

Fostering Cooperation in Structured Populations Through Local and Global Interference Strategies

The Anh Han¹, Simon Lynch¹, Long Tran-Thanh², Francisco C. Santos³

¹School of Computing, Media and the Arts, Teesside University

²School of Electronics and Computer Science, University of Southampton

³INESC-ID and Instituto Superior Tecnico, Universidade de Lisboa

{t.han,s.c.lynch}@tees.ac.uk, ltt08r@ecs.soton.ac.uk, franciscocsantos@tecnico.ulisboa.pt

Abstract

We study the situation of an exogenous decision-maker aiming to encourage a population of autonomous, self-regarding agents to follow a desired behaviour at a minimal cost. The primary goal is therefore to reach an efficient trade-off between pushing the agents to achieve the desired configuration while minimising the total investment. To this end, we test several interference paradigms resorting to simulations of agents facing a cooperative dilemma in a spatial arrangement. We systematically analyse and compare interference strategies rewarding local or global behavioural patterns. Our results show that taking into account the neighbourhood's local properties, such as its level of cooperativeness, can lead to a significant improvement regarding cost efficiency while guaranteeing high levels of cooperation. As such, we argue that local interference strategies are more efficient than global ones in fostering cooperation in a population of autonomous agents.

1 Introduction

The problem of understanding and predicting collective behaviours of agents in evolving multiagent systems is a well studied research topic in evolutionary game theory, as it can be found in a variety of real-world situations, ranging from ecosystems to human organisations and social networks [Santos *et al.*, 2006; Sigmund *et al.*, 2001; Tuyls and Parsons, 2007; Raghunandan and Subramanian, 2012; Han *et al.*, 2017a; Sahraei *et al.*, 2014]. In this context, cooperation is typically assumed to emerge from the combined actions of participating agents within the system. However, in many scenarios, such behaviours are advocated and promoted by an *exogenous agent*, which is not part of the system, calling for a new set of heuristics capable of *engineering* a desired collective behaviour in a self-organised multiagent system. For instance, consider a wildlife management organisation (e.g., the WWF) that aims to maintain a desired level of biodiversity in a particular region. To do so, the organisation, not being part of the region's eco-system, has to decide whether to modify the current population of some species,

and if so, then when, and in what degree to *interfere* in the eco-system (i.e., to modify the composition and the biodiversity of the population) [Levin, 2000]. Note that since more efficient population controlling typically implies more physical actions, which requires higher (monetary) expenses in both human resources and equipment, the organisation has to achieve a balance between efficient wildlife management and a low total investment cost. Moreover, due to the evolutionary dynamics of the eco-system (e.g., frequency and structure dependence [Santos *et al.*, 2006]), undesired behaviours can reoccur over time, for example when the interference was not sufficiently strong in the past. Given this, the decision-maker also has to take into account the fact that she will have to repeatedly interfere in the eco-system in order to sustain the level of biodiversity over time. That is, she has to find an efficient sequential interference strategy that leads to her desired goals, while also minimising the total cost of interference.

Although the (sequential) decision-making literature provides a number of techniques to tackle similar problems where the goal is to provide a sequence of decisions that lead to optimal behaviour of the system (e.g., the desired level of biodiversity) while minimising the total cost of making such decisions [Madani *et al.*, 2004; Guha and Munagala, 2007; Bachrach *et al.*, 2009; Tran-Thanh *et al.*, 2012; Ding *et al.*, 2013], these approaches typically ignore the fact that the agents, with whom the decision-maker has to interact, also have their own strategic behaviours that together drive the (evolutionary) dynamics of the system. Given this, we argue that such solutions will not be able to exploit the system characteristics, and thus, will fail in providing efficient performance in achieving the desired system status (e.g., the status quo between the fighting opponents, or the desired diversity of population). On the other hand, game theoretic literature, which deals with these strategic behaviours, typically focuses on the extremes of the problem. In particular, researchers either assume that the system is fully closed (i.e., there are no outsider decision-makers) or the decision-maker has full control over the behaviour of the agents. Typical models for the former are classical (both non-cooperative and coalitional) game theoretical models [Bachrach *et al.*, 2009; Aziz *et al.*, 2010; Aadithya *et al.*, 2011]. The latter includes models from mechanism design, where the decision-

maker is the system designer, and can define some set of norms and penalties such that the agents are incentivised not to deviate from the desired behaviour [Endriss *et al.*, 2011; Wooldridge, 2012; Levit *et al.*, 2013; Harrenstein *et al.*, 2014]. Given this, these works are not suitable to meet our current aims either.

