

# Artificial Evolution for the Detection of Group Identities in Complex Artificial Societies

Corrado Grappiolo, Julian Togelius  
Center for Computer Games Research  
IT University of Copenhagen, Denmark  
Email: cogr@itu.dk, juto@itu.dk

Georgios N. Yannakakis  
Department of Digital Games  
University of Malta, Malta  
Email: georgios.yannakakis@um.edu.mt

**Abstract**—This paper aims at detecting the presence of group structures in complex artificial societies by solely observing and analysing the interactions occurring among the artificial agents. Our approach combines: (1) an unsupervised method for clustering interactions into two possible classes, namely *in-group* and *out-group*, (2) reinforcement learning for deriving the existing levels of collaboration within the society, and (3) an evolutionary algorithm for the detection of group structures and the assignment of group identities to the agents. Under a case study of static societies — i.e. the agents do not evolve their social preferences — where agents interact with each other by means of the Ultimatum Game, our approach proves to be successful for small-sized social networks independently on the underlying social structure of the society; promising results are also registered for mid-size societies.

**Keywords**—Group Identity Detection, Evolutionary Computation, Artificial Societies, Emergence of Complexity.

## I. INTRODUCTION

When populations of socially-driven individuals are let to interact with each other, in absence of any form of centralised control ruling their comportment, complex dynamics occur yielding *emergent* global patterns such as culture and friendship. Such patterns occur due to the ability of the individuals to *adapt* their behaviour based on the experience gathered from the interaction with the other individuals (i.e. self-organisation) [1]. Societies are Complex Adaptive Systems (CAS) [2] which means that they are hard to model. In particular, the modelling of a society's *Emergence of Complexity* — i.e. the transition from the local perspective (i.e. individuals and their interactions) to the global perspective (e.g. friendship social networks) [2] — is far from being trivial. The key reason that makes the modelling challenging is the recurrent influence the global structures have on the behaviour of the individuals [1].

This paper focuses on the automatic detection of friendship-based group structures as a form of global pattern formation. The approach we take exploits the recursive structure of the CAS-society (see Figure 1): we aim to infer existing groups by analysing how the individuals behave when they interact with each other. We claim that the influence the group structures have on the individuals can be detected by observing the different levels of altruism manifested during the interactions, given that collaboration can be interpreted as reciprocal altruism [1], and that individuals belonging to the same group (*in-group*) tend to collaborate more with each other than when they interact with individuals belonging to other

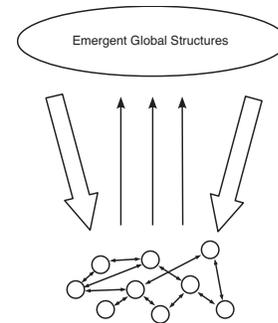


Fig. 1. The canonical representation of a Complex Adaptive System [2]. The emergent global structures are not only the result of the interactions among the CAS-society (i.e. *Emergence of Complexity*); they also influence the behaviours of the society's individuals.

groups (*out-group*) [3], [4]. Therefore, the aim of our research is the definition of a computational framework — hereafter called Group Modelling (GM) framework — capable of measuring the level of collaboration existing among the individuals and assign group identities to them. The GM framework we envision should be capable of operating in real-time, so that it can provide insight of global-level behaviours and dynamics, such as group formation and evolution. The GM framework we propose is composed of three pipelined modules, as depicted in Figure 2: (1) at first, the continuous flow of interactions occurring in the society is analysed by an *Interaction Classifier* ( $\mathcal{I}$ ), which labels the individuals as being either *in-group* or *out-group*; (2) subsequently, a Collaboration Learning module ( $\mathcal{C}$ ) interprets the *in/out-group* labels as rewards which are then used by a Reinforcement Learning [5] update rule in order to represent the existing levels of altruism among the agents; (3) finally, a Group Identity Detection module ( $\mathcal{G}$ ) interprets the learned collaboration values as a weighted, directed and fully connected network and aims to identify community structures within it, thus, leading to the assignment of group identities to the agents.

As a case study, the GM framework presented in this paper aims to detect consolidated group identities in complex artificial societies of believable, human-like agents — manifesting social preferences such as altruism and reciprocity [6] when interacting with each other — in the Ultimatum (or Bargain Social Dilemma) Game [7]. Consolidated groups are groups of static agents which have reached an equilibrium in the game.

The task of optimal network partitioning into community

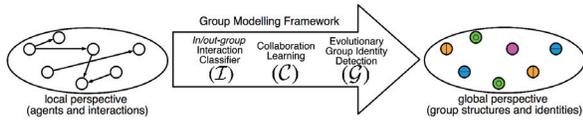


Fig. 2. The schematic representation of the GM framework proposed.

structures is known to be computationally hard (NP-complete) over the set of all graphs of a given size [8], [9]; the novelty of this study is the successful application of evolutionary computation to derive an approximate solution for the problem in fully connected, directed and weighted networks.

The work presented here builds upon and extends an earlier study of the authors [10] which focused on the  $\mathcal{I}$  module. In this paper we compare the performance of an evolutionary algorithm (EA) against the performance obtained by a hierarchical Complete Link clustering algorithm [10] for correctly inferring the group identities of the agents. Results obtained on two society sizes (20 and 50 agents), under different distributions of social relationship statuses (friend, acquaintance and stranger), showcase that the EA approach is able to detect the *true* group structures quickly and efficiently for small-sized societies independently of the underlying social structures of the society. Moreover, the key findings demonstrate the promise of the GA for mid-sized societies.

