Intelligence Integration in Distributed Knowledge Management

Dariusz Król
Wroclaw University of Technology, Poland

Ngoc Thanh Nguyen
Wroclaw University of Technology, Poland

INFORMATION SCIENCE REFERENCE
Hershey · New York

Director of Editorial Content: Kristin Klinger
Managing Development Editor: Kristin M. Roth
Senior Managing Editor: Jennifer Neidig
Managing Editor: Jamie Snavely
Assistant Managing Editor: Carole Coulson
Copy Editor: Lanette Ehrhardt
Typesetter: Jeff Ash
Cover Design: Lisa Tosheff
Printed at: Yurchak Printing Inc.

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue, Suite 200
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

and in the United Kingdom by
Information Science Reference (an imprint of IGI Global)
3 Henrietta Street
Covent Garden
London WC2E 8LU
Tel: 44 20 7240 0856
Fax: 44 20 7379 0609
Web site: http://www.eurospanbookstore.com

Library of Congress Cataloging-in-Publication Data

Intelligence integration in distributed knowledge management / Dariusz Krol and Ngoc Thanh Nguyen, editors.

p. cm.

Includes bibliographical references and index.

Summary: "This book covers a broad range of intelligence integration approaches in distributed knowledge systems, from Web-based systems through multi-agent and grid systems, ontology management to fuzzy approaches"--Provided by publisher.

ISBN 978-1-59904-576-4 (hardcover) -- ISBN 978-1-59904-578-8 (ebook)

1. Expert systems (Computer science) 2. Intelligent agents (Computer software) 3. Electronic data processing--Distributed processing. I. Krol, Dariusz. II. Nguyên, Ngoc Thanh.

QA76.76.E95I53475 2009

006.3--dc22

2008016377

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Chapter XII
Norm Emergence in Multi-Agent Societies

Bastin Tony Roy Savarimuthu
University of Otago, New Zealand

Maryam Purvis
University of Otago, New Zealand

Stephen Cranefield
University of Otago, New Zealand

ABSTRACT

Norms are shared expectations of behaviours that exist in human societies. Norms help societies by increasing the predictability of individual behaviours and by improving cooperation and collaboration among members. Norms have been of interest to multi-agent system researchers, as software agents intend to follow certain norms. But, owing to their autonomy, agents sometimes violate norms, which needs monitoring. In order to build robust MAS that are norm compliant and systems that evolve and adapt norms dynamically, the study of norms is crucial. Our objective in this chapter is to propose a mechanism for norm emergence in artificial agent societies and provide experimental results. We also study the role of autonomy and visibility threshold of an agent in the context of norm emergence.

INTRODUCTION

Norms are behaviours that are expected by the members of a particular society. These expected behaviours are common in human societies and sometimes even in animal societies (Clutton-Brock & Parker, 1995). The human society follows norms such as tipping in restaurants, exchange of gifts at Christmas, dinner table etiquette and driving vehicles on the left or right hand side of the road. Some of the well-established norms may become laws. The norms are of interest to researchers because they help to improve the predictability of the society. Norm adherence enhances coordination and cooperation among the members of the society (Axelrod, 1986; Shoham & Tennenholtz, 1995). Norms have been of interest in different areas of research such as sociology, economics, psychology and computer science (Elster, 1989). Sociologists and economists are divided on their view of norms based on the theories of *homo economicus* and *homo sociologicus* (Elster, 1989). Sociologists consider that the norms are always used for the overall benefit of the society. Economists, on the other hand, state that the norms exist because they cater to the self-interest of every member of the society and each member is thought to be rational (Gintis, 2003). A more integrated view of norms from sociology and economics point of view is provided by Conte and Castelfranchi (1999). Applying social theories in multi-agents is synergetic, as agents are modeled using some of the social concepts such as autonomy and speech act theory. Both disciplines complement each other as agents serve as a platform to design, test and validate social theories. Some researchers (Boman, 1999; Verhagen, 2000, 2001) have undertaken agent-based simulations of social theories. Even though researchers in different fields have been trying to answer questions such as why agents follow certain norms and the implications of not following these norms, there has been limited work on mechanisms that propose the emergence of these norms. In this chapter, we explain a mechanism for norm emergence and discuss the role of autonomy and visibility threshold of an agent in an agent society.

BACKGROUND

In this section, we describe different types of norms and the treatment of norms in multi-agent systems. We also describe the work related to norm emergence.

