

Social Norm Emergence in Virtual Agent Societies

Bastin Tony Roy Savarimuthu, Maryam Purvis, Martin Purvis,
and Stephen Cranefield

Department of Information Science, University of Otago, Dunedin, P.O. Box 56,
Dunedin, New Zealand

{tonyr,tehrany,mpurvis,scanefield}@infoscience.otago.ac.nz

Abstract. The advent of virtual environments such as SecondLife call for a distributed approach for norm emergence and spreading. In open virtual environments, monitoring various interacting agents (avatars), using a centralized authority might be computationally expensive. The number of possible states and actions of an agent could be huge. An approach for sustaining order and smoother functioning of these environments can be facilitated through norms. Agents can generate norms based on interactions. In particular, those social norms that incur certain cost to an individual agent but benefit the whole society are more interesting than those benefit both the agent and the society. The problem is that the selfish agents might not be willing to share the norm adherence cost. In this work, we experiment with notion proposed by Axelrod that social norms are best at preventing small defections where the cost of enforcement is low. We also study how common knowledge can be used to facilitate the overall benefit of the society. We believe our work can be used to facilitate norm emergence in virtual online societies.

1 Introduction

Norms are expectations of an agent about the behaviour of other agents in the society. The human society follows norms such as tipping in restaurants and exchange of gifts at Christmas. Norms are of interest to researchers because they help to improve the predictability of the society. They also reduce the computations required by an agent to make a decision. Norms have been of interest in different areas of research such as Sociology, Economics, Psychology and Computer science [1]. Norms have been shown to facilitate co-ordination and co-operation among the members of the society [2, 3]. Some of the well established norms may become laws.

While the discussion on how norms emerge and spread remains a research issue among scientists in Sociology, the advent of new ways of human interactions proxied through software agents in virtual 3D worlds such as SecondLife [4] have created interest among researchers in MAS to work on the applicability of the concept of norms in these digital societies.

We believe that software agents that operate autonomously or on behalf of human users in these virtual worlds cannot be effectively monitored and controlled

through centralized policing mechanisms. The explosion of possible action states for an agent in an open environment is huge. It would be computationally expensive and taxing for a centralized monitor to enforce behavioural regularities to ensure smoother functioning of these systems. We believe that an alternative approach based on norms could be effectively used in such scenarios where norms can be derived and built using a bottom-up approach through interactions between the agents. Our objective in this paper is to experiment with the bottom-up approach where we observe whether a norm against littering a park spreads in an agent society.

The paper is organized as follows. Section 2 provides a brief overview on the work on norms with a focus on the norm emergence addressed in the field of multi-agent systems. The experimental set up used for conducting simulations on norm emergence is described in Section 3. The results are presented in Section 4. The discussion and future work are presented in Section 5.

2 Related Work on Norms

Due to multi-disciplinary interest in norms, several definitions for norms exist. Habermas [5], a renowned sociologist, identified norm regulated actions as one of the four action patterns in human behaviour. A norm to him means *fulfilling a generalized expectation of behaviour*, which is a widely accepted definition for social norms. When members of a society violate the societal norms, they may be punished (and even ostracized in some cases). Many social scientists have studied why norms are adhered. Some of the reasons for norm adherence include a) fear of authority b) rational appeal of the norms and c) feelings such as shame, embarrassment and guilt that arise because of non-adherence.

Social scientists have categorized norms into several categories [1]. One such categorization is based on the the cost-benefit analysis of a norm from a perspective of an individual agent and how that relates to the society as the whole, proposed by Horne [6]. There are four categories of norms according to Horne. There are norms which benefit both the individual agents and the society and some norms incur cost to both the agent and the society. There are some norms that benefit the agent but cost the society. Some norms cost the agent but are beneficial to the society. The last category of norms where the whole society is benefitted when an agent incurs a cost, is more interesting than the others. For these norms to be established there should be agents in the society that will help in the enforcement process. In the realm of digital societies such as that of SecondLife, it is important that there are these distributed enforcing agents that are helpful in regulating social order which is beneficial to the society.

2.1 Normative Multi-agent Systems

Norms have been of interest to multi-agent system researchers for over a decade [3, 7, 8]. Norms in multi-agent systems are treated as constraints on behaviour, goals to be achieved or as obligations [9]. The research on norms in sociology and multi-agent systems complement each other. Researchers in multi-agent systems

have borrowed ideas from sociology such as speech act theory and autonomy to model software agents. Sociologists on the other hand have used multi-agent systems for testing their theories through simulations.

