Norm Establishment Constrained by Limited Resources (Extended Abstract)

Samhar Mahmoud, Simon Miles, Adel Taweel, Brendan Delaney, Michael Luck
King's College London, London, UK
{samhar.mahmoud, simon.miles, adel.taweel, brendan.delaney, michael.luck}@kcl.ac.uk

ABSTRACT

Peer punishment has been an important instrument in enabling social norms to emerge. However, it is usually assumed that unlimited resources are available to agents to cope with the resulting enforcement costs. In this paper, we use a modified version of the metanorm model [1] to investigate this, and show that it allows norm emergence only in limited cases under bounded resources.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

Keywords

Metanorm, Limited Enforcement Cost

1. INTRODUCTION

Many have been concerned with the development of mechanisms to ensure the emergence of social norms. In particular, researchers from many scientific areas have considered *punishment* as a key motivating element for norms to be established (e.g., [2, 7]). Here, punishment is a monetary incentive, typically incurring an enforcement cost for the punisher, but bringing a potential benefit to the population as a whole. Work that uses punishment as a means for social norms to emerge has assumed that agents applying such punishment have unlimited resources, allowing them to bear the enforcement cost. This assumption is significant in real world settings in which resources are limited and require more careful exploitation. For example, sensors in wireless networks have limited energy and thus need to optimise their use of it.

In response, this paper seeks to address such limitations by integrating the constraint of limited resources within the metanorm model originally proposed by Axelrod [1] and adapted by Mahmoud et al. [5, 6]. The metanorm model has been shown to be capable of regulating distributed computational systems under various settings. This paper investigates the limitations of the static punishment mechanism of this model under limited resource constraints, and makes some suggestions to resolve such issues.

2. METANORM MODEL

In Axelrod's metanorm model [1], a population of agents play a game in which each agent has to decide between cooperation and defection. Defection brings a reward for the defecting agent called *temptation*, and a penalty to all other agents called *hurt*, but each defector risks being observed by the other agents and punished as a result. These other agents thus decide whether to punish agents that were observed defecting, with a low penalty for the punisher known as an *enforcement cost*, and a high penalty for the punished agent known as a *punishment cost*. Agents that do not punish those observed defecting risk being observed themselves, and potentially incur metapunishment. Thus, each agent decides whether to metapunish agents observed to spare defecting agents. The strategy of each agent in determining whether to defect and whether to punish others is based on two different attributes, *boldness (B)* (encouraging agents to defect) and *vengefulness (V)* (encouraging them to punish and metapunish others), which are distinct for each agent. The agent population evolves through a number of iterations, with a mechanism whereby successful behaviour (as measured by the scoring system) tends to be replicated and unsuccessful behaviour tends to be discarded. However, a major problem with Axelrod's model is due to the evolutionary approach adopted (as identified in [3]). In consequence, this original approach was replaced with a reinforcement learning algorithm that limits accessibility to global information, and instead allows agents to learn from their own experience [6]. Moreover, in order to capture a key feature of computational systems with structural relations between their components such as on-line virtual communities, Axelrod's classic model has been adapted by introducing a topological structure [5] that determines observability among agents, so that an agent's neighbours are the only witnesses of its interactions.

3. LIMITING RESOURCES

As mentioned above, in real world settings, agents usually have limited resources for enforcement. Thus, once an agent i is in a position to apply punishment to a violator agent j, the punishment can only take place if sufficient resources exist to supplement the enforcement cost that can result from the punishment. First, i needs to be able to identify the amount of resources available for punishment on all agents. In addition, i needs to estimate the resources (enforcement cost) required to apply this particular punishment. Having verified that a punishment is possible, i can then punish j, which results in the resources of i decreasing by the relevant enforcement cost. This affects future decisions that i can make with regard to punishment and metapunishment. Now, since resources are in general renewable, we consider them to be limited only over a particular period of time. With regard to the metanorm model, this means that once a particular agent expends all of its resources, it cannot apply any form of punishment or metapunishment until the restriction period has passed, after which resources are renewed. Since the model is round based, we assume that resources

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

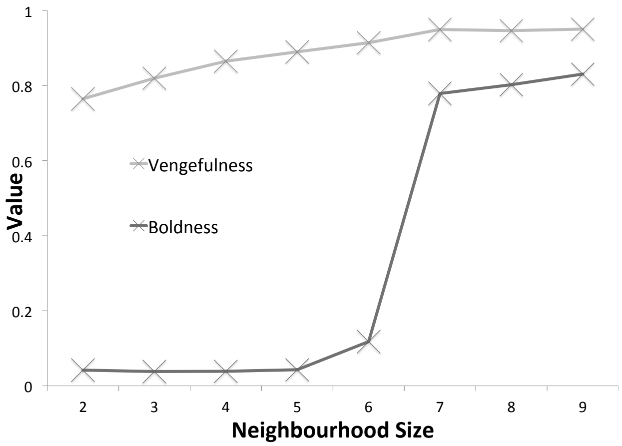

Figure 1: impact of limited resources of 12 units on final B and V and various neighbourhood size

are renewed every round for each agent. In what follows, we show the effect of this new restriction on norm establishment using the metanorm model.