## II. RELATED WORK

There is a number of studies investigating groups of agents and their behaviours in artificial societies, virtual environments, and multiplayer games, but also in real life scenarios. Nowak et al. [11], among others, focus on the evolution of collaboration by evolving the policies of artificial agents. The approach, however, neglects the impact of collaboration on group formation. Similar research was conducted by Hammond and Axelrod [12], though it was focused on the evolution of ethnocentrism. Among the studies on collective behaviour, Lerman and Galstyan [13] create mathematical models, through differential equations, of the collective behaviour of simple multi-agent systems, such as social insects. Their approach differs from ours in that they aim to build a generic model of an agent, based on observations, and then devise a mathematical model of it. Martinez et al. [14] investigated the use of rule-learning algorithms to predict group behaviours in artificial societies. Their method is based on historical data and even though we share a common goal (i.e. modelling of group dynamics), that study does not aim to model group identities in real time.

The use of social network in Agent-Based Modelling is not novel [15]. For instance, Xianyu [16] utilises spatial networks in order to regulate the interactions among the agents and allow them to manifest inequity aversion-based social preferences. Similar approaches to the use of complex network has also seen, among others, in the work of Liu et al. [17]. Our work differ from the assumption that our agents are spatially close to each other and interact with any of the other agents in the society; moreover, the use we make of the network is to represent their social relationships, i.e. friendship.

With respect to modelling community structures in real-life situations, Palla et al. [18] adopt a data-driven approach

for the detection of real-life community structures, by means of clique percolation, of phone call networks. Similarly, Eagle et al. [19] rely on phone data to detect community structures of reciprocal friendship, and cross-validate them against the self-reports gathered from the participants of their experiment. That study is possibly the most relevant to this paper since self-reports gathered in that study could correspond to our notion of *true* group structures. However, the data-driven social networks built in [19] — i.e. based on the phone calls — are undirected and unweighted; furthermore, their analysis is based on the frequency and proximity of the interactions, rather than on the investigation of their qualities, such as collaboration.

EAs have been used for the detection of community structures, such as in the studies on Pizzuti [20] and Tasgin [21], and Liu et al. [22], who propose Multi-objective Evolutionary Computation for the detection of overlapping community structures. Nevertheless, their work did not examine directed weighted networks. Farmer and Fotheringham [9], on the other hand, aim to identify community structures in such complex networks. While we are inspired by that study, we use principles of evolutionary computation (instead of spectral partitioning) for identifying group structures.

## III. INTERACTION PROTOCOL: THE ULTIMATUM GAME

The artificial agents interact with each other exclusively by means of the Ultimatum, or Bargain Social Dilemma, Game [7]. We decided to implement this game as our interaction protocol due to its popularity in many fields, such as Behavioural and Social Sciences, Economics and Evolutionary Computation (see [7], [24]–[26] for a non exhaustive list). Moreover, we believe that such resource-based interaction protocol can be easily converted into game-based mechanics, given our aim to model group structures in collaborative multiplayer games.

The social dilemma within this game is embedded as follows: within a population  $P = \{a_0, \dots, a_{n-1}\}$  of  $n$  artificial agents, at each round of interaction  $t$  (Episode), an agent  $a_i \in P$  is randomly selected — through uniform distribution — to become the provider agent of the bargain game, whilst the remaining  $n-1$  agents take the role of receiver agents. The provider has now to bargain, with each  $n-1$  receiver agents, an equal endowment  $e$ .  $a_i$  makes an offer  $0 \leq o_{i,j} \leq e$  to each receiver agent  $a_j \in P \setminus \{a_i\}$ . The receiver agent has now two options: it can accept the bargain ( $d_{j,i} = 1$ ), in which case  $a_j$  will gain  $o_{i,j}$  and  $a_i$  will gain  $e - o_{i,j}$ , or refuse the bargain ( $d_{j,i} = 0$ ), in which case both agents will gain zero. We will hereafter refer to the whole set of  $n-1$  bargains made by a single provider agent  $a_i$  as a *one-to-many* interaction [10].

## IV. *True* GROUP STRUCTURES AND AGENT BEHAVIOUR

The central behavioural feature we aim to investigate in our research is the capability of each artificial agent in the society to manifest social preferences [6], here intended as the ability to be differently altruistic based on which other agents it interacts with.

In this study we consider three social relationship types: two agents can either be Strangers (S), Acquaintances (A) or Friends (F). All relationship types are *reciprocal*, meaning that e.g. if agent  $a_i$  considers agent  $a_j$  as a friend, the same is being

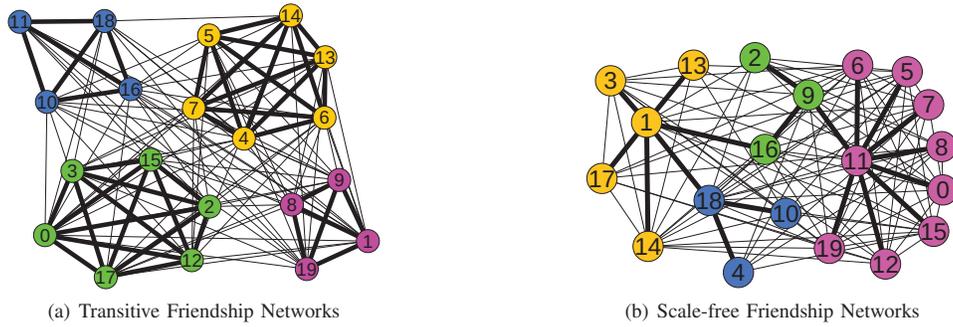


Fig. 3. An example two social networks  $SN = SN_F + SN_A$  considered in our study, for a population of 20 agents. In both cases,  $SN_A$  was built with  $p_A = 0.5$ . The group identities are represented by the colours of the vertices. Figure 3(a) depicts four groups built upon transitive friendship networks ( $SN_F$ ), Figure 3(b) depicts four groups built upon scale-free friendship networks. The group identities were built my means of the Clauset Newman Moore algorithm [23], on  $SN_F$  only.

perceived by  $a_j$  towards  $a_i$ . The whole set of  $\frac{1}{2}[n(n-1)]$  reciprocal relationship values can be interpreted as a fully connected, weighted, discrete social network  $SN = (P, R)$ , in which vertex  $i \in P$  represents agent  $a_i$ , and the weighted edge  $w_{i,j} = \{0, 1, 2\} \in R$  represents, respectively, the S, A or F relationship type existing between agents  $a_i$  and  $a_j$ . An alternative way to imagine  $SN$  is as the sum of two mutually exclusive friendship and acquaintanceship networks,  $SN = SN_F + SN_A$ .