The definition of normative multi-agent systems as described by the researchers involved in Normas 2008 workshop is as follows [10]: *A normative multi-agent system is a multi-agent system organized by means of mechanism to represent, communicate, distribute, detect, create, modify and enforce norms, and mechanism to deliberate about norms and detect norm violation and fulfillment.*

While some aspects of normative multi-agent systems such as normative system architectures, norm representations, norm adherence and violation detections [11, 12, 13] have received a lot of interest, areas such as norm creation and modification have not been explored in-depth. Our work in this paper borders the area of norm distribution (i.e. norm spreading) and the detection of the norm that has been created (through emergence).

2.2 Norm Spreading and Emergence

The concepts of norm spreading and norm emergence are related. A norm emerges in a society as a result of norm spreading mechanism. According to Boyd and Richerson [14], there are three ways by which a social norm can be propagated from one member of the society to another. They are a) Vertical transmission (from parents to offspring), b) Oblique transmission (from a leader of a society to the followers) and c) Horizontal transmission (from peer to peer interactions). Norm propagation is achieved by spreading and internalization of norms. Boman and Verhagen [7, 15] have used the concept of normative advice (advice from the leader of a society) as one of the mechanisms for spreading and internalizing norms in an agent society. The work done by Savarimuthu et al. [16] uses a distributed approach for normative advice based on the notion of leadership. Another recent development is the consideration of the role of network topologies on norm emergence [17, 18]. Sen et al. [19] have experimented with the emergence of traffic norm using social learning. In his well known work, Axelrod [2] has shown the role of metanorms to be effective in norm emergence. He also discusses several other approaches that might be useful for norm establishment which include the role of power, reputation, internalization and punishment. The contribution of this paper to this area are two fold. Firstly, we investigate Axelrod's statement that social norms are better suited for preventing smaller defections when the enforcement costs are low using social simulations in the context of norm emergence. Secondly, we introduce the notion of common knowledge that can help sustain norms in agent societies.

3 Experimental Setup and Parameters

Multi-Agent Based Simulation (MABS) is an inter-disciplinary research area which brings together researchers within the Agent-Based Social Simulation community (ABSS) [20] and the Multiagent Systems community (MAS). MABS

researchers use simulation as a mechanism to experiment with and validate new hypotheses (both social and computational theories). Simulations are used as tools that provide explanations for behaviours exhibited by complex systems. Simulation tools help researchers to investigate the implications of a strategies adopted by participating agents by running simulations starting from different initial conditions [21]. Adopting this approach, we have set up a simulation environment with multiple agents that interact in a social context.

We model agents in our virtual society as particles moving in a 2 dimensional space of linear size L . This virtual environment can be considered as a communal region such as a park. The agents explore and enjoy the park by moving around. Collisions of these particles in the virtual space represent interactions between agents in a social space as shown in figure 1. Each collision corresponds to two agents observing each other's actions. When two agents interact (when they meet each other within certain area of the park) they can observe each other performing one of the two actions, Litter (L) or Not Litter (NL), i.e. keep the environment clean. The payoff matrix that corresponds to the littering scenario is given in table 1.

An agent in our society starts with a score (s) of 100. When an agent litters, it gets a payoff of 0.5 while the cost associated with non littering is -0.5. These are the payoffs to an individual agent. When an agent litters, it pollutes the area




Fig. 1. Simulation of the park scenario (100 agents)

Table 1. Payoff matrix

	L	NL
L	0.5, 0.5	0.5,-0.5
NL	-0.5, 0.5	-0.5,-0.5

that belongs to the commons. So, this can be considered defecting the entire society. This impacts the productivity of the society. Productivity refers to the benefits that the park users will receive in using the park. This has been modeled using a variable *prod* that holds values from 0 to 1. The value of 1 represents the clean state of the park while 0 represents a littered, unusable park. Initially the productivity of the park is set to 1. When an agent litters, the productivity of the park goes down. For every X littering actions the productivity value is reduced by 0.01. The value of X was set to 10 in our system. The final payoff value is the sum of the individual payoff and the productivity of the system.

Let us now assume that the society does not have a law against littering and hence there is no centralized policing mechanism. In this scenario, any agent that believes that there should be no littering in the society, might choose to punish the other agent whom it observes littering. A punishment cost (P_{cost}) is incurred by the non-litterer when punishing a littering agent. Every agent in the society is initialized with an autonomy value from 0 to 9 based on a uniform distribution. Autonomy refers to the stubbornness of the agent. This value governs the number of punishments required by an agent to move from L to NL (change of strategies).