This paper presents an alternative approach to address the problem of the external decision-maker by combining the decision-making process design with an evolutionary game theoretic perspective. While the former captures the strategic behaviour of the decision-maker, the latter can be used to formalise the evolutionary dynamics of the multiagent system and the strategic decisions within large populations exhibiting arbitrary population structures and social tensions. In particular, we consider a multiagent system where the agents, *distributed in a network*, interact with their neighbours via the one-shot Prisoner’s Dilemma (PD), where uncooperative behaviour is preferred over cooperation [Sigmund *et al.*, 2001; Santos *et al.*, 2006]. As an outsider decision-maker, we aim to promote cooperation by interfering in the system, rewarding particular agents in the population at specific moments;

The research question here is to identify when and how much to invest (i.e., pay the agents) at each time step, in order to achieve our desired ratio of cooperation within the system such that the total cost of interference is minimised. We investigate two general classes or approaches of interference strategies; the first is based on the current composition of the population while the second is based on local neighbourhood properties, in particular, for any agent being considered for investment we examine its current cooperativeness level (the number of its neighbours who will cooperate). These two classes represent a large fraction of the sequential interference strategies in real-world scenarios, requiring different levels of information in the decision making process. While the first requires as input the overall population cooperativeness, the second requires more detailed, local information regarding each neighbourhood to make a decision. They represent two different approaches of governing and incentivising institutions: a large global institution overseeing the whole population (such as the United Nations) as opposed to multiple local institutions (such as regional or group-wide ones) [Vasconcelos *et al.*, 2013]. We are particularly interested in whether taking into account neighbourhood information will lead to more efficient interference and how this may be achieved.

The remainder of the paper is structured as follows: the next section provides a brief overview of the related work, this is followed by a detailed description of our model, its results and a final discussion.

2 Related Work

In the literature of evolution of cooperation, different mechanisms of cooperation have been studied, including direct and indirect reciprocity [Nowak and Sigmund, 2005], kin and group selections [Traulsen and Nowak, 2006], network reciprocity [Santos *et al.*, 2006], punishment and rewarding [Sigmund *et al.*, 2001; Han, 2016], and pre-commitments [Han *et al.*, 2017b]. These mechanisms are incorporated as part of individual strategic behaviours, in order to study how they

evolve and whether their evolution promotes a better outcome for cooperative behaviour. These works, however, do not consider the influence of an exogenous decision-maker, but rather rely on the changes made by internal agents. In contrast, our interference strategies are external, i.e. they are not incorporated into the individual strategy. In addition, the aim of our strategies is to minimise the cost of interference while guaranteeing high levels of cooperation, contrary to past literature where the cost optimisation is often omitted. In this respect, our work is also different from the modelling works of institutional incentives to encourage cooperation through costly reward and punishment (see e.g., [Sigmund *et al.*, 2010; Vasconcelos *et al.*, 2013]) as well as through enforcing agreements [Han *et al.*, 2017b].

Closely related to the current work is the analysis in [Han *et al.*, 2015; 2014] where cost-efficient interference is studied in well-mixed populations (i.e. having a fully connected graph structure). Thus, the interference therein only depends on the composition of the population. In the more complex scenario of structured populations we consider in this paper, agents might have different types of neighbourhood. As shown below, taking into account neighbourhood properties leads to more cost-efficient interference strategies than that based solely on the population composition.

Works on cooperation in social networks also assume that changes are initiated from inside the system [Raghunandan and Subramanian, 2012; Sahraei *et al.*, 2014; Franks *et al.*, 2013; 2014]. Among them, more relevant to our paper is the recent work by Franks *et al.* [Franks *et al.*, 2013; 2014], which has explored the use of influencers on complex networks. However, these influencers are also part of the system and thus, similar to the cases mentioned above, this work does not consider external interference mechanisms. Given this, it does not address similar decision-making problems that we examine here. Apart from these examples, some other works apply schemes of mechanism design to control the behaviour of the system. Endriss *et al.* [Endriss *et al.*, 2011] investigated how to tax games to incentivise certain behaviours in system equilibrium points and Wooldridge *et al.* [Wooldridge, 2012] discussed the options to manipulate games in order to achieve desired behaviours. These works, however, assume that the decision-maker has full control on the agents within the systems, which allows her to both reward the good agents (i.e., those who follow the desired behaviours) and penalising the deviant ones. With our approach, the decision-maker has little or no control on the agents so they can rely only on rewarding schemes, nudging agents to adopt a given strategy. Moreover, mechanism design work typically does not focus on the cost of maintaining the desired behaviour, whereas cost optimisation is one of our objectives.