4. EXPERIMENTAL EVALUATION

The most desirable results are those with high vengefulness and low boldness, which are referred to as *norm establishment*. This is because low boldness means that agents defect rarely, and high vengefulness means that agents are generally willing to punish an agent that defects. Other results involving midrange or high levels of boldness are referred to as *norm collapse*, since they involve a high number of defections. Topologies are also an important component, and have been shown to have different effects on norm establishment [5]. For the purposes of this paper, we use a simple lattice topology for which norm establishment has been achieved previously.

The interaction model of agents involves the following sequence of actions. For every defection opportunity that an agent i has, all of i's neighbours have a chance to punish i. If a neighbour j decides to spare i from punishment, then all of j's neighbours have the chance of metapunishing j. Assuming that every agent has 4 distinct neighbours, this means that for every single defection, 4 punishment decisions need to be taken, and if all these punishment decisions result in sparing the defector, $4 \times 4 = 16$ metapunishments can arise. Based on this, it is clear that agents will invest most of their resources on punishment and metapunishment as a result of the outcome of the first few defections, with scarce resources left to regulate the behaviour of the remaining agents. This explains the results obtained from an experiment where each agent is provided with 12 resource units that are renewed every round, and each agent has a neighbourhood size of 8. The results of this experiment show that the model fails to establish the norm with the average boldness of agents remaining very high, and reflecting a very high rate of defection. The surprise here is that the average vengefulness is also high, which is due to two factors. First, for the first few occurrences of defection, sufficient resources remain available for metapunishment. Second, resources run out quickly, so no more enforcement costs are paid by agents to cause vengefulness to drop. This last factor also explains the high boldness, with insufficient punishment taking place to deter defecting agents by outweighing the temptation gained.

The above analysis suggests that the number of neighbours of each agent plays a major role in the obtained results. Therefore, a further set of experiments were conducted in which neighbourhood size was varied. The results reported in Figure 1 are for a lattice topology with neighbourhood size varying between 2 and 9, and a limited resource of 12 units. Each point on the graph represents an average of $1,000$ runs with a particular neighbourhood size. We can see that with limited resources of 12 units, we establish the norm up to a neighbourhood size of 6, but not after that. A similar outcome is found using other amounts of limited resources. For example, with limited resources of 6 units, norm establishment is observed up to a neighbourhood size of 3, and up to a neighbourhood size of 4 with limited resources of 8 units.

5. DISCUSSION

This paper has studied the effect of integrating a limited resource constraint within the well established metanorm model. The experimental results show that the static punishment mechanism of the metanorm model fails to establish the norm, with a clear relationship between the available resources and the neighbourhood size of the network topology. This is mainly due to the lack of consideration of available resources when a punishment is applied, which suggests that an adaptive punishment mechanism may be better suited to achieving norm establishment. Such adaptive punishment has been proposed in previous work [4], in which agents calculate an appropriate punishment to deter a defector from future violations, based on past behaviour of the defector. Initial experimental results show that norm establishment is improved with such a mechanism. However, norm establishment is still limited, since the proposed adaptive punishment mechanism does not consider the amount of resources available to the agent when applying the punishment.

REFERENCES

[1] R. Axelrod. An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111, 1986.

[2] D. Helbing, A. Szolnoki, M. Perc, and G. Szabo. Punish, but not too hard: how costly punishment spreads in the spatial public goods game. *New Journal of Physics*, 12(8), 2010.

[3] S. Mahmoud, N. Griffiths, J. Keppens, and M. Luck. An analysis of norm emergence in Axelrod's model. In *NorMAS'10: Proceedings of the Fifth International Workshop on Normative Multi-Agent Systems*. AISB, 2010.

[4] S. Mahmoud, J. Keppens, N. Griffiths, and M. Luck. Efficient norm emergence through experiential dynamic punishment. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 576–581. IOS Press, 2012.

[5] S. Mahmoud, J. Keppens, M. Luck, and N. Griffiths. Establishing norms for network topologies. In *Coordination, Organizations, Institutions and Norms in Agent Systems IV*, volume 7254 of *Lecture Notes in Computer Science*, pages 203–220. Springer Berlin Heidelberg, 2012.

[6] S. Mahmoud, J. Keppens, M. Luck, and N. Griffiths. Overcoming omniscience for norm emergence in axelrod's metanorm model. In *Coordination, Organizations, Institutions and Norms in Agent Systems IV*, volume 7254 of *Lecture Notes in Computer Science*, pages 186–202. Springer Berlin Heidelberg, 2012.

[7] D. Villatoro, G. Andrighetto, J. Sabater-Mir, and R. Conte. Dynamic sanctioning for robust and cost-efficient norm compliance. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, pages 414–419. IJCAI/AAAI, 2011.