Types of Norms

Due to multidisciplinary interest in norms, several definitions for norms exist. Habermas (1985), one of the renowned sociologists, 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. Researchers have divided norms into different categories. Tuomela (1995) has categorized norms into the following categories.

- r-norms (rule norms)
- s-norms (social norms)
- m-norms (moral norms)
- p-norms (prudential norms)

Rule norms are imposed by an authority based on an agreement between the members (e.g., one has to pay taxes). Social norms apply to large groups such as a whole society (e.g., one should not litter). Moral norms appeal to one's conscience (e.g., one should not steal or accept bribe). Prudential norms are based on rationality (e.g., one ought to maximize one's expected utility). When members of a society violate the societal norms, they may be punished. Many social scientists have studied why norms are adhered. Some of the reasons for norm adherence include:

- Fear of authority;
- Rational appeal of the norms; and
- Feelings such as shame, embarrassment and guilt that arise because of nonadherence.

Elster (1989) categorizes norms into consumption norms (e.g., manners of dress), behaviour norms (e.g., norm against cannibalism), norms of reciprocity (e.g., gift-giving norm), norms of cooperation (e.g., voting and tax compliance) and so forth.

Normative Multi-Agent Systems

The research of norms in multi-agent systems is recent (Boman, 1999; Conte, Falcone, & Sartor, 1999; Shoham & Tennenholtz, 1995). Norms in multi-agent systems are treated as constraints on behaviour, goals to be achieved or as obligations (Castelfranchi, 1995). There are two main research branches in normative multi-agent systems. The first branch focuses on normative system architectures, norm representations and norm adherence and the associated punitive or incentive measures. The second branch of research is related to emergence of norms.

Lopez and Marquez (2004) have designed an architecture for normative BDI agents and Boella and Torre (2006) have proposed a distributed architecture for normative agents. Some researchers are working on using deontic logic to define and represent norms (Boella & Torre, 2006; Garcia-Camino, Rodriguez-Aguilar, Sierra, & Vasconcelos, 2006). Several researchers have worked on mechanisms for norm compliance and enforcement (Aldewereld et al., 2006; Axelrod, 1986; Lopez, Luck, & Inverno, 2002). A recent development is the research on emotion-based mechanism for norm enforcement by Fix, Scheve, and Moldt (2006).

Related Work on Emergence of Norms

The second branch focuses on two main issues. The first issue is on norm propagation within a particular society. According to Boyd and Richerson (1985), there are three ways by which a social norm can be propagated from one member of the society to another. They are:

- Vertical transmission (from parents to offspring);
- Oblique transmission (from a leader of a society to the followers); and
- Horizontal transmission (from peer to peer interactions).

Norm propagation is achieved by spreading and internalization of norms. Boman and Verhagen (Boman, 1999; Verhagen, 2000, 2001) have used the concept of normative advice (advise from the leader of a society) as one of the mechanisms for spreading and internalizing norms in an agent society. Their work focuses on norm spreading within one particular society and does not address how norms emerge when multiple societies interact with each other. The concept of normative advice in their context assumes that the norm has been accepted by the top-level enforcer, the Normative Advisor, and the norm does not change. But, this context cannot be assumed for scenarios where norms are being formed (when the norms undergo changes). So, the issue that has not received much attention is the emergence of norms in multi-agent societies. But, there are lots of literature in the area of sociology on why norms are accepted in agent societies and how they might be passed on. Karl-Dieter Opp (Opp, 2001) has proposed a theory of norm emergence based on sociological concepts. Epstein (2001) has proposed a model of emergence based on the argument that the norms reduce individual computations and has provided some results. Our objective in this chapter is to propose a mechanism for norm emergence based on the concept of oblique norm transmission in artificial agent societies. We also provide our experimental results.

PROPOSED MECHANISMS

In this section, we will describe the mechanisms that help norm emergence when different agent societies with different norms interact with each other. Assume that two agent societies with different norms inhabit a particular geographical location. When these societies are co-located, interactions between them are inevitable. When they interact with each other, their individual societal norms might change. The norms may tend to emerge in such a way that it might be beneficial to the societies involved. Our working hypothesis is *Interactions between agent societies with different norms in a social environment (with a shared context), results in the convergence of norms. Norm convergence might result in the improvement of the average performance of the societies.* To demonstrate our hypothesis, we have experimented with agents that play the Ultimatum game (Slembeck, 1999). The shared context of interaction is the knowledge of the rules of the game. This game has been chosen because it is claimed to be sociologists' counter argument to the economists' view on rationality (Elster, 1989).