Last but not least, Bachrach *et al.* [Bachrach *et al.*, 2009] investigated how much investment into a cooperative game is needed in order to ensure that a certain coalition structure is stable. This work, however, only considers non-evolving systems with a single time step, and thus, interference can only be applied once. In contrast, in our system, due to its evolutionary dynamics, interference is repeatedly carried out over time (see the next section for more detail). Thus, the related work (above) cannot be applied directly to our approach.




Figure 1: Evolution of strategies over time (red for D, blue for C) and per-generation cost: (Two top rows) POP for different values of p_C (with $\theta = 4.5$); and (Two bottom rows) NEB for different values of n_C (with $\theta = 5.5$). We only show 50 generations for the sake of clear presentation. For POP: for small p_C , the population enters cyclic patterns while when it is sufficiently large (close to 100% in general), the population quickly converges to 100% of cooperation, avoiding cyclic behaviours. For NEB: when $n_C \leq 2$, the population converges to a stable level of cooperation that needs to be sustained by constantly interfering in every generation. When $n_C \geq 3$, the population quickly converges to 100% of cooperation, thus avoiding the cost of maintaining a stable state. Parameters: $b = 1.8$, $L = 100$.

3 Models and Methods

3.1 Prisoner's Dilemma on Square Lattice Networks

We consider a population of agents on a square lattice of size $Z = L \times L$ with periodic boundary conditions—a widely adopted population structure in population dynamics and evolutionary games (for a survey, see [Szabó and Fath, 2007]). We focus our analysis on the efficiency of various interference strategies in spatial settings, adopting an agent-based model directly comparable with the setup of recent lab experiments on cooperation [Rand *et al.*, 2014].

Initially each agent is designated either as a cooperator (C) or defector (D) with equal probability. Agents' interaction is modelled using the one-shot Prisoner's Dilemma game, where mutual cooperation (mutual defection) yields the reward R (penalty P) and unilateral cooperation gives the cooperator the sucker's payoff S and the defector the temptation T . As a popular interaction model of structured populations [Szabó and Fath, 2007], we adopt the following scaled payoff matrix of the PD: $T = b$, $R = 1$, $P = S = 0$ ¹ (with $1 < b \leq 2$).

At each time step or generation, each agent plays the PD with its (four) immediate neighbours. The score for each agent is the sum of the payoffs in these encounters. At the start of the next generation, each agent's strategy is changed to that of its highest scored neighbour [Nowak and May, 1992; Szabó and Fath, 2007]. Our analysis will be primarily based on this deterministic, standard evolutionary pro-

cess in order to focus on understanding the cost-efficiency of different interference strategies. However, we confirm that all our conclusions remain valid for a stochastic update rule (see Section 4.4 for more details). Namely, instead of copying the highest scored neighbour, at the end of each generation an agent A with score f_A chooses to copy the strategy of a randomly selected neighbour agent B with score f_B with a probability given by the Fermi rule [Traulsen *et al.*, 2006]: $(1 + e^{(f_A - f_B)/K})^{-1}$, where K denotes the amplitude of noise in the imitation process [Szabó and Fath, 2007]. Varying K allows capturing a wide range of update rules and levels of stochasticity, including those used by humans, as measured in lab experiments [Zisis *et al.*, 2015; Rand *et al.*, 2013].

We simulate this evolutionary process until a stationary state or a cyclic pattern is reached. Similarly to [Nowak and May, 1992], all the simulations in this work (described in next section) converge quickly to such a state. For the sake of a clear and fair comparison, all simulations are run for 200 generations. Moreover, for each simulation, the results are averaged from additional 50 generations after that. Furthermore, to improve accuracy, for each set of parameter values, the final results are obtained from averaging 50 independent realisations.

Note that we do not consider mutations or explorations in this work. Thus, whenever the population reaches a homogeneous state (i.e. when the population consists of 100% of agents adopting the same strategy), it will remain in that state regardless of interference. Hence, whenever detecting such a state, no further interference will be made. It is noteworthy that the results remain robust assuming a sufficiently small

¹The results obtained in this paper remain robust when P is a small positive value, resulting in a strict PD [Nowak and May, 1992].

mutation rate, allowing the population to reach a stable state before any new mutants arise.

