanh, L., and Santos, F. C. (2018). Fostering cooperation in structured populations through local and global interference strategies. In *IJCAI-ECAI*, pages 289–295. AAAI Press.

Han, T. A., Pereira, L. M., and Santos, F. C. (2012). Corpus-based intention recognition in cooperation dilemmas. *Artificial Life journal*, 18(4):365–383.

Han, T. A. and Tran-Thanh, L. (2018). Cost-effective external interference for promoting the evolution of cooperation. *Scientific reports*, 8.

Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard U. Press, Cambridge.

Perc, M., Jordan, J. J., Rand, D. G., Wang, Z., Boccaletti, S., and Szolnoki, A. (2017). Statistical physics of human cooperation. *Physics Reports*, 687:1–51.

Traulsen, A., Nowak, M. A., and Pacheco, J. M. (2006). Stochastic dynamics of invasion and fixation. *Phys. Rev. E*, 74:11909.