Ultimatum Game

The Ultimatum game (Slembeck, 1999) is an experimental economics game in which two parties interact anonymously with each other. The game is played for a fixed sum of money (say x dollars). The first player proposes how to divide the money with the second player. Say, the first player proposes y dollars to the second player. If the second player rejects this division, neither gets anything. If the second accepts, the first gets (x-y) dollars and the second gets y dollars. For example, assume that each game is played for a sum of 100 dollars by two agents, A and B. Assume that A offers 40 dollars to B. If B accepts the offer, then A gets 60 dollars and B gets 40 dollars. If B rejects the offer both of them do not get any money.

Description of the Multi-Agent Environment

An agent society is made up of a fixed number of agents. For our experiments we have designed two kinds of societies, namely selfish and benevolent societies, as shown in Figure 1. Society 1 and Society 2 correspond to selfish and benevolent societies, respectively. Society 1 is modeled after the materialistic world where agents try to maximize their personal income. Selfish agents propose the least amount of money and accept any non-zero amount. The second kind of society

Figure 1. Architecture of the experimental framework

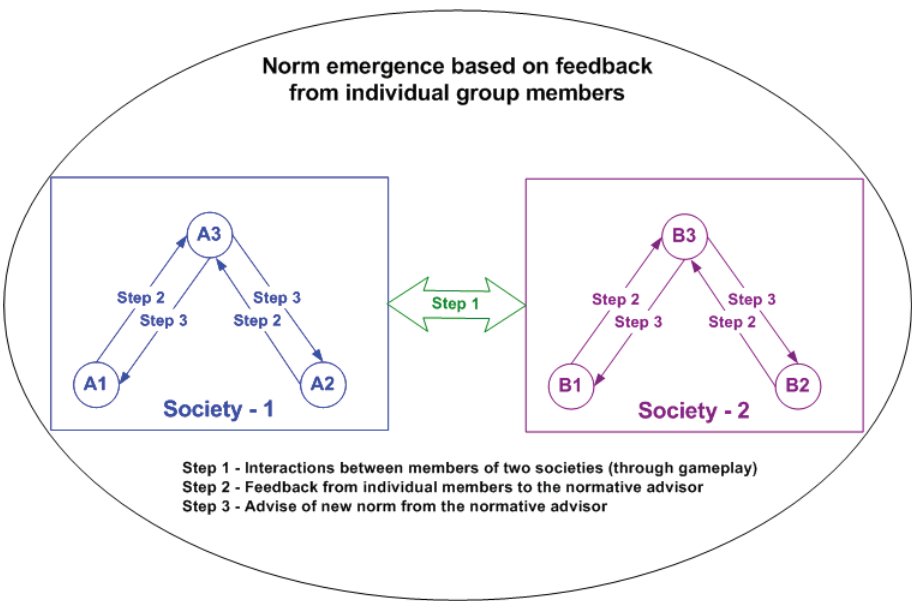

is the benevolent society such as the Ika tribe of Ethiopia (Elster, 1989). The benevolent agents are generous agents. They propose more than the fair share. But, they expect nothing less than the fair share. They also reject high offers. Each agent has two types of norms:

- Group norm (G norm); and
- Personal norm (P norm).

The G norm is shared by all the members of the society. The P norm is internal to the agent and it is not known to any other member. Autonomy is an important concept associated with choosing either a G norm or a P norm when an agent interacts with another agent. When an agent is created, it has an autonomy value uniformly distributed between 0 and 1. Depending upon the autonomy value, an agent chooses either the G norm or the P norm. For example, if the autonomy of an agent is .4, it chooses P norm 4 times and the G norm 6 times out of 10 games. Normative Advisor is one of the agents in the society, which is responsible for collecting the feedback from the individual agents. It modifies the G norm of the society and advises the change to all the members of the society. As shown in Figure 1, the Normative Advisor agents of the two societies are A3 and B3, respectively.

Experimental Parameters

The G norm and P norm are made up of two sub norms, namely the proposal norm and the acceptance norm. The proposal norm corresponds to the range of values (minimum and maximum values) that an agent is willing to propose to other agents. The acceptance norm corresponds to the range of values that an agent is willing to accept from other agents. A sample G norm for a selfish agent looks like the following where min and max are the minimum and maximum values when the game is played for a sum of 100 dollars.