3.2 Cost-Efficient Interference in Networks

As already stated, we aim to study how one can efficiently interfere in a structured population to achieve high levels of cooperation while minimising the cost of interference. An investment in a cooperator consists of a cost $\theta > 0$ (to the external decision-maker/investor). In particular, we investigate whether global interference strategies (where investments are triggered based on network level information) or their local counterparts (where investments are based on local neighbourhood information) lead to successful behaviour with better cost efficiency. To do so, we consider two main classes of interference strategies based on the global composition of the population and the neighbourhood information.

1. Population composition based (POP): In this class of strategies the decision to interfere (i.e. to invest on all cooperators in the population) is based on the current composition of the population (we denote x_C the number of cooperators currently in the population). Namely, they invest when the number of cooperators in the population is below a certain threshold, p_C (i.e. $x_C \leq p_C$), for $1 \leq p_C \leq Z$. They do not invest otherwise (i.e. when $x_C > p_C$). The value p_C describes how widespread defection strategy should be to trigger the support of cooperators' survival against defectors.

2. Neighbourhood based (NEB): In this class of strategies, the decision to invest in a given cooperator is based on the cooperativeness level in that cooperator's neighbourhood. While POP can be seen as a *global* interference strategy, NEB adopts a *local* interference paradigm. Namely, the decision-maker invests in a cooperator when the number of its cooperative neighbours is below a certain threshold, n_C , for $0 \leq n_C \leq 4$; otherwise, no investment is made.

In addition to these primary strategy classes, we consider two additional advanced local strategies, assuming that the external decision-maker can assess individual scores of players in a neighbourhood (e.g. an institution might have access to individuals' income records).

(NEB-i) Maintaining-C: for every cooperator, if the best scoring neighbour is a D which has a payoff larger than that of C, then invest an amount equal to the difference of the payoff of best scoring neighbour to that of C, plus a small amount $\epsilon > 0$.

(NEB-ii) Maintaining-C plus influencing D-neighbours. For every C we proceed as in (NEB-i). We then perform an additional second step in which we lure D neighbours of C to copy its strategy. To achieve this, if the highest scoring neighbour of D is a defector, then invest an additional amount to ensure that C has a higher payoff than that of the most fit D, by a small amount $\epsilon > 0$.

Note that these interference strategies are increasingly more subtle, requiring more information and so more detailed observations of the population. In particular, POP requires knowledge of overall cooperativeness of the population, while NEB requires information about local cooperativeness in each neighbourhood. NEB-i and NEB-ii need to access the fitness levels of all players in each neighbourhood.




Figure 2: Population-based (POP: top row) vs neighbourhood-based (NEB: bottom row) interference, for varying per-individual cost of investment θ as well as the threshold of population cooperation (p_C) or the threshold number of cooperators in a neighbourhood (n_C), respectively. The left column reports the frequency of cooperation while the right one reports the total cost required (normalised by dividing by 10^7). Parameters: $b = 1.8$, $L = 100$; $k = 4$ (node degree).

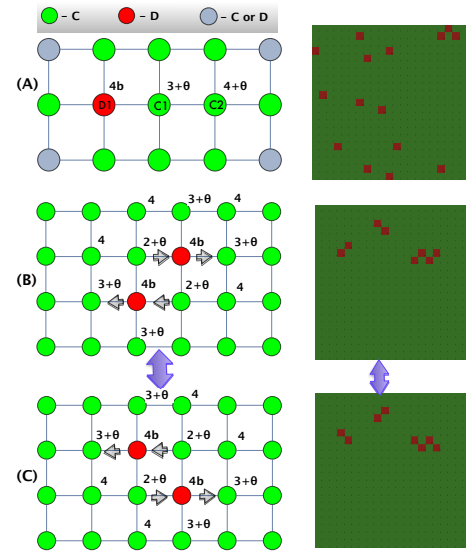


Figure 3: Typical cyclic patterns and stable configurations observed in simulations. (A) POP strategies. D1 copies C1 if $\theta > 4b - 3$ otherwise if $4b - 4 < \theta \leq 4b - 3$ neither D1 nor C1 copies each other, leading to a stable configuration, cf. the snap shot of a stable population ($b = 1.8$, $\theta = 4.1$). When $4b - 4 > \theta$ then C1 copies D1. (B-C) NEB strategies where $n_C = 3$. Cyclic pattern between (B) and (C) configurations occurs when $4b - 3 < \theta < 4b - 2$, cf. the snap shot of the cyclic pattern ($b = 1.8$, $\theta = 5.1$). When $\theta \geq 4b - 2$ defectors copy C, escaping the cycle.