In this paper we considered two methods of generating  $SN_F$  and one method for generating  $SN_A$ . With respect to the former, the first method instantiates transitive networks (i.e. if  $w_{i,j} = 2, w_{i,k} = 2 \Rightarrow w_{j,k} = 2$ ) by firstly defining the number of groups structures in the artificial society, according to normal distribution  $N(\mu_n, \sigma_n)$ , and then by assigning group identities to each agents uniformly; the second method builds scale-free friendship networks by starting from two friend agents and iteratively adding one node and one friendship link, in accordance to Barabási and Albert [27], for  $n - 2$  loops.

The generation of  $SN_A$ , instead, is done stochastically as follows: each non-friend agent couple has a probability  $p_A$  of being an acquaintance and  $1 - p_A$  of being a stranger.

Figures 3(a) and 3(b) depict, respectively, an example of transitive and scale-free friendship networks, and relative group identities. We define the *True* group structures of the society as the community structures, obtained by means of the Clauset Newman Moore algorithm [23], on  $SN_F$  only, whilst  $SN_A$  will have the function to introduce noise in the observations of the social preferences.

The three relationship types and the group structures influence the altruistic behaviour of the agents when interacting with each other, causing the agents to manifest social preferences [6]. Although in this study we assume consolidated groups and therefore no agent adaptation, we still aim to study believable agents with social preferences simulating human-like behaviours. The social preferences are driven by the findings of Marzo et al. [28]: (1) each agent maintain a social perception  $s_{i,j} \in [0, 1]$  of each other agent  $a_j \in P$ ;  $a_i$  will consider  $a_j$  a stranger if  $0 \leq s_{i,j} \leq 1/3$ , acquaintance if  $1/3 < s_{i,j} \leq 2/3$ , friend if  $2/3 < s_{i,j} \leq 1$ ; (2) the offers are

generated as follows:

$$o_{i,j} = 1/2 e (s_{i,j} + r) \quad (1)$$

where  $1/2$  indicates that the provider agent always aims to maximise its own gain, and  $r$  is uniformly sampled within the  $[-0.1, 0.1]$  interval; (3) the offer decision making policy is the following:

$$d_{i,j} = \begin{cases} 1 & \text{if } o_{j,i} \leq s_{i,j} \\ 0 & \text{else} \end{cases} \quad (2)$$

## V. GROUP MODELLING FRAMEWORK

The GM framework proposed (see Figure 2) is composed of three connected modules, described in details in their respective subsections, which are executed sequentially at the end of each Episode  $t$ . First, an *Interaction Classifier* module ( $\mathcal{I}$ , Subsection V-A), analyses the flow of interactions and for each provider-receiver agent pair decides whether they belong to the same group (*in-group*) or not (*out-group*) [3]. Second, a *Collaboration Learning* module ( $\mathcal{C}$ , Subsection V-B) interprets  $\mathcal{I}$ 's *in/out-group* labels as rewards in order to learn up-to-date collaboration values of each agent pairs in the society via Reinforcement Learning-inspired update rules [5]. Third, a *Group Identity Detection* module ( $\mathcal{G}$ , Subsection V-C), processes the learned collaboration values computed by  $\mathcal{C}$  via artificial evolution [29] to partition the society into group structures and assign group identities to each agent.

### A. Interaction Classifier

In this study we use the one-to-many *Agent Preference* probabilistic classifier [10]: given the whole set of  $n - 1$  offers formulated by provider agent  $a_i$ :

$$O_i = \{o_{i,0}, \dots, o_{i,i-1}, o_{i,i+1}, \dots, o_{i,n-1}\} \quad (3)$$

$\mathcal{I}$  firstly performs a min-max normalisation of all the offers:

$$\hat{o}_{i,j} = \frac{o_{i,j} - \min\{O_i\}}{\max\{O_i\} - \min\{O_i\}} \quad (4)$$

then, it interprets  $\hat{o}_{i,j}$  as the probability that  $a_j$  is *in-group* with  $a_i$ ; given a uniform-distribution sampled random number  $r \in [0, 1]$ ,  $\mathcal{I}$ 's output labels are obtained as follows:

$$\mathcal{I}(i, j) = \begin{cases} \text{in-group} & \text{if } r \leq \hat{\delta}_{i,j} \\ \text{out-group} & \text{otherwise} \end{cases} \quad (5)$$

### B. Collaboration Learning

Given the CAS perspective of the problem, we assume that the interactions occurring in the society have the Markov property — i.e. the agents perform actions based on their current internal *state*, which represents their social relationships perceptions and group identities. As a consequence, the whole society has the Markov property. With respect to the canonical Reinforcement Learning terminology, hence,  $\mathcal{C}$  aims to learn the collaboration-function values,  $C_{i,j}$ , of the whole society.

At the end of the  $t$ -th flow of interactions,  $\mathcal{C}$  performs collaboration learning by means of the *constant- $\alpha$  Monte Carlo* update rule for non-stationary environments [5], [10]:

$$C_{i,j}(t) = C_{i,j}(t-1) + \alpha [R_{i,j}(t) - C_{i,j}(t-1)] \quad (6)$$

where  $0 < \alpha < 1$  is a constant step-size parameter regulating the learning rate, and  $R_{i,j}(t)$  is the immediate reward of the interaction between  $a_i$  and  $a_j$  after the  $t$ -th episode, where  $R_{i,j}(t) = 1$  if  $\mathcal{I}(i, j) = \text{in-group}$  and  $R_{i,j}(t) = 0$  if  $\mathcal{I}(i, j) = \text{out-group}$ .