- G-Proposal norm (min=1, max=30)
- G-Acceptance norm (min=1, max=100)

The representations given above indicate that the group proposal norm of the selfish agent ranges from 1 to 30 and the group acceptance norm of the agent ranges from 1 to 100. A sample P norm for a selfish agent might look like the following.

- P-Proposal norm (min=10, max=40)
- P-Acceptance norm (min=20, max=100)

Initially, the G norm of a society is assigned with a particular value, which will be shared by all the members of the society. The personal norms will vary from one agent to another. An agent can accept or reject a proposal based on the norm it chooses (which is based on its autonomy).

Collective Feedback Mechanism for Norm Emergence

In this section, we describe our mechanism for norm emergence that is based on collective feedback of individual agent experiences when playing the Ultimatum game against agents in the other society. The agents have a common G norm to start with. They also have an internal P norm. Both norms continuously evolve based on social learning to maximize the benefit of the society. In the context of the Ultimatum game, the goal is to improve the performance of the overall society while maximizing their own benefit. In one iteration, every agent in a society plays an equal number of games against all the agents in the other society. After the end of each game the agents record the history of interactions (both successes and failures). At the end of each iteration, all the agents submit their successful proposal and acceptance values to the Normative Advisor Agent of their society.

The Normative Advisor Agent uses the average successful values submitted by all the agents in a society and derives the new G norm value for the group. In each iteration the Normative Advisor Agent fractionally increases or decreases G norm values for a society so that it can accommodate the norms of the other society. This mechanism will reduce the overall losses and increase the overall income. After each iteration, the group norm will be propagated to all the agents in the society. Similar to the G norm, P norm of an agent will also change continuously. While G norm changes only at the end of each iteration, P norm changes within each iteration. When an agent chooses P norm over G norm, the outcome of that game determines whether the P norm will change or not. For example, when an agent's proposal that is based on a P norm is rejected *n* consecutive times, the agent modifies its P norm. The agent modifies its P norm fractionally so that it moves closer to the G norm.

EXPERIMENTATION AND RESULTS

The agents in our experiments are built on Otago Agent Platform (Purvis et al., 2002) and they communicate using FIPA ACL messages ("Foundation for Intelligent Physical Agents (FIPA)," 2007). Our experimental set up is made up of two societies with fixed number of agents in each society. In each iteration an agent plays the ultimatum game with all the players in the other group. The games were played over a fixed number of iterations (5 to 5000). In the first experiment, the agents do not use the designed mechanisms. In the second experiment, the agents use designed mechanism. At the end of each experiment, we observe whether norms emerge (whether the proposal norms stabilize or not). In the third and the fourth experiments, we explore the role of autonomy and the visibility threshold, respectively, on norm emergence.

The initial G norms associated with the three experiments are given below.

- G-Proposal norm for selfish society (min=1, max=30)
- G-Acceptance norm for selfish society (min=1, max=100)
- G-Proposal norm for benevolent society (min=55, max=70)
- G-Acceptance norm for benevolent society (min=45, max=55)

In our experimental setup the minimum and maximum values are parameterized and can be changed easily. We have chosen these sample values to demonstrate the results that we obtained.

Experiment 1: Societies that Resist Changes

Assume that the two societies that play the Ultimatum game resist changes to their G norms and P norms. In this scenario, the G norms are the same across all agents in one society. The P norms will be different from one agent to another. The agents do not change their G or P norms over all iterations.

The results of the average game money won by both societies in this scenario are shown in Figure 2. It can be observed that the performance of both societies are well below what could be achieved by both groups if they were rational such as the Utopian Society. Utopian Society, in its most common and general meaning, refers to a hypothetical perfect society. It is synonymous to a fair society where the average income for the Ultimatum game will be 50. When sociologists conducted Ultimatum game experiments in modern societies, many of the societies proposed the fair 50-50 split. This indicates that the norm of fairness had evolved in these societies (Elster, 1989). The performance of the selfish society in this experiment is better than the benevolent society because the selfish agents accept any non-zero proposal.