Figure 4: Compare the total costs for different NEB strategies: NEB-4 with $n_C = 4$, NEB-3 with $n_C = 3$; and also when fitness observation is possible: NEB-i and NEB-ii. We vary ϵ , which represents the extra amount to the threshold value of θ which guarantees population convergence to 100% of cooperation for each strategy. In general, NEB-3 is significantly more cost-efficient than that with $n_C = 4$. NEB-ii is better than NEB-i. Parameters: $b = 1.8$, $L = 100$.

4 Results

4.1 Cost-Efficient Interference: POP vs NEB

We compare the population-based and neighbourhood based interference strategies, i.e. POP vs NEB, namely their abilities to promote cooperation while maintaining a minimal total cost. In Figure 1, we investigate the evolution over time of the frequency of cooperation and the corresponding cost required as per-generation for POP and NEB strategies, and for different values of p_C and n_C , respectively. We observe that for POP, the population converges to cyclic patterns for a small p_C , while when p_C is sufficiently large (close to 100% in general), the population quickly converges to 100% of cooperation, avoiding cyclic behaviours. Cyclic patterns occur since interference can lead to an increase in the frequency of Cs; when this rate becomes larger than p_C , it causes the interference by POP to stop, which in turn leads to the decrease of the frequency of cooperative actions and then causes interference to restart.

On the other hand, for NEB, when $n_C \leq 2$, the population converges to a stable level of cooperation nurtured by constant interference. When $n_C \geq 3$, the population quickly converges to 100% of cooperation, thus avoiding the cost of maintaining a stable state. Most interestingly, comparing $n_C = 4$ and $n_C = 3$, we observe that the per-generation cost increases for the former until reaching 100% of cooperation, while it decreases for the latter. The reason is that there are increasingly more cooperators requiring costly investment in the former case while such a number decreases in the latter. As such, NEB with $n_C = 3$ has a significantly lower total cost than NEB with $n_C = 4$ while guaranteeing similar high levels of cooperation (see Figure 4 for other parameter values of θ).

Next, for each interference approach (POP and NEB) and threshold (p_C and n_C), we relate the stationary cooperation level with the cost spent per-individual (θ) (see Figure 2). In

general, the C frequency obtained by POP mostly depends on p_C – it requires a high p_C ($> 85\%$ of the whole population) to reach a high level of cooperation, turning POP a costly strategy. For NEB, similar high levels of cooperation can be achieved when $n_C \geq 3$ (comparing panels A and C). Considering the total cost minimisation (panels B and D), NEB has a significantly larger area where a low cost is required to achieve high levels of cooperation when compared to POP. Thus, adopting a local interference strategy is largely more efficient than a global population-wide paradigm.

4.2 Conditions for Cost-Efficient Interference

From Figure 2, we also observe that both POP and NEB strategies require a certain threshold of θ to be successful. Next, we derive a closed-form expression for such critical thresholds, for arbitrary b .

From Figure 2B we observe that, for POP to obtain an efficient cost while ensuring high levels of cooperation (the black area near top right corner), θ must be sufficiently high. Our analysis of Figure 3A shows that whenever $\theta \leq 4b - 3$, different cyclic patterns or stable mixed configurations are eventually reached. Particularly, when θ is slightly less than $4b - 3$, the population always ends up in a stable pattern, characterised by a number of standalone D individuals surrounded by C neighbours. When $\theta > 4b - 3$ (i.e. when $b = 1.8$, $\theta > 4.2$), cyclic patterns can be escaped, leading to a small total cost and high levels of cooperation.

Similarly, from Figure 2D, for NEB to obtain an efficient cost while ensuring high levels of cooperation (the black area near top right corner), θ must be sufficiently high. We have derived the threshold of θ for $n_C = 3$: $\theta \geq 4b - 2$ (i.e. when $b = 1.8$, $\theta \geq 5.2$), see Figure 3 (panels B and C). Moreover, since NEB with $n_C = 4$ is equivalent to POP with $p_C = Z$ – both mean always invest in all cooperators – the same condition is obtained for these two limiting cases (i.e. $\theta > 4b - 3$).

We can see that the threshold of θ for NEB with $n_C = 4$ is lower than the one obtained for NEB with $n_C = 3$ ($4b - 3$ compared to $4b - 2$), which is also confirmed by the numerical results in Figure 2D. When θ lies between $4b - 3$ and $4b - 2$, NEB with $n_C = 4$ is more cost-efficient.

