Using Reinforcement Learning and Artificial Evolution for the Detection of Group Identities in Complex Adaptive Artificial Societies

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

ABSTRACT

We present a computational framework capable of inferring the existence of groups, built upon social networks of reciprocal friendship, in Complex Adaptive Artificial Societies (CAAS). Our modelling framework infers the group identities by following two steps: first, it aims to learn the ongoing levels of cooperation among the agents and, second, it applies evolutionary computation, based on the learned cooperation values, to partition the agents into groups. Experimental investigations, based on CAAS of agents who interact with each other by means of the Ultimatum Game, show that a cooperation learning phase, based on Reinforcement Learning, can provide highly promising results for minimising the mismatch between the existing and the inferred groups, for two different society sizes under investigation.

Categories and Subject Descriptors

I.2.m [Computing Methodologies]: Artificial Intelligence— Miscellaneous

General Terms

Algorithms

Keywords

Group Identity Detection; Artificial Societies; Ultimatum Game; Reinforcement Learning; Evolutionary Computation.

1. INTRODUCTION

When a population of socially-driven individuals are let to interact with each other in absence of any form of centralised control, global patterns arise, such as the *emergence* of friendship [13], culture [11], or other forms of group structures. The patterns have an influence on the behaviours of each individuals; these are manifested by the concept of group identity — i.e. the ability of the individuals to identify themselves as part of a group —, meaning that individuals belonging to the same group (in-group) will tend to be more collaborative, intended as reciprocally altruistic [9], than when they interact with individuals belonging to other groups (out-group) [4]. A straightforward, direct approach

Copyright is held by the author/owner(s).

GECCO'13 Companion, July 6–10, 2013, Amsterdam, The Netherlands. ACM 978-1-4503-1964-5/13/07.

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

for the detection of groups is represented by the direct querying of each individual. However, this might not always be possible, making thus the task of group detection/modelling problematic. Nevertheless, due to the aforementioned relationship existing between the collaborative nature of the interactions and the presence of group identities, alternative indirect approaches could be adopted. The overall aim of our research is to provide a computational framework of indirect inference — hereafter called Group Modelling (GM) framework — capable of detecting the presence of group structures, and assign relative group identities, in Complex Adaptive Society Systems. Whilst our ultimate goal is the application of our framework to virtual environments populated by human-controlled avatars, e.g. collaborative multiplaver games, the focus given in this paper is on Complex Adaptive Artificial Societies (CAAS) [8].

2. EXPERIMENTAL SETUP

The agents forming the CAAS manifest relational-based social preferences when interacting with each other by means of the Ultimatum (or Bargain) Social Dilemma Game (UG). Our choice to use UG for our research is motivated by its extensive application in many akin studies, among which we remark those which put an emphasis on social preferences [2, 8, 10]. The complex global pattern under investigation is the existence of group structures, based on the notion of reciprocal friendship, perceived by the adaptive agents of the society at any time throughout the experiments. The evolution of friendship is implemented as follows: at regular intervals of interaction episodes, a new friendship link is added to the society's social network, \mathcal{F} , by applying the well-established Barabási-Albert algorithm for the creation of Scale-free networks [1, 6]. As soon as a new edge is added to \mathcal{F} , the two involved agents are notified, their behaviours are affected (i.e. they start behaving more altruistically with each other) and their own social perception towards each other moves towards friendship. The true group identities are obtained by partitioning \mathcal{F} by means of the well-established Clauset-Newman-Moore algorithm [3]. The process ends when all the agents have at least one friend.

3. GROUP MODELLING FRAMEWORK

Our GM framework is composed of two interconnected modules: the first, called Cooperation Learning Module (CL), is in charge of learning the ongoing levels of cooperation existing among the artificial agents; the second module, called Group Identity Detection Module (GID), partitions the agents



Figure 1: 20-Agent Societies

into group structures, based on the learnt cooperation values, and assigns a group identity to each one of the identified structures.

The main focus of this paper is the investigation of appropriate and efficient techniques to be used for the CL module, being its computation the fundamental feature for the correct computation of the GID module. For the GID module, instead, we rely on a Evolutionary Algorithm (EA) following the principles introduced in our previous study on static societies [6].