### C. Group Identity Detection

Collaboration matrix  $C$  is finally processed by  $\mathcal{G}$  in order to partition the agents into group structures and assign group identities.  $C$  can be viewed as a fully-connected, directed, weighted and continuous network, and the partitioning task can be interpreted as the detection of community structures. However, the task of optimal network partitioning into community structures is computationally hard, known to be NP-complete over the set of all graphs of a given size [8], [9]. Given the complexity, we decide to derive an approximate solution for this problem by means of evolutionary computation [29]. Evolutionary algorithms have already been used to detect community structures in networks [20], [21]; though, to the best of our knowledge, there has not been any attempt to use them on fully connected, directed and weighted and continuous networks, which  $C$  represents a good instance.

The fitness function of our EA is a modularity measure, introduced by Leich and Newman for weighted networks [8] and then successfully extended to directed weighted networks by Farmer and Fotheringham [9]:

$$f(x) = \frac{1}{w} \sum_{i,j} \left( C_{i,j} - \frac{w_i^{\text{in}} w_j^{\text{out}}}{w} \right) \delta(c_i, c_j) \quad (7)$$

where  $w = \sum_{i,j} C_{i,j}$  is the total sum of collaboration values of the matrix/network,  $w_i^{\text{in}}$  is the in-degree of vertex  $i$ /agent  $a_i$ ,  $w_j^{\text{out}}$  is the out-degree of vertex  $j$ /agent  $a_j$ ,  $c_i$  is the label of the community/group identity to which vertex  $i$ /agent  $a_i$  is assigned, and  $\delta(c_i, c_j)$  is the Kronecker delta symbol, for which  $\delta(c_i, c_j) = 1$  if  $c_i = c_j$  and  $\delta(c_i, c_j) = 0$  otherwise.

The EA we implemented considers a population of  $m$  genotypes; each genotype  $x_i$  has  $n$  chromosomes: each chromosome  $k$  of  $x_i$ ,  $x_i^k$ , represents the group identity  $c_k$  of agent  $a_k$ . The possible values of the chromosomes are taken from an alphabet  $G$ , which holds the labels of possible group structures. The set of different  $c \in G$  symbols in  $x_i$ ,  $\hat{G}_i \subseteq G$ , represents the group structures detected by  $x_i$ .

At the beginning of each experimental setup,  $G$  is initialised as  $G = \{0, 1\}$  and the genotypes are initialised with chromosomes values sampled uniformly within  $G$ . In other words, our GM framework initially assumes the existence of two group structures within society  $P$ .

At the end of Episode  $t$ , once  $C$  has been updated, the EA is reinitialised by following these five steps: (1) the EA recalculates the fitness values of its genotypes, and sorts the population by descending fitness, since the updates of  $C$  performed by the  $\mathcal{C}$  module might have generated changes in the genotype ranking of the EA; (2) only the genotypes with the highest possible fitness,  $\{x_0, \dots, x_j\}$  remain in the population, the remaining  $m - (j+1)$  genotypes are discarded; (3) the EA calculates the average number of group structures detected by the individuals who survived the previous step,  $\mu_{\hat{G}} = (\sum_{i=0}^j \hat{G}_i) / (j+1)$  (4) the EA initialises  $G$  as  $G = \{0, \dots, \mu_{\hat{G}}/2\}$ ; (5) the EA fills the population by generating  $m - (j+1)$  genotypes which chromosomes are uniformly sampled within  $G = \{0, \dots, \mu_{\hat{G}}/2\}$ . The fitness of the  $m - (j-1)$  genotypes is calculated and the whole population of  $m$  genotypes is sorted by decreasing fitness.

At this point, the evolutionary process starts. Each evolutionary generation is organised as follows: (1) the first half of the population survives the evolution, the second half is replaced by  $m/2$  offsprings; (2) each new offspring  $x_h$  is generated by means of uniform crossover and mutation: first, two genotypes,  $x_i$  and  $x_j$ , are chosen, among those who survived the previous step, according to rank selection; uniform crossover is then applied: each chromosome  $x_h^k$ ,  $k = 0 \dots n-1$ , is uniformly chosen between either  $x_i^k$  or  $x_j^k$  with probability 0.5; mutation is finally applied iteratively: as long as a uniformly sampled random value  $r \in [0, 1]$ , is less than a mutation probability  $p_{mut}$ , a random chromosome of  $x_h$  is uniformly chosen and its value is replaced by one of the symbols of  $G$  uniformly chosen; (3) the EA calculates the fitness values of its new  $m/2$  offsprings and sorts the whole population of  $m$  genotypes in a descending order.

At the end of each evolutionary generation, the average fitness value of the first half of the genotypes is calculated. If the average is not improving for  $l$  consecutive generations (local limit),  $G$  is updated by inserting a new symbol; if the average is not improving for  $g$  consecutive generations (global limit), the EA stops and the genotype with the highest scored fitness is the one chosen by  $\mathcal{G}$  module to represent the group structures and identities of the artificial society.

