Aspects of Active Norm Learning and the Effect of Lying on Norm Emergence in Agent Societies

Bastin Tony Roy Savarimuthu, Rexy Arulanandam, and Maryam Purvis

University of Otago, Dunedin, P.O. Box 56, Dunedin, New Zealand
{tonyr,tehrany}@infoscience.otago.ac.nz,
rexi1010@gmail.com

Abstract. Norms have facilitated smoother functioning in human societies. In the field of normative multi-agent systems researchers are interested in investigating how the concept of social norms can be used to facilitate social order in electronic agent societies. In this context, the area of norm emergence has attracted a lot of interest among researchers. The objectives of this paper are two-fold. First, we discuss the norm learning approaches in agent societies and discuss the three aspects of active norm learning (experiential, observational and communication-based learning) in agent societies. Using an example we demonstrate the usefulness of combining these three aspects of norms learning. Second, we investigate the effect of the presence of liars in an agent society on norm emergence. Agents that lie distort truth when they are asked about the norm in an agent society. We show that lying has deleterious effect on norm emergence. In particular, using simulations we identify conditions under which the norms that have emerged in a society can be sustained in the presence of liars.

1 Introduction

Norms have been of interest to researchers in multi-agent systems because they enable cooperation and coordination among software agents. They are also light-weight mechanisms for enabling social control. Agents that know about norms in agent societies do not need to recompute what the norms of the society are and also do not often need to spend time in contemplating actions that are forbidden and obliged as they are aware of these norms. Also, agents that are aware of norms know that violating them will have consequences for them. However, this is true only when the agents know what the norms are. A new agent joining an open society may not know what the norms of the society are. This agent will need to be equipped with some mechanism for learning the norms in the society.

Researchers have employed several mechanisms for the learning of norms. These include imitation, normative advice from leaders, machine learning and data-mining [15]. Section 2 provides a brief background on norm research in the field of multi-agent systems. Section 3 discusses different mechanisms used by researchers for norm learning. In Section 4 we discuss three aspects of active learning namely experiential learning, observational learning and communication-based learning and compare the use of these three aspects in the existing works on norm learning. We also discuss how these three aspects can be integrated in the context of a simple example. In Section 5 we

D. Kinny et al. (Eds.): PRIMA 2011, LNAI 7047, pp. 36–50, 2011.
© Springer-Verlag Berlin Heidelberg 2011

investigate the effect of lying on norm emergence when communication-based learning is used. We identify the tipping points (i.e. number of liars in the society) after which the norms are no longer sustained in an agent society.

Thus, this paper makes two contributions to normative multi-agent systems. First, it discusses how active learning on the part of the agent helps in the process of expediting norm learning in agent society. Second, it discusses the impact of lying on norm emergence in an agent society.

2 Background

Researchers in multi-agent systems have studied how the concept of norms can be applied to artificial agents. Norms are of interest to multi-agent system (MAS) researchers as they help in sustaining social order and increase the predictability of behaviour in the society. Researchers have shown that norms improve cooperation and collaboration [21, 25]. Epstein has shown that norms reduce the amount of computation required to make a decision [7]. However, software agents may tend to deviate from norms due to their autonomy. So, the study of norms has become important to MAS researchers as they can build robust multi-agent systems using the concept of norms and also experiment on how norms may evolve in response to environmental changes.

Research in normative multi-agent systems can be categorized into two branches. Researchers have worked on both prescriptive (top-down) and emergent (bottom-up) approaches to norms.The first branch focuses on normative system architectures, norm representations, norm adherence and the associated punitive or incentive measures. Several architectures have been proposed for normative agents (refer to [13] for an overview). Researchers have used deontic logic to define and represent norms [10]. Several researchers have worked on mechanisms for norm compliance and enforcement such as sanctioning mechanisms [2] and reputation mechanisms [5].

The second branch of research is related to emergence of norms [20, 21]. In the bottom-up approach, the agents infer a norm through learning mechanisms [20, 21] and cognitive approaches [1]. This paper contributes to this branch of research by providing an overview of norm learning mechanisms.

3 Approaches to the Learning of Norms

Researchers have employed four types of mechanisms for an individual agent to learn norms: imitation, machine learning, data mining and advice-based learning. The learning mechanisms identified in this paper are extensions to the learning mechanisms discussed in the categorization presented in a previous work [15]. Since the imitation, machine learning and advise-based learning mechanisms (also called as leadership mechanisms) are explained in the previous work, we only provide a brief summary of these three approaches and provide a longer discussion on the data mining approaches.