Figure 2. Performance of societies based on initial societal norms

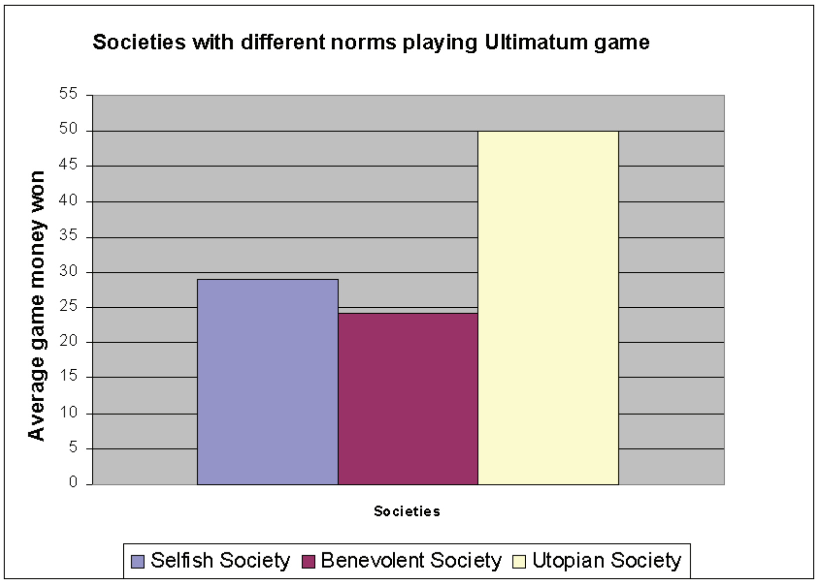

Experiment 2: Societies that Use Collective Feedback from Agents

In this experiment, both societies use the collective feedback mechanism. Figure 3 shows the G-Proposal norm changes of the benevolent as well as the selfish societies over 100 iterations. It can be observed that both groups are continuously changing their G-Proposal norm to accommodate the G-Proposal norm of the other group. Initially, the G-Proposal norm values for the benevolent group decrease because the Normative Advisor Agent changes the norm closer to the selfish societies' G-Proposal norm (based on the collective feedback). For the same reason the G-Proposal norm values for the selfish society increase (until iteration 32). Then, the norms in both societies oscillate to move closer to each other. When, one societies' maximum and minimum values are closer to the other, the G proposal norms start to converge (around iteration 80). These experiments show that the overall performance of the societies have improved as a result of norm emergence, as shown in Figure 3. It can also be observed that the ideal values are not reached as the agents are autonomous and may choose to ignore the G norm, particularly when the autonomy values are high. But, when the number of iterations increased to 5000, the outcomes were closer to the norm of fairness.

Experiment 3: Effect of Autonomy on Norm Convergence in an Agent Society

Unlike previous experiments where norm emergence was observed when two societies come together, in this experiment we observe the effect of autonomy on norm emergence in a single agent society.

The objective of the experiment was to study the effect of autonomy on norm emergence. There were 20 agents in a society and the agents played the Ultimatum game. The experiments were conducted over 20, 50 and 100 iterations. These experiments were carried out using two values of autonomy for all agents (0.2 and 0.8) representing lower and higher autonomy values.

It can be observed from Figure 4 that, when the autonomy of an agent is high (0.8), the con-