## VI. EXPERIMENTAL SETUP AND RESULTS

As previously mentioned, the work proposed in this paper extends our earlier work on the detection of consolidated group identities [10] by considering evolutionary computation as a more efficient way to implement the Group Identity Detection component. In that study the group identities where detected

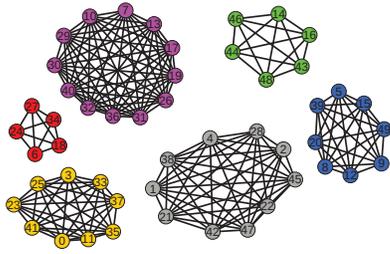


Fig. 4. The six *true* group structures for  $n = 50$ ,  $p_A = 0$  and transitive friendship networks for which the EA framework obtained the smallest *nme* error of 0.1, by detecting the whole red group of agents as being part of the grey group of agents.

via an agglomerative Complete-Link (CL) algorithm with the elbow-rule as stopping criterion [10]. We will therefore consider the performance obtained by the previous best GM configuration as a baseline performance for this paper. The two approaches share the *Agent Preference* interaction classifier (see (5)) and the *constant- $\alpha$  Monte Carlo* update rule (see (6)). In all the experimental setups presented here, we used the following parameters for the GM frameworks:  $\alpha = 0.1$  for  $\mathcal{C}$ ; genotype population size of  $m = 40$ , local and global limit  $l = 10$  and  $g = 50$  generations respectively, mutation probability  $p_{mut} = 0.8$  for EA's  $\mathcal{G}$  component.

The performance evaluation of the GM frameworks was calculated as the normalised mismatch error  $nme(T, I)$ , between the *True* ( $T$ , see Section IV) and the inferred ( $I$ , see Subsection V-C) group identities, calculated as follows:

$$nme(T, I) = \frac{n - h(T, I)}{n}$$

where  $h(T, I)$  is the maximum assignment score obtained by means of Kuhn's Hungarian algorithm [30]. We have considered two artificial society sizes, of  $n = 20$  and  $n = 50$  agents, with  $SN_A$  built according to two probabilities:  $p_A = 0.5$  and  $p_A = 0$ . Finally, the following sample means and sample standard deviations of the number of group structures, generated by the Scale-free method for  $SN_F$  throughout the whole set of experiments conducted, were used to generate the number of transitive groups:  $N(\mu_{20} = 4.43, \sigma_{20} = 0.70)$  and  $N(\mu_{50} = 7.44, \sigma_{50} = 1.10)$ ;

The results presented in this paper are based on the average performance obtained by the repeated execution of the same experimental setup for 10 independent runs. The graphs also illustrate the best performance obtained, by EA and CL depending on the experimental setup, across the 10 execution runs.

#### A. Detection of Group Identities in Transitive Friendship Networks

Figure 5 depicts the average and best performance of the EA and CL algorithms when transitive  $SN_F$  networks are used for both 20 and 50 agents and under two different probabilities of acquaintanceship ( $p_A$ ) between non-friend agents. The first, straightforward finding, is that CL manages to reach zero misclassification error in all 40 total runs, hence independently on  $n$  and  $p_A$ . A near-optimal performance is matched by EA

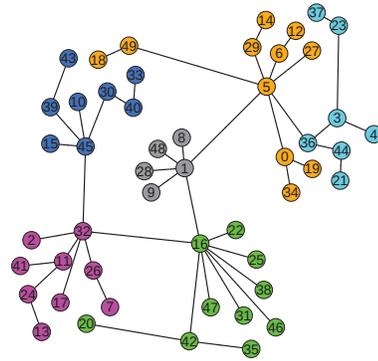


Fig. 6. Scale-free *true* group structures for  $n = 50$  and  $p_A = 0$  for which the EA framework obtained the smallest *nme* error of 0.16.

for the  $n = 20$  case ( $nme = 0$  in 19 out of 20 runs) but not for  $n = 50$ , for which the average *nme* error settles at about 0.26 across the two  $p_A$  settings ( $nme$  is 0.254 for  $p_A = 0.5$  and 0.268 for  $p_A = 0$  — see Figure 5(b) and Figure 5(d)).

We argue that the evolutionary approach of EA is sound and that the relatively poorer performance for  $n = 50$ , independently on the values of  $p_A$ , is mostly due to suboptimal settings for the evolutionary computation parameters. To support this claim, Figure 4 depicts the *true* group structures obtained from transitive friendship networks, and the relative group identities, for the experimental setup in which EA obtained the best performance, which is depicted in Figure 5(d). In this society structure the EA considered the *true* red group (i.e. agents  $a_6, a_{18}, a_{24}, a_{27}$  and  $a_{34}$ ) as belonging to the grey group, hence yielding an error of  $nme = 0.10$ . Two possible reasons would have led to this suboptimal solution: either matrix  $C$  is not well representing the ongoing collaboration values among the society, or matrix  $C$  is correct but  $\mathcal{G}$  does not thoroughly search for the best partitioning. The fitness value of the *true* group structures calculated on the learned  $C$  matrix gives a Modularity Measure of 0.314, whilst the inferred group structures give a Modularity Measure of 0.299. Moreover, the fitness value calculated for the *true* group structures was the highest we managed to obtain. Such findings support that the  $C$  matrix is accurate. In addition, Figure 8 depicts the average landscape of the highest fitness values recorded at the end of the evolutionary process for each Episode for  $n = 50$  under the two different settings for  $p_A$ , across the 10 experimental runs. The EA reaches an average highest fitness of 0.26 for  $p_A = 0$ , and we know empirically — at least for unweighted directed networks [8] — that a Modularity Measure of 0.3 indicates good partitioning into community structures. In summary, indications suggest that the evolutionary process would have reached the global optima, providing a more thorough investigation of the parameter space.