We confirm by performing additional simulations that the two conditions hold for other values of b and are also robust for different population sizes Z .

4.3 Comparing with Fitness Observation based Strategies: NEB-i and NEB-ii

We compare the cost efficiency of NEB with $n_C = 3$ and then with $n_C = 4$ where the conditions for their efficiency (derived above) are satisfied, namely, $\theta \geq 4b - 2$ and $\theta > 4b - 3$, respectively. Abusing notations, let us assume that $\theta = (4b - 2) + \epsilon$ and $\theta = (4b - 3) + \epsilon$ for the two cases, respectively. This allows us to conveniently compare these strategies with the more advanced ones, NEB-i and NEB-ii. The parameter ϵ represents the extra amount to the threshold value of θ which guarantees (close to) 100% of cooperation.

Figure 4 shows the total costs for the four interference strategies for varying ϵ . They all guarantee high levels of cooperation (100%) for $\epsilon > 0$. In general, NEB-3 is always sig-




Figure 5: Stochastic update: evolution of strategies (red for D, blue for C) and per-generation cost for NEB with $n_C = 3$ (left column) and POP with $p_C = Z$ (right column). Note that NEB with $n_C = 4$ is equivalent to POP with $p_C = Z$. Similarly to the deterministic update, NEB with $n_C = 3$ has a decreasing per-generation cost while that with $n_C = 4$ leads to an increasing tendency. As a result, the former leads to a significantly lower total cost of investment than the latter one (namely, 2.2×10^5 compared to 2.1×10^6). Parameters: $b = 1.8$, $\theta = 3$, $L = 100$, $K = 0.3$ (the value of K was chosen based on suggestion from previous works; similar results were observed for other values of K).

nificantly more cost-efficient than that with $n_C = 4$. Moreover, assuming that it is possible to assess individual scores of players in a neighbourhood, more cost-efficient strategies are obtained. Namely, both NEB-ii and NEB-i are slightly better than NEB-3.

4.4 Stochastic Update

We confirm that our results are robust with respect to the stochastic update rule (see again Section 3.1). In particular, similarly to the deterministic case, we observe that the local interference approach (NEB) leads to more cost-efficient strategies than the global interference one (POP). Namely, NEB with $n_C = 3$ is highly cost-efficient (while guaranteeing high levels of cooperation) – see Figure 5 for a typical evolution of strategies and per-generation costs (and the total costs) for NEB with $n_C = 3$ and POP with $p_C = Z$. Interestingly, due to the stochastic effect (where there is a chance that a higher scored individual copies the lower scored neighbour’s strategy), it is easier to escape cyclic configurations and reach full cooperation in the population. As a result, a lower threshold of θ is required to reach good performance of the interference strategies (e.g. in Figure 5, $\theta = 3$ was sufficient to reach full cooperation, compared to $\theta \geq 5.2$ in the deterministic case). However, due to this stochastic effect, it takes longer to converge to a homogeneous state (compared Figures 5 and 1 where $n_C = 3$ and $n_C = 4$, noting that NEB with $n_C = 4$ is equivalent to POP with $p_C = Z$).

5 Conclusions and Future Work

In summary, in this paper, we seek to determine how to efficiently interfere in a spatially distributed system (i.e. a structured population) of agents to achieve a desired collective state (i.e. high levels of cooperation). In addition to understanding the mechanisms underlying the emergence of cooperation in a multiagent system, we aim at identifying the most

effective strategy to foster cooperative scenarios. In particular, the cost of interference measured as the consumption of (monetary) resources, and the greater impact we want to make, the higher the cost we have to pay. To tackle this problem, we have combined different interference paradigms with an evolutionary game theoretic model in structured populations, showing that local strategies (even their simplest versions) outperform global strategies in promoting cooperation in networked communities.

In particular, we have systematically studied two major interference approaches: a global paradigm which oversees the status of the whole population when deciding whether to interfere and a local approach which considers local properties of a neighbourhood when deciding to intervene. We have shown that local approaches lead to more cost-efficient interference strategies. Notably, we have shown that a policy which only invests in a cooperator if it has at least a defective neighbour (i.e. NEB-3), is highly efficient, promoting substantial levels of cooperation. We have also derived the analytical condition of optimal thresholds above which it is profitable to invest and confirmed the results are robust for other stochastic update rules. Furthermore, we have demonstrated that more advanced strategies based on neighbourhood properties can be achieved by assessing individual scores of players in a neighbourhood.