Imitation Mechanisms. The philosophy behind an imitation-based learning mechanism is *When in Rome, do as the Romans do* [7]. Models based on imitation are characterised by agents first observing and then mimicking the behaviour of what the majority of the agents do in a given agent society (following the crowd). Epstein's main argument [7] for an imitation mechanism is that individual thought (i.e. the amount of computing needed by an agent to infer what the norm is) is inversely related to the strength of a social norm. This implies that when a norm becomes entrenched the agent can follow it without much thought. An issue for debate is whether imitation-based behaviour (solely) really leads to norms as there is no notion of generalized expectation. Imitation mechanism also does not consider sanctions or rewards thereby focus only on conventions and not norms[1]. The direct utility from conforming to a particular behaviour is not modelled in some cases (i.e. blindly imitating what the crowd does without carefully considering the impact of the actions either for the agent or the society).

Works Based on Machine Learning. Several researchers have experimented with agents finding a norm based on learning on the part of an agent when it interacts with other agents in the society by performing some actions [20, 21, 25]. Researchers have used simple reinforcement algorithms for norm learning. The reinforcement learning algorithms identify a strategy that maximizes an agent's utility and the chosen strategy is declared as the norm. Since all agents in the society make use of the same algorithm, the society stabilises to an uniform norm. Agents using this approach cannot distinguish between a strategy and a norm. These agents accept the strategy that maximizes its utility as its norm. However, the agents do not have a notion of normative expectation associated with a norm (i.e. agents do not expect certain behaviour on the part of other agents). Another weakness is that agents in machine learning approach do not have a mental notion of norms (i.e. the ability to reason about why norms have to be followed and the consequences for not following norms) as they are mainly utilitarian agents. These limitations are addressed by the works that employ cognitive approaches where the norms learnt affect an agent's future decision making by influencing its beliefs, intentions and goals [1].

Advice-Based Learning. Boman [3] has used a centralised approach, where agents consult with a normative advisor before they make a choice on actions to perform. Verhagen [23] has extended this notion of normative advice to obtaining normative comments from a centralized normative advisor (e.g. the leader of the society) on an agent's previous choices. Savarimuthu et al. [16] have adopted a distributed approach approach for normative advice. In their mechanism, there could be several normative advisors (called role models) from whom other agents can request advice. Hoffmann [9] has experimented with the notion of norm entrepreneurs who think of a norm that might be beneficial to the society. An entrepreneur can recommend a norm to a certain percentage of the population (e.g. 50%) which leads to varying degrees of establishment of a norm. The models based on advice assume that a powerful authority is present in the society and all agents in the society acknowledge the power of such agents. Both centralised and distributed notions of norm spreading using *power* have been employed. The centralised approach is suitable for closed societies. However, this might not work well for open, flexible and dynamic societies. Distributed approaches for norm spreading and emergence are promising because the computational cost required to spread, monitor and control a norm is distributed to all the members of the society.

[1] Many sociologist consider sanctions and/or rewards a core part of the norm.

Data Mining Mechanisms. Agents can use data mining approach to identify norms in agent societies. Agents in open agent societies can learn norms based on what they infer based on their observations of the society. The repository of an agent's observations can be mined for patterns of behaviour. There has been a proposal of an agent architecture for normative systems to employ data mining for citizens of a country to find information and norms from official documents [22]. However, the work does not describe what types of norms are discovered and also the mechanisms used in the identification of norms.

Savarimuthu et al. [17, 18] have proposed an architecture for norm identification which employs association rule mining, a data mining approach. The architecture makes use signals (sanctions and rewards) as the starting points for norm identification. Mechanisms for identifying two types of norms, prohibition norms and obligations norms have been studied. The details on how an agent identifies a prohibition norm are explained in the context of a public park scenario, where the norm against littering is identified by the agent. The obligation norm inference is explained in the context of a tipping norm in a restaurant scenario. They have demonstrated that an agent using the proposed architecture can dynamically add, remove and modify norms based on mining the interactions that take place between agents. They have shown that agents can identify co-existing norms. The agents can also identify conditional norms (e.g. identification of *normative pre-conditions* which are conditions that have to be true for the norm to hold).

In the work of Lotzmann et al. [12] an agent constructs a decision tree of events that take place. It learns norms by considering the occurrence probabilities of those events that take place. For example, an agent participating in a traffic scenario either as a pedestrian or a car driver, decides about which action to perform, based on the probability of events represented as nodes of a decision tree. Based on these probabilities, a pedestrian agent learns that if it jaywalks instead of using the pedestrian crossing, it has a high probability of being run over by a car. A car driver learns to