Another parameter that we have defined in the system is the minimum Survival score (S_{score}). When an agent's score, s goes below S_{score} or the productivity of the system goes below 0.5 ($prod < 0.5$), it changes its strategy (moves from L to NL). S_{score} is set to 50 in our experiments.

4 Experiments and Results

4.1 Role of Punishments with Low Enforcement Cost

In the first experiment there are 100 agents, 50 agents of type L and 50 of type NL. In every society there will be certain percentage of agents that are vengeful enough to punish another agent when they observe certain behaviour that they consider to be inappropriate. Let us assume that there are certain percentage of non-littering agents that are punishers ($p=0.05, 0.10$ and 0.25). P_{cost} is kept low (0.01) in these experiments.

In each iteration two agents are randomly chosen to interact. We conducted experiments over 6000 iterations. At the end of the simulation, we observe whether a littering or non-littering norm emerges. In our experiments we consider a norm to have emerged if all the agents are either of type L or NL (100% norm emergence). In other works, the value for norm emergence has varied from 70% to 100%.




Fig. 2. Emergence of norm against littering in an agent society

Figure 2 shows 6 different lines. For each p value, there are two lines, one representing the number of litterers and the other representing the number of non-litterers. As the number of litterers decreases, the number of non-litterers increases (and vice versa). For this reason, these lines for a given p value, are the mirror images of each other. It is of interest to observe whether the littering or non-littering group reaches the value of 100. All the 6 lines start from a value of 50. Note that non-litterers are represented using hollow symbols while the litterers are shown using solid symbols.

It can be observed from figure 2 that as the number of punishers increases, the norm against littering is established. When the values of p are 0.10 and 0.25, the system converges to a norm against littering. When the value of p is 0.05, there aren't enough punishers in the system. Hence, the productivity of the society drops gradually and falls below the minimum level which results in the establishment of littering norm. Once there are adequate number of punishers, the cost of punishment is spread across the non-littering punishers, hence their individual scores do not reach the minimum threshold and they are successful in converting the litterers to become non-litterers.

Wikipedia is a good example of a system where more number of enforcers are needed to maintain the quality of the articles at a reasonable level. Members of Wikipedia have been successful in establishing a norm of collaboration using a peer to peer mechanism based on careful scrutiny of the content.

Figure 3 shows three different productivity lines that correspond to the results reported in figure 2. It can be observed that when there were adequate number of punishers $p = 0.1, 0.25$, the productivity of the society gradually improved

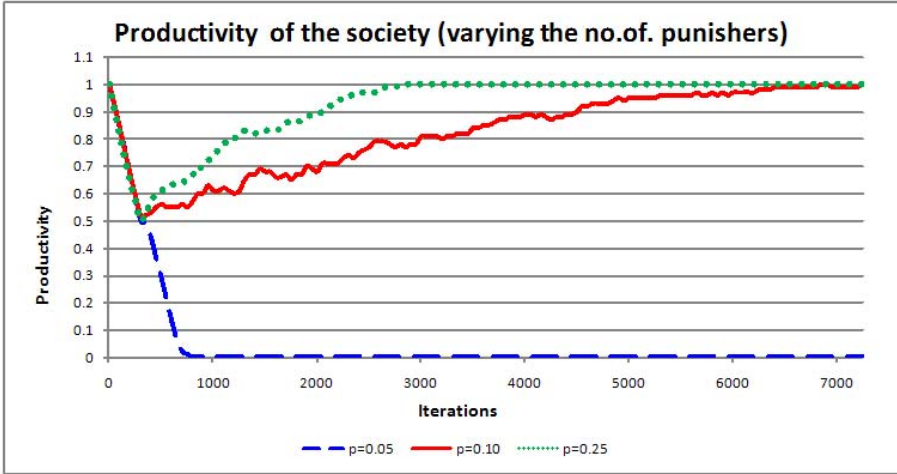


Fig. 3. Productivity of the society

and was sustained ($prod = 1$) in the end. When the litterers took over ($p = 0.05$) due to lack of enough punishers, the value of productivity plummeted to 0.

So, the important characteristics that governs this change are the autonomy of the individual agents (littering agents), the minimum score of productivity and the minimum threshold for survival (S_{score}). If a society has large autonomy values and high threshold for survival or productivity, the system will end up with litterers.