Figure 3. Emergence of norms based on collective feedback mechanism

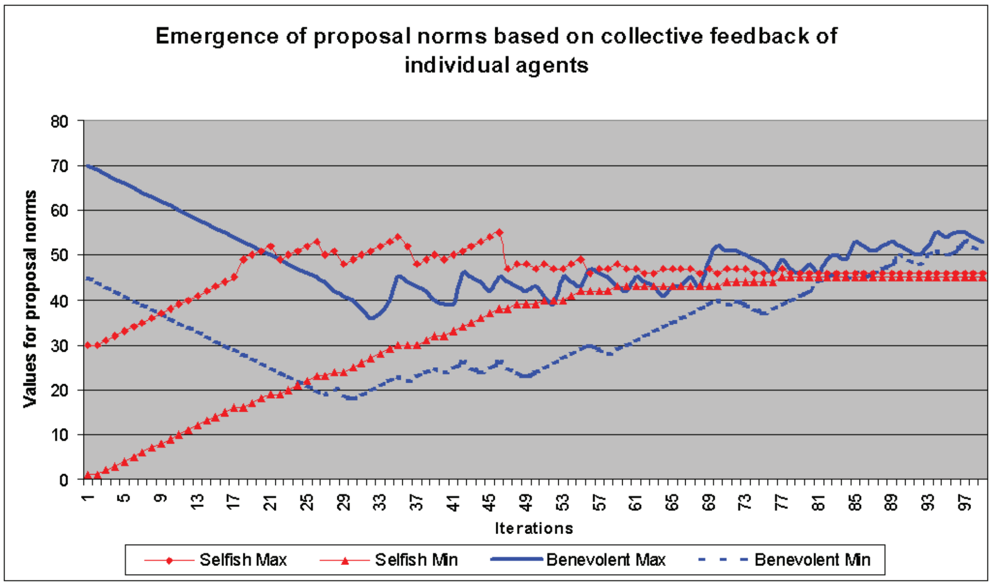

Figure 4. Effect of autonomy on norm emergence

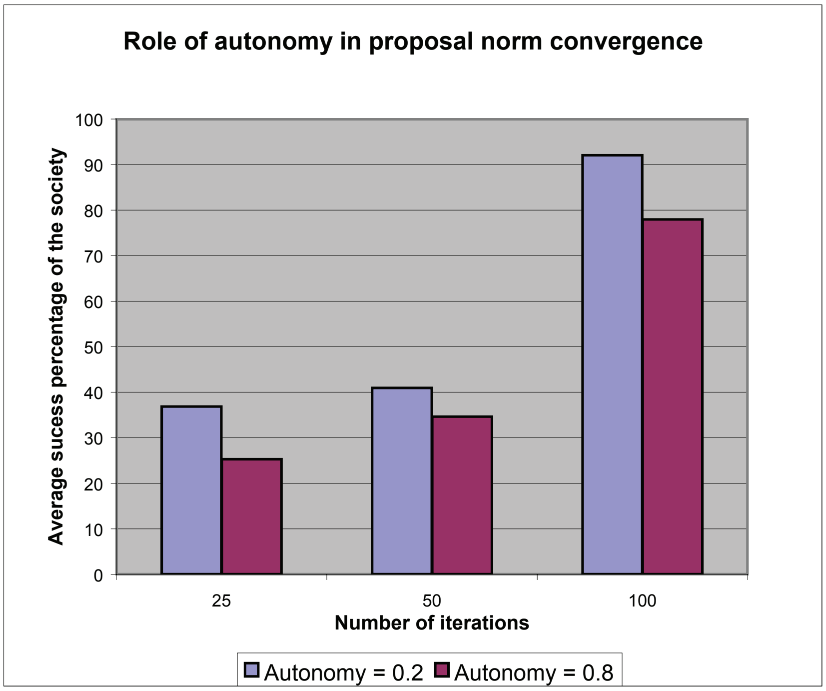

Figure 5. Effect of visibility threshold on norm emergence

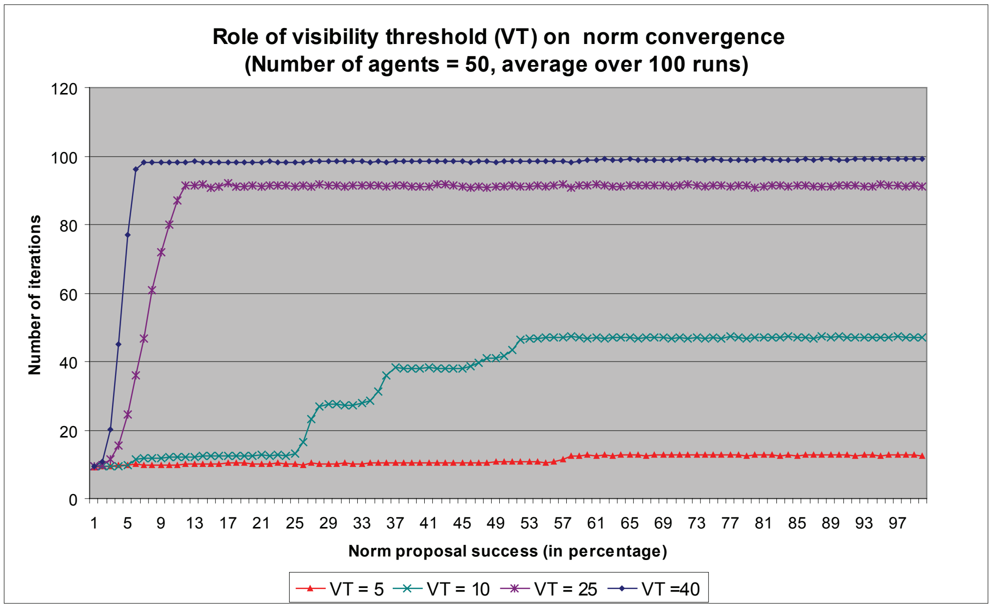

vergence of the norm is low. This indicates the negative effect of autonomy on the system. This result indicates that societies that have more autonomous agents will adopt or evolve norms slower than agent societies that have less autonomous or cooperative agents. This is because the agents that have higher autonomy tend to resist changes to their norms. After obtaining the feedback from the Normative Advisor agent, they move close to the advisor's norm depending upon their autonomy. If the autonomy is higher they do not readily adopt the recommendations provided through normative advice.