Our findings provide useful insights for decision-makers in a large variety of application domains, ranging from drug prevention programmes (e.g., the Unplugged programme of Mentor-Adepis in the UK which aims to promote behaviours that lead to the decrease of alcohol and drug usage among the youth) and international educational aids to developing countries (e.g., the Global Partnership for Education fund), to wildlife conservation initiatives (e.g., WWF) and environmental governance (i.e., how to promote and enforce behaviours that lead to sustainable environment among both industrial and domestic actors). According to our results, it is better to focus on strategies which look at each individual’s neighbourhood and support cooperation based on local information, instead of using global observables as the primary trigger for decision-making.

Note that this work may foster future studies of even more efficient interference strategies. In particular, it would be interesting to consider more adaptive strategies, which modify the amount and the frequency of investment dynamically depending on the current state of the system. The analysis of the resulting systems, however, is not straightforward, as it remains unclear whether to increase or decrease the amount/frequency of investments will lead to more efficient performance. Also, we aim to extend our investigation to systems with other, more complicated scenarios such as heterogeneous networks and multi-player games [Santos *et al.*, 2008] where more behavioural equilibria and degree-dependent strategies are expected [Duong and Han, 2015].

References

- [Aadithya *et al.*, 2011] K. V. Aadithya, T. P. Michalak, and N. R. Jennings. Representation of coalitional games with algebraic decision diagrams. In *AAMAS '11, Taipei, Taiwan, May 2-6, 2011, Volume 1-3*, pages 1121–1122, 2011.