This type results can be observed in many social interactions. For example, when you go to a restaurant, if you were the first ones, you might be polite and keep your voice low when interacting with your friends. As the restaurant becomes crowded, you might observe the noise levels rising. As the noise level increases, there is no incentive for you to keep your voice down. Moreover, you might be forced to speak out loud as that is the only way you might be heard by others. This case is similar to the litterers becoming non-litterers after certain threshold is surpassed.

4.2 Role of Punishments with High Enforcement Cost

This experiment shows that when the enforcement cost increases the system drives all the agents to litter. This experiment shows the behaviour of the system for three different values of punishment cost ($P_{cost} = 0.1, 1, 10$). There are 25% punishers in a society for all the three experiments. Again, note that non-litterers are represented using hollow symbols while the litterers are shown using solid symbols.

It can be observed from figure 4 that lower enforcement costs ($P_{cost} = 0.1, 1$) resulted in a non-littering norm while the higher cost ($P_{cost} = 10$) resulted in littering norm. When P_{cost} was set to 10 the society initially had more number of non-litterers (till iteration 1000). Soon, the survival score (S_{score}) of punishers




Fig. 4. Emergence of norm against littering in an agent society ($p=0.25$)

fell below the minimum threshold due to high punishment cost. The non-littering punishers then became litterers and hence the littering norm was established in the society.

From this experiment it can be inferred that social norms can successfully be established and sustained against smaller defections when the costs of enforcements are low. But for norms that require larger costs of punishment (e.g. honour killing), social norms might not be very useful. In those cases, institutionalized mechanisms such as laws would be best suited.

4.3 Conditional Punishment Based on Common Knowledge

In human societies we tend to gather information about the state of the world through certain common knowledge sources such as newspapers, television, radio and even from some influential, well connected people. At any point of time, not everyone in the society might know about the current state of an issue or a problem. But, once some information is available through the common knowledge sources, it can be assumed that there would be an increase in the awareness of the situation in the society. Chwe [22] describes the role of common knowledge in solving coordination problems. In this section we describe our experiment on how common knowledge helps with in the context of social littering.

Let us assume that a common knowledge source is available (e.g. a newspaper). This common knowledge source periodically informs the agents about the state of the park. The agents in the society can choose to look at the information available from the knowledge source periodically. Based on the information available, the agents can choose to react. For example, whenever the park's




Fig. 5. Comparison of emergence of norm against littering - with and without Common Knowledge (CK)

productivity is less than certain value ($prod < 0.75$ in our experiment), the non littering agents can choose to punish. For example, say there were only 5% of the non litterers were punishers originally. After the information is known to all the other non litterers (remaining 45% of non-litterers), can choose to punish the litterers based on a conditional probability which is their vengefulness value. Each non-littering punisher agent has a vengefulness value (V) which is similar to the autonomy value and this is initialized at the start of the experiments. A non-littering punisher with vengefulness value of 8 will punish a litterer 8 out of 10 times.

Figure 5 shows a comparison of punishment mechanisms with and without the use of common knowledge keeping $P_{cost}=0.1$ as a constant. The figure shows four lines that correspond to the establishment or fading of the non-littering norm. We have omitted the lines that show the trend lines of the littering norm for the sake of clarity of the diagram.

When $p=0.05$, the punishment mechanism that makes use of common knowledge results in a non-littering norm (hollow triangles) while the mechanism that does not use common knowledge results in a littering norm (solid triangles). When $p=0.10$, the punishment mechanism that uses common knowledge (hollow squares) converges faster than the one that does not use it (solid squares). So, it can be inferred that the availability of common knowledge has increased the rate of establishment of a norm against littering in one case ($p=0.10$) and has resulted in the emergence of a new norm in another ($p=0.05$).

5 Discussion and Future Work

The results presented in this paper demonstrate that norms can be established through a bottom-up process based on a distributed, peer to peer punishment mechanism. The results obtained in our experiments are in agreement with Axelrod's statement that the norms are best in preventing smaller defections when the cost of enforcement is low. Also, it was shown that common knowledge can be used as a mechanism for improving norm establishment when used in conjunction with the punishment mechanism.

We agree that our results are preliminary. However, this work is aimed towards experimenting with mechanisms that might be suitable for generating norms in a bottom-up approach than a prescriptive top-down approach. In particular, our work is relevant for virtual online societies where behavioural norms should be derived by the agents themselves rather than adhering to an enforced law.