#### B. Detection of Group Identities in Scale-free Friendship Networks

Following the same experimental protocol as before, graphs in Figure 7 depict the average and best performance of the EA and CL algorithms when friendship networks with a scale-free property are used. Not surprisingly, the CL fails to detect the *true* group structures, which suggests that the algorithm is

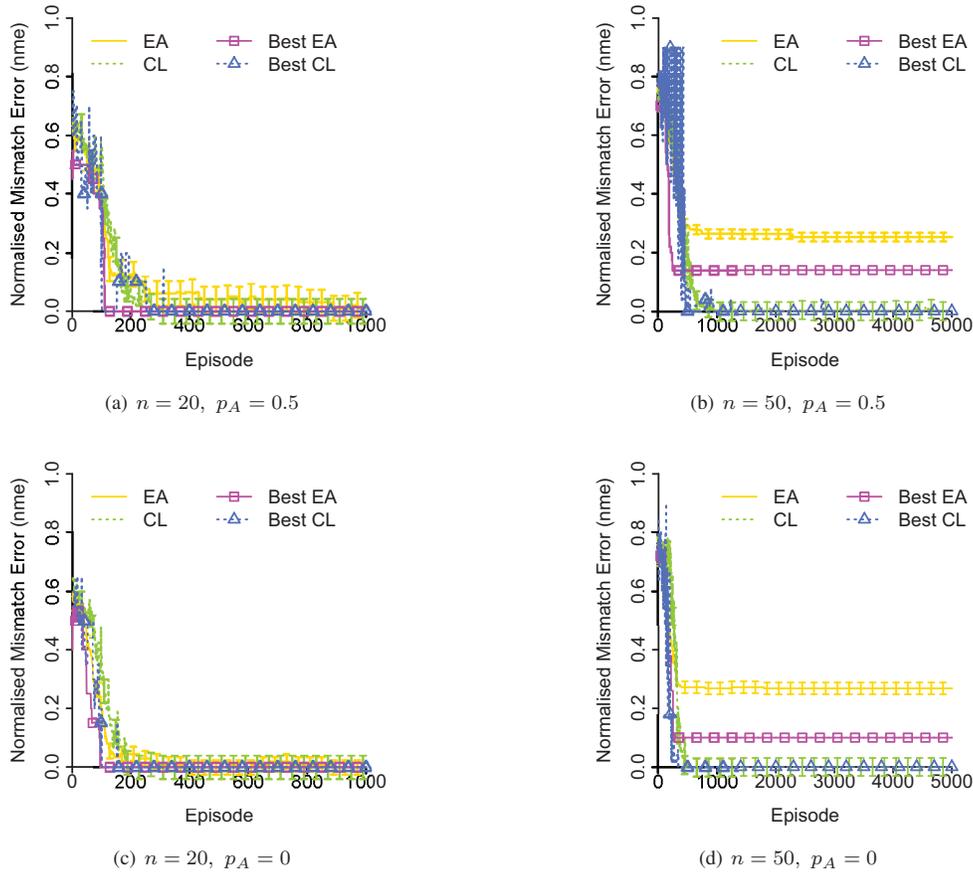


Fig. 5. Best and average performance across 10 runs of the two  $\mathcal{G}$  modules (EA and CL) in case of transitive friendship networks.

not robust across different friendship network structures. This is due to the very nature of the clustering algorithm, which hierarchically builds concentric groups, unlike the nature of the community structures, which define group identities by considering within community and between communities vertex degrees. On the other hand, the EA manages to reach very low misclassification errors. Since the Modularity Measure used as the EA's fitness function is the extension of the one used to build the *true* group structures, and since the  $C$  matrix proved to well represent the existing collaboration values among the agents, we argue that the evolutionary process is beneficial for the accurate inference of the *true* group identities.

Nonetheless, EA reaches  $nme = 0$  only three times out of 20 for  $n = 20$  and never for  $n = 50$ . We argue that the consideration made in Subsection VI-A still holds: we are confident that better settings for the evolutionary process parameters would lead to lower misclassification error. In addition, insights for the performance of the EA can be found when analysing the Modularity Measure (7) with respect to group structures inferred by the EA. The network ( $SN = SN_F$ ) of the best performing EA depicted in Figure 6 corresponds to an obtained misclassification error of 0.16 (see Figure 5(d)). The colour of the agents represent the *true* group identities, whilst the spatial organisation depicts the inferred group identities. Hence, agents  $a_{18}$  and  $a_{49}$  are identified as belonging to the

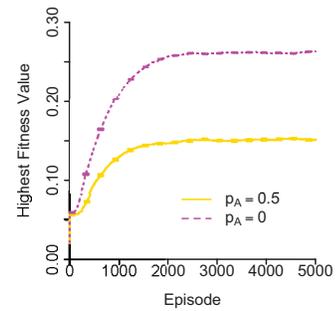


Fig. 8. Highest fitness values obtained for  $n = 50$  within transitive friendship networks and across the two  $p_A$  values considered.

blue group,  $a_{20}$  as belonging to the purple group,  $a_{23}$  and  $a_{37}$  as belonging to the orange group, and finally  $a_0$ ,  $a_{19}$  and  $a_{34}$  as belonging to the cyan group. The fitness value of the *true* and the inferred group structures is 0.127 and 0.129, respectively, suggesting that the EA manages to find a better solution for a community structure, given the  $C$  matrix. In other words, the causes for the EA's misclassification error are more likely to be found in the previous modules of the learning framework, i.e.  $\mathcal{I}$  and  $\mathcal{C}$ .