- [Aziz *et al.*, 2010] H. Aziz, F. Brandt, and P. Harrenstein. Monotone cooperative games and their threshold versions. In *AAMAS '10*, pages 1107–1114, 2010.
- [Bachrach *et al.*, 2009] Y. Bachrach, E. Elkind, R. Meir, D. Pasechnik, M. Zuckerman, J. Rothe, and J. Rosenschein. The cost of stability in coalitional games. In *Algorithmic Game Theory*, volume 5814 of *LNCS*, pages 122–134. 2009.
- [Ding *et al.*, 2013] W. Ding, T. Qin, X.-D. Zhang, and T.-Y. Liu. Multi-armed bandit with budget constraint and variable costs. In *AAAI '13*, pages 232–238, 2013.
- [Duong and Han, 2015] Manh Hong Duong and The Anh Han. On the expected number of equilibria in a multi-player multi-strategy evolutionary game. *Dynamic Games and Applications*, pages 1–23, 2015.
- [Endriss *et al.*, 2011] U. Endriss, S. Kraus, J. Lang, and M. Wooldridge. Incentive engineering for boolean games. *IJCAI '11*, pages 2602–2607, 2011.
- [Franks *et al.*, 2013] H. Franks, N. Griffiths, and A. Jhumka. Manipulating convention emergence using influencer agents. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 26(3):315–353, 2013.
- [Franks *et al.*, 2014] H. Franks, N. Griffiths, and S. S. Anand. Learning agent influence in mas with complex social networks. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 28(5):836–866, 2014.
- [Guha and Munagala, 2007] S. Guha and K. Munagala. Approximation algorithms for budgeted learning problems. *STOC '07*, pages 104–113, 2007.
- [Han *et al.*, 2014] The Anh Han, Long Tran-Thanh, and Nicholas R Jennings. The cost of interference in evolving systems. In *The 17th International Workshop on Coordination, Organisations, Institutions and Norms (COIN 2014)*, 2014.
- [Han *et al.*, 2015] The Anh Han, Long Tran-Thanh, and Nicholas R Jennings. The cost of interference in evolving multiagent systems. In *AAMAS'2015*, pages 1719–1720, 2015.
- [Han *et al.*, 2017a] The Anh Han, Luís Moniz Pereira, and Tom Lenaerts. Evolution of commitment and level of participation in public goods games. *Autonomous Agents and Multi-Agent Systems*, 31(3):561–583, 2017.
- [Han *et al.*, 2017b] The Anh Han, Luís Moniz Pereira, Luis A Martinez-Vaquero, and Tom Lenaerts. Centralized vs. personalized commitments and their influence on cooperation in group interactions. In *AAAI*, pages 2999–3005, 2017.
- [Han, 2016] The Anh Han. Emergence of social punishment and cooperation through prior commitments. In *AAAI'2016*, pages 2494–2500, Phoenix, Arizona, USA, 2016.
- [Harrenstein *et al.*, 2014] P. Harrenstein, P. Turrini, and M. Wooldridge. Hard and soft equilibria in boolean games. In *AAMAS '14*, pages 845–852, 2014.
- [Levin, 2000] Simon A Levin. Multiple scales and the maintenance of biodiversity. *Ecosystems*, 3(6):498–506, 2000.
- [Levit *et al.*, 2013] V. Levit, T. Grinshpoun, A. Meisels, and A. L. C. Bazzan. Taxation search in boolean games. In *AA-MAS '13, Saint Paul, MN, USA, May 6-10, 2013*, pages 183–190, 2013.
- [Madani *et al.*, 2004] O. Madani, D. J. Lizotte, and R. Greiner. The budgeted multi-armed bandit problem. *ICOLT '04*, pages 643–645, 2004.
- [Nowak and May, 1992] Martin A Nowak and Robert M May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.
- [Nowak and Sigmund, 2005] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437(1291-1298), 2005.
- [Raghunandan and Subramanian, 2012] M. A. Raghunandan and C. A. Subramanian. Sustaining cooperation on networks: an analytical study based on evolutionary game theory. In *AAMAS'12*, pages 913–920, 2012.
- [Rand *et al.*, 2013] David G. Rand, Corina E. Tarnita, Hisashi Ohtsuki, and Martin A. Nowak. Evolution of fairness in the one-shot anonymous ultimatum game. *Proc. Natl. Acad. Sci. USA*, 110:2581–2586, 2013.
- [Rand *et al.*, 2014] David G Rand, Martin A Nowak, James H Fowler, and Nicholas A Christakis. Static network structure can stabilize human cooperation. *Proc Natl Acad Sci USA*, 111(48):17093–17098, 2014.
- [Sahraei *et al.*, 2014] B. R. Sahraei, H. Bou-Ammar, D. Bloembergen, K. Tuyls, and G. Weiss. Evolution of cooperation in arbitrary complex networks. In *AAMAS '14, Paris, France, May 5-9, 2014*, pages 677–684, 2014.
- [Santos *et al.*, 2006] F. C. Santos, J. M. Pacheco, and T. Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc Natl Acad Sci USA*, 103:3490–3494, 2006.
- [Santos *et al.*, 2008] F. C. Santos, M. D. Santos, and J. M. Pacheco. Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454(7201):213, 2008.
- [Sigmund *et al.*, 2001] K. Sigmund, C. Hauert, and M. Nowak. Reward and punishment. *Proc. Natl Acad Sci USA*, 98(19):10757–10762, 2001.
- [Sigmund *et al.*, 2010] Karl Sigmund, Hannelore De Silva, Arne Traulsen, and Christoph Hauert. Social learning promotes institutions for governing the commons. *Nature*, 466(7308):861, 2010.
- [Szabó and Fath, 2007] György Szabó and Gabor Fath. Evolutionary games on graphs. *Physics reports*, 446(4-6):97–216, 2007.
- [Tran-Thanh *et al.*, 2012] L. Tran-Thanh, A. Chapman, A. Rogers, and N. R. Jennings. Knapsack based optimal policies for budget-limited multi-armed bandits. *AAAI '12*, pages 1134–1140, 2012.
- [Traulsen and Nowak, 2006] A. Traulsen and M. A. Nowak. Evolution of cooperation by multilevel selection. *Proc Natl Acad Sci USA*, 103(29):10952, 2006.
- [Traulsen *et al.*, 2006] A. Traulsen, M. A. Nowak, and J. M. Pacheco. Stochastic dynamics of invasion and fixation. *Phys Rev E*, 74(1):011909, 2006.
- [Tuyls and Parsons, 2007] K. Tuyls and S. Parsons. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416, 2007.
- [Vasconcelos *et al.*, 2013] V. V. Vasconcelos, F. C. Santos, and J. M. Pacheco. A bottom-up institutional approach to cooperative governance of risky commons. *Nature Climate Change*, 3(9):797–801, 2013.
- [Wooldridge, 2012] M. Wooldridge. Bad equilibria (and what to do about them). *ECAI '12*, pages 6–11, 2012.
- [Zisis *et al.*, 2015] Ioannis Zisis, Sibilla Di Guida, The Anh Han, Georg Kirchsteiger, and Tom Lenaerts. Generosity motivated by acceptance - evolutionary analysis of an anticipation games. *Scientific reports*, 5(18076), 2015.