We are currently extending our simulation scenario to include agents of different personality types. Another extension to our system is to experiment with the emergence of different norms among different sub-groups within an agent society. To achieve that we need an application domain that has more states than a simple co-ordination game. Another important extension is to test the model on network topologies as agents evolve norms based on the influence from agents that they are connected to. The concept of distributed norm emergence is applicable to many applications in agent societies (e.g. buyer-seller scenarios in supply chain management, file sharing). A fertile ground for the study and experimentation of new mechanisms for norm emergence are the social networking applications.

Furthermore, norm emergence in digital societies can be studied using real data collected from social networking sites. New mechanisms for norm emergence can be experimented and applied on services such as file sharing on the Internet. Further investigations on norm emergence mechanisms are important because we believe that the future of the digital world will be played in digital environments where the software agents that act on behalf of human users might want to emerge norms rather than accepting an enforced norm.

References

1. Elster, J.: Social norms and economic theory. *The Journal of Economic Perspectives* 3(4), 99–117 (1989)
2. Axelrod, R.: An evolutionary approach to norms. *The American Political Science Review* 80(4), 1095–1111 (1986)
3. Shoham, Y., Tennenholtz, M.: On social laws for artificial agent societies: Off-line design. *Artificial Intelligence* 73(1-2), 231–252 (1995)
4. Second life, <http://secondlife.com/>
5. Habermas, J.: *The Theory of Communicative Action: Reason and the Rationalization of Society*, vol. 1. Beacon Press (1985)
6. Horne, C.: Sociological perspectives on the emergence of norms. In: Hechter, M., Opp, K.D. (eds.) *Social Norms*, pp. 3–34 (2001)

7. Boman, M.: Norms in artificial decision making. *Artificial Intelligence and Law* 7(1), 17–35 (1999)
8. Conte, R., Falcone, R., Sartor, G.: Agents and norms: How to fill the gap? *Artificial Intelligence and Law* 7(1), 1–15 (1999)
9. Castelfranchi, C., Conte, R.: *Cognitive and social action*. UCL Press, London (1995)
10. Boella, G., Torre, L., Verhagen, H.: Introduction to the special issue on normative multiagent systems. *Autonomous Agents and Multi-Agent Systems* 17(1), 1–10 (2008)
11. López y López, F., Márquez, A.A.: An architecture for autonomous normative agents. In: *Fifth Mexican International Conference in Computer Science (ENC 2004)*, pp. 96–103. IEEE Computer Society, Los Alamitos (2004)
12. García-Camino, A., Rodríguez-Aguilar, J.A., Sierra, C., Vasconcelos, W.: Norm-oriented programming of electronic institutions. In: *Proceedings of The Fifth International Joint Conference on Autonomous Agents and Multi Agent Systems, AAMAS*, pp. 670–672. ACM Press, New York (2006)
13. Boella, G., van der Torre, L.: An architecture of a normative system: counts-as conditionals, obligations and permissions. In: *Proceedings of The Fifth International Joint Conference on Autonomous Agents and Multi Agent Systems, AAMAS*, pp. 229–231. ACM Press, New York (2006)
14. Boyd, R., Richerson, P.J.: *Culture and the evolutionary process*. University of Chicago Press, Chicago (1985)
15. Verhagen, H.: *Norm Autonomous Agents*. PhD thesis, Department of Computer Science, Stockholm University (2000)
16. Savarimuthu, B.T.R., Purvis, M., Cranefield, S., Purvis, M.: Mechanisms for norm emergence in multi-agent societies. In: *Sixth International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2007)*, pp. 1097–1099 (2007)
17. Pujol, J.M.: *Structure in Artificial Societies*. PhD thesis, Software Department, Universitat Politècnica de Catalunya (2006)
18. Savarimuthu, B.T.R., Cranefield, S., Purvis, M., Purvis, M.K.: Role Model Based Mechanism for Norm Emergence in Artificial Agent Societies. In: *Sichman, J.S., Padget, J., Ossowski, S., Noriega, P. (eds.) COIN 2007. LNCS*, vol. 4870, pp. 203–217. Springer, Heidelberg (2008)
19. Sen, S., Airiau, S.: Emergence of norms through social learning. In: *Proceedings of Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, Hyderabad, India, pp. 1507–1512. MIT Press, Cambridge (2006)
20. Davidsson, P.: Agent based social simulation: a computer science view. *Journal of Artificial Societies and Social Simulation* 5 (2002)
21. López-Sánchez, M., Noria, X., Rodríguez, J.A., Gilbert, N.: Multi-agent based simulation of news digital markets. *IJCSA* 2(1), 7–14 (2005)
22. Chwe, M.: *Rational Ritual: Culture, Coordination and Common Knowledge*. Princeton University Press, Princeton (2001)