Experiment 4: Effect of Visibility Threshold on Norm Emergence

Assume that the collective feedback mechanism is modified in such a way that an agent can choose to seek advice from a local normative advisor agent as opposed to a centralized normative advisor agent. In this modified mechanism, an agent can choose another agent as its normative advisor whose successful proposal norm is within a limit represented by Visibility Threshold (VT). For example, if VT = 5 and an agent's successful proposal average is 80%, then the agent can choose another agent whose successful proposal average is between 80 and 85%.

We have conducted experiments using a society of 50 agents and varying the values for VT (5, 10, 25, 50). It can be observed from Figure 5 that, as the visibility threshold increases, the rate of norm emergence increases. When VT increases, an agent gets to choose a normative advisor within a broader spectrum and the probability of choosing a highly successful role model is high. So, convergence is faster for larger values of VT.

DISCUSSION

The experiments described in this chapter are our initial efforts in the area of norm emergence. Verhagen's thesis (Verhagen, 2000) focuses on the spreading and internalizing of norms. This assumes that a norm is agreed or chosen by a top-level entity (say, a Normative Advisor) and this G norm does not change. The G norm is spread to the agents through the normative advice using a top-down approach. Our work differs from this work, as we employ a bottom-up approach through the collective feedback mechanism. Another distinction is that our work focuses on norm emergence across societies, while the former concentrates on norm propagation in one particular society. In our work both P norm, as well as G norm, evolve continuously. In their work, P norm changes to accommodate the predetermined G norm.

The success of norm emergence using the proposed mechanisms can be explained by the theory of instrumentality proposition proposed by Karl-Dieter Opp (Opp, 2001). The four positive criteria for norm emergence specified by Karl are given below.

1. ***Homogeneity of goals G*** - In our experiments, the goal of an agent was to maximize its personal and societal income.
2. ***Knowledge that a norm N leads to G*** - The agents in our system worked toward establishing a norm that leads to an increase in overall score of the society.
3. ***Knowledge that behaviour B leads to N*** - The agents are aware that by reporting their experience to the Normative Advisor Agent, they can help to achieve the group goal.
4. ***Incentives to perform B*** - The agents know that they can increase their own personal score by providing feedback and receiving the advice. Another incentive for an agent to report experiences is its eagerness to predict other agents' behaviour (e.g., knowing the acceptance range of the other agent).

This emerging area of research on norm emergence offers interesting avenues for further research. In the real world, people are not related to each other by chance. They are related to each other through the social groups that they are in, such as the work group, church group, ethnic group and the hobby group. Information tends to percolate among the members of the group through interactions. People seek advice from a close group of friends and hence information gets transmitted between the members of the social network. Therefore, it is important to experiment our mechanism for norm emergence on top of social networks. In our recent work, we have investigated the role of topologies such as random networks and scale-free networks (Savarimuthu, Cranefield, Purvis, & Purvis, 2007a, 2007b). We have also demonstrated how the role model agent mechanism for norm emergence works on top of dynamically changing network topologies (Savarimuthu et al., 2007a, 2007b). These dynamically changing network topologies represent the social space in which agents can join and leave the network at any time.

An interesting problem in the context of norm emergence mechanism is to experiment with attaching weights to the advice provided by others. The weights of the edges (links) should be considered when the agent makes a decision on whom to choose as advisor agents. We plan to incorporate these ideas in our future experiments.

CONCLUSION

We have explained a mechanism for norm emergence in artificial agent societies. The mechanism used collective feedback of individual agent experiences. We have demonstrated the use of oblique norm transmission in these mechanisms for norm emergence. Through the experimental results, we have shown that norms emerge in agent societies when two different societies are brought together, and this norm might be beneficial to the societies as a whole. We have demonstrated the role of autonomy and visibility threshold of an agent on norm emergence. We have also discussed our future work.