The CL module aims to learn the Society's ongoing collaboration values by solely processing the *offers* of the agents, given the assumption that friendship influences altruism or cooperation [7, 10]. We consider three different learning rules for the CL module: (1) Reinforcement Learning's (RL) α -constant Monte Carlo (α -cMC) update rule [12] for nonstationary environments with Agent-Preference reward function (AP) [7]; (2) Temporal-Difference (TD) update rule [12] with Entropy-Agent-Preference reward function (HAP) [7]; (3) pheromone evaporation update rule used in Ant Colony Optimisation [5].

4. RESULTS AND DISCUSSION

In order to examine which of the CL three modules, combined with the evolutionary GID module, would better infer the existence of *emergent* group structures of reciprocal friendship we have conducted experiments on two scenarios: 20 and 50 agent societies. Each scenario was iterated for 30 experimental runs; the three GM frameworks were evaluated against the same experimental setup, in order to grasp a better, more generic understanding of their performance, and to limit random fluctuations introduced by the interaction classifiers and the stochastic nature of the EA. The performance measure we consider in this study is the normalised mismatch error [6] between the *true* community structures built upon \mathcal{F} and those detected by the GM framework.

The results gathered (see Figure 4, the grey vertical bar indicates the interaction episode beyond which no new reciprocal friendship links are perceived by the agents — consolidation of friendship) seem to suggest that the three techniques used for the CL module are equally good. Nonetheless α -cMC with AP reward function appears to perform generally better than the other two, especially with respect to the consolidation of the friendship structures and 50 agent societies, though statistical significance of such better performance is registered only occasionally throughout the episodes [8].

5. FUTURE WORK AND CONCLUSIONS

The rationale behind the success of α -cMC, and in part TD0, is their ability to perform the transition from the raw interactions to the two *in/out-group* classes. This was achieved via the AP and HAP reward functions [7]. These measure altruism as a more general property of the agents, which cannot be successfully estimated by isolating each single agent interaction [7]. The CL module provides an augmented representation of \mathcal{F} , by transforming the dichotomy friend-vs-non friend into a continuous one by means of the update rules. This opens to a plethora of improvement possibilities, two of which are proposed here: (1) a new module could be embedded after the CL module in order to refine the collaboration values so that the GID module would further improve its detection performance [6]; (2) GM could be enriched by a new interactive component, giving the possibility to the agents to self-report, at any time, their own perceptions of friendship, and thus directly affect the learnt collaboration values. Other future work will consider the further investigation of alternative techniques for the CL and GID modules, plus the application of the GM framework to collaborative multiplayer games.

6. **REFERENCES**

- A. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999.
- [2] Y. Chen and S. Li. Group Identity and Social Preferences. *The American Economic Review*, pages 431–457, 2009.
- [3] A. Clauset, M. Newman, and C. Moore. Finding Community Structure in Very Large Networks. *Physical Review E*, 70(6):066111, 2004.
- [4] R. M. Dawes and D. M. Messick. Social Dilemmas. International Journal of Psychology, 2(35):111–116, 2000.
- [5] M. Dorigo and K. Socha. An Introduction to Ant Colony Optimization. *Handbook of Approximation Algorithms and Metaheuristics*, pages 26–1, 2006.
- [6] C. Grappiolo, J. Togelius, and G. N. Yannakakis. Artificial Evolution for the Detection of Group Identities in Complex Artificial Societies. In *Proc. of IEEE ALife*, pages 126–133, 2013.
- [7] C. Grappiolo and G. N. Yannakakis. Towards Detecting Group Identities in Complex Artificial Societies. In *Proc. of SAB*, pages 421–430, 2012.
- [8] C. Grappiolo, J. Togelius, and G. N. Yannakakis. Interaction-based Group Identity Detection via Reinforcement Learning and Artificial Evolution. In *Proc. of EcoMass*, (to appear), 2013.
- [9] S. J. Lansing. Complex Adaptive Systems. Annual Review of Anthropology, 32:183–204, 2003.
- [10] F. Marzo, B. J. Grosz, and A. Pfeffer. Social preferences in Relational Contexts. In *Proc. of Collective Intentionality*, 2005.
- [11] H. Situngkir. On Selfish Memes: Culture as Complex Adaptive System. *Journal of Social Complexity*, 2(1):20–32, 2004.
- [12] R. Sutton and A. Barto. Reinforcement Learning: An Introduction. Cambridge University Press, 1998.
- [13] E. Wolf. Kinship, Friendship, and Patron-client Relations in Complex Societies. *The Social Anthropology of Complex Societies*, pages 1–22, 1966.