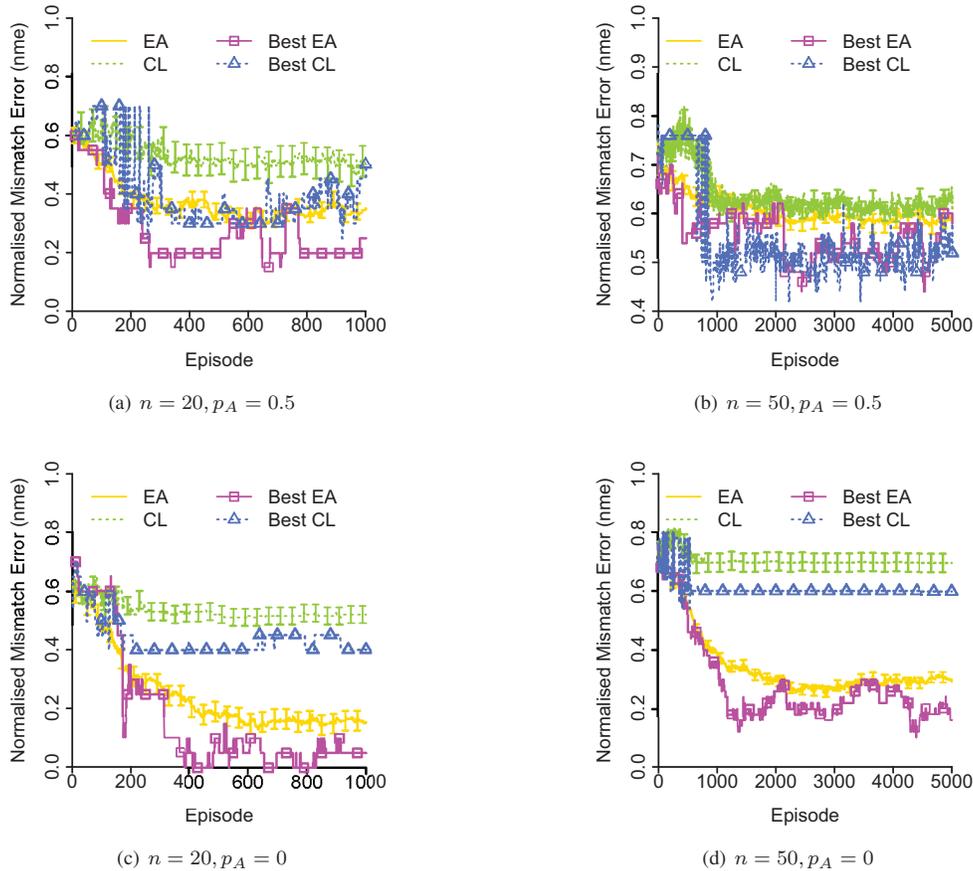


Fig. 7. Best and average performance across 10 runs of the two  $\mathcal{G}$  modules (EA and CL) in case of scale-free friendship networks.

## VII. DISCUSSION AND FUTURE WORK

The promising results presented in Section VI allow us to state that the GM framework, with a  $\mathcal{G}$  module based on evolutionary computation is well robust across friendship network types for small societies (i.e.  $n = 20$  agents), though it needs further investigation, both under the form of fine parameter tuning, and by considering further modifications of the evolutionary algorithm for increasing and maintain genetic diversity, such as fitness sharing and island models.

As previously hinted, the fact that the performance of EA is dependent on the probabilities of acquaintanceship,  $p_A$ , opens for the further investigation of whether the GM framework is capable of predicting properties of future group dynamics (e.g. the migration of agent from one group to another) as Figure 6 might seem to suggest. The peculiar property of  $\mathcal{G}$  to update its alphabet of groups  $G$  based on the results obtained in the previous Episode suggests that the framework would be robust with changes in the emergent structures of the society, though further investigation of the most appropriate learning rules (i.e. those defining the  $\mathcal{C}$  module) would be required.

Our evolutionary approach aims also to contribute to the more general research field of network analysis. The Modularity Measure used as fitness function (7) allows for the use of EAs for the detection of community structures in directed, weighted and continuous networks. An additional step of our

research would require a comparison of performances between our algorithm and e.g. to those of Pizzuti [20], Tasgin [21], and Liu et al. [22]; we will also compare the performance of our evolutionary approach against some other heuristics, such as Spectral Partitioning [9].

The framework makes use of an *one-to-many* Interaction Classifier,  $\mathcal{I}$ , which is focused solely to the different altruism levels of the provider agent. Therefore, we argue that our framework can be easily extended to a vast plethora of resource-exchange-based scenarios, such as the Dictator Game [4], the Favour Game [31], and possibly also to non-game theory games such as collaborative multiplayer games [32]. Moreover, its ability to detect group structures in real-time allows for the application of the GM framework within the Experience-Driven Procedural Content Generation framework [33] for the creation of adaptive multiplayer games [32].

Future work will pursue three concurrent paths: first, we will investigate the ability of our framework to detect the formation and evolution of group structures within dynamic societies; second, we will extend the field of application of our framework to interaction scenarios alternative to the Ultimatum Game, such as the Favour Game [31]; third, we will implement a multiplayer game which mechanics of interactions among the players will require some form of resource exchange/giving,

similarly to the Favour Game [31] and the Response Game [4].

In fact, we believe that collaborative games would enhance gaming experience features — e.g. agency, proximity and contextualisation of the interactions within an existing environment — this would not only increase the level of simulations of real-life scenarios [34], we believe it would also facilitate a proper *emergence* of group structures, as opposed to a more artificial instantiation of them observed in relevant game theory literature [4]. Moreover, the real-time capabilities of the GM framework could be combined with Experience-Driven Procedural Content Generation mechanisms [33] in order to, for instance, investigate the effects that the changes in the environment have on group dynamics.

### VIII. CONCLUSIONS

This paper presented an evolutionary computation approach for the detection of consolidated group structures, built on existing networks of reciprocal friendship, in complex artificial societies of agents interacting with each other by means of the Ultimatum, or Bargain Social Dilemma, Game. The framework proposed is capable of detecting the *true* group structures by solely observing the different levels of altruism a provider agent manifests depending on which other agent it interacts with. The evolutionary algorithm designed is able to successfully detect community structures in fully connected, directed and weighted networks. Such characteristics collectively can be found among the most complex network structures in the field of network analysis.

### ACKNOWLEDGMENTS.

This work has been supported, in part, by the FP7 ICT project SIREN (project no: 258453).

### REFERENCES

- [1] S. J. Lansing, "Complex Adaptive Systems," *Annual Review of Anthropology*, vol. 32, pp. 183–204, 2003.
- [2] R. Lewin, *Complexity: Life at the Edge of Chaos*. University of Chicago Press, 2000.
- [3] R. M. Dawes and D. M. Messick, "Social Dilemmas," *International Journal of Psychology*, vol. 2, no. 35, pp. 111–116, 2000.
- [4] Y. Chen and S. Li, "Group Identity and Social Preferences," *The American Economic Review*, pp. 431–457, 2009.
- [5] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*. Cambridge University Press, 1998, vol. 1, no. 1.
- [6] E. Fehr and U. Fischbacher, "Why Social Preferences Matter – the Impact of Non-selfish Motives on Competition, Cooperation and Incentives," *The Economic Journal*, vol. 112, no. 478, pp. 1–33, 2002.
- [7] R. Forsythe, "Fairness in Simple Bargaining Experiments," *Games and Economic Behavior*, vol. 6, no. 3, pp. 347–369, 1994.
- [8] E. A. Leicht and M. E. J. Newman, "Community Structure in Directed Networks," *Physical Review Letters*, vol. 100, no. 11, p. 118703, 2008.
- [9] C. Farmer and A. Fotheringham, "Network-based Functional Regions," *Environment and Planning A*, vol. 43, no. 11, pp. 2723–2741, 2011.
- [10] C. Grappiolo and G. N. Yannakakis, "Towards Detecting Group Identities in Complex Artificial Societies," in *Proc. of SAB*, 2012, pp. 421–430.
- [11] M. A. Nowak, C. E. Tarnita, and T. Antal, "Evolutionary Dynamics in Structured Populations," *Philosophical Transactions of the Royal Society*, vol. 365, no. 1537, pp. 19–30, 2010.
- [12] R. A. Hammond and R. Axelrod, "The Evolution of Ethnocentrism," *Journal of Conflict Resolution*, vol. 50, no. 6, pp. 926–936, 2006.
- [13] K. Lerman and A. Galstyan, "Automatically Modeling Group Behavior of Simple Agents," in *Proc. of AAMAS*, 2004.
- [14] V. Martinez, G. I. Simari, A. Sliva, and V. Subrahmanian, "CONVEX: Similarity-Based Algorithms for Forecasting Group Behavior," *IEEE Intelligent Systems*, vol. 23, pp. 51–57, 2008.
- [15] M. Janssen and E. Ostrom, "Empirically Based, Agent-based Models," *Ecology and Society*, vol. 11, no. 2, p. 37, 2006.
- [16] B. Xianyu, "Social preference, incomplete information, and the evolution of ultimatum game in the small world networks: An agent-based approach," *Journal of Artificial Societies and Social Simulation*, vol. 13, no. 2, p. 7, 2010.
- [17] J. Liu, H. Abbass, W. Zhong, and D. Green, "Local-Global Interaction and the Emergence of Scale-free Networks with Community Structures," *Artificial Life*, vol. 17, no. 4, pp. 263–279, 2011.
- [18] G. Palla, A. Barabasi, and T. Vicsek, "Quantifying Social Group Evolution," *Nature*, vol. 446, no. 7136, pp. 664–667, 2007.
- [19] N. Eagle, A. Pentland, and D. Lazer, "Inferring Friendship Network Structure by Using Mobile Phone Data," *Proc. of NAS*, vol. 106, no. 36, pp. 15 274–15 278, 2009.
- [20] C. Pizzuti, "GA-Net: A genetic Algorithm for Community Detection in Social Networks," *Parallel Problem Solving from Nature*, vol. 5199, pp. 1081–1090, 2008.
- [21] M. Tasgin, A. Herdagdelen, and H. Bingol, "Community Detection in Complex Networks Using Genetic Algorithms," *arXiv:0711.0491*, 2007.
- [22] J. Liu, W. Zhong, H. Abbass, and D. Green, "Separated and overlapping community detection in complex networks using multiobjective Evolutionary Algorithms," in *Proc. of CEC*, 2010, pp. 1–7.
- [23] A. Clauset, M. Newman, and C. Moore, "Finding Community Structure in Very Large Networks," *Physical Review E*, vol. 70, no. 6, p. 066111, 2004.
- [24] G. E. Bolton, E. Katok, and R. Zwick, "Dictator game giving: Rules of fairness versus acts of kindness," *International Journal of Game Theory*, vol. 27, pp. 269–299, 1998.
- [25] E. Fehr and K. M. Schmidt, "A Theory Of Fairness, Competition, And Cooperation," *The Quarterly Journal of Economics*, vol. 114, no. 3, pp. 817–868, 1999.
- [26] J. H. Kagel, C. Kim, and D. Moser, "Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs," *Games and Economic Behavior*, vol. 13, no. 1, pp. 100–110, 1996.
- [27] A. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [28] F. Marzo, B. J. Grosz, and A. Pfeffer, "Social preferences in Relational Contexts," in *Proc. of Collective Intentionality*, 2005.
- [29] T. Back, D. Fogel, and Z. Michalewicz, *Handbook of Evolutionary Computation*. IOP Publishing Ltd., 1997.
- [30] H. Kuhn, "The Hungarian Method for the Assignment Problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 2006.
- [31] M. Jackson, T. Rodriguez-Barraquer, and X. Tan, "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *The American Economic Review*, vol. 102, no. 5, pp. 1857–1897, 2012.
- [32] G. N. Yannakakis, J. Togelius, R. Khaled, A. Jhala, K. Karpouzis, A. Paiva, and A. Vasalou, "Siren: Towards Adaptive Serious Games for Teaching Conflict Resolution," in *Proc. of ECGBL*, 2010, pp. 412–417.
- [33] G. N. Yannakakis and J. Togelius, "Experience-Driven Procedural Content Generation," *IEEE Transactions on Affective Computing*, vol. 2, pp. 147–161, 2011.
- [34] G. W. Greenwood and D. Ashlock, "Evolutionary Games and the Study of Cooperation: Why Has So Little Progress Been Made?" in *Proc. of WCCI*, 2012, pp. 680–687.