REFERENCES

Aldewereld, H., Dignum, F., Garcia-Camino, A., Noriega, P., Rodriguez-Aguilar, J. A., & Sierra, C. (2006). Operationalisation of norms for usage in electronic institutions. In *Paper presented at the AAMAS*, Hakodate, Japan.

Axelrod, R. (1986). An evolutionary approach to norms. *The American Political Science Review, 80*(4), 1095-1111.

Boella, G., & Torre, L. v. d. (2006). An architecture of a normative system: Counts-as conditionals, obligations and permissions. In *Paper presented at the AAMAS*, New York.

Boman, M. (1999). Norms in artificial decision making. *Artificial Intelligence and Law, 7*(1), 17-35.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.

Castelfranchi, R. C. C. (1995). *Cognitive and social action.* London: UCL Press.

Clutton-Brock, T. H., & Parker, G. A. (1995). Punishment in animal societies. *Nature, 373*, 209-216.

Conte, R., & Castelfranchi, C. (1999). From conventions to prescriptions—Towards an integrated view of norms. *Artificial Intelligence and Law, 7*(4), 323-340.

Conte, R., Falcone, R., & Sartor, G. (1999). Agents and norms: How to fill the gap?. *Artificial Intelligence and Law, 7*(1), 1-15.

Elster, J. (1989). Social norms and economic Theory. *The Journal of Economic Perspectives, 3*(4), 99-117.

Epstein, J. M. (2001). Learning to be thoughtless: Social norms and individual computation. *Computer Economy, 18*(1), 9-24.

Fix, J., Scheve, C. v., & Moldt, D. (2006). Emotion-based norm enforcement and maintenance in multi-agent systems: Foundations and petri net modeling. In *Paper presented at the AAMAS*.

Foundation for Intelligent Physical Agents (FIPA). (2007). Retrieved April 3, 2008, from http://www.fipa.org

Garcia-Camino, A., Rodriguez-Aguilar, J. A., Sierra, C., & Vasconcelos, W. (2006). Norm-oriented programming of electronic institutions. In *Paper presented at the AAMAS*, New York.

Gintis, H. (2003). Solving the puzzle of prosociality. *Rationality and Society, 15*(2), 155-187.

Habermas, J. (1985). *The theory of communicative action: Reason and the rationalization of society* (Vol. 1). Beacon Press.

Lopez, F., Luck, M., & Inverno, M. (2002). Constraining autonomy through norms. In *Paper presented at the Proceedings of the First International Joint Conference on Autonomous Agents and Multi Agent Systems AAMAS'02*.

Lopez, F. L. Y., & Marquez, A. A. (2004). An architecture for autonomous normative agents. In *Paper presented at the Fifth Mexican International Conference in Computer Science (ENC'04)*, Los Alamitos, CA, USA.

Opp, K.-D. (2001). How do norms emerge? An outline of a theory. *Mind and Society, 2*(1), 101-128.

Purvis, M., Cranefield, S., Nowostawski, M., Ward, R., Carter, D., & Oliveira, M. (2002). Agent cities interaction using the opal platform. In *Paper presented at the Proceedings of the Workshop on Challenges in Open Agent Systems, AAMAS*.

Savarimuthu, B. T. R., Cranefield, S., Purvis, M., *& Purvis, M. (2007a). Mechanisms for norm emergence in multi-agent societies.* In *Paper presented at the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Honolulu, HI, USA.

Savarimuthu, B. T. R., Cranefield, S., Purvis, M., & Purvis, M. (2007b). Role model based mechanism for norm emergence in artificial agent societies. In *Paper presented at the International Workshop on Coordination, Organization, Institutions and Norms (COIN) at AAMAS 2007*, Honolulu, HI, USA.

Shoham, Y., & Tennenholtz, M. (1995). On social laws for artificial agent societies: Off-line design. *Artificial Intelligence, 73*(1-2), 231-252.

Slembeck, T. (1999). *Reputations and fairness in bargaining--experimental evidence from a repeated ultimatum game with fixed opponents* (Experimental). EconWPA.

Tuomela, R. (1995). *The importance of us: A philosophical study of basic social notions.* Stanford, CA: Stanford Series in Philosophy, Stanford University Press.

Verhagen, H. (2000). *Norm autonomous agents.* Department of Computer Science, Stockholm University.

Verhagen, H. (2001). Simulation of the learning of norms. *Social Science Computer Review, 19*(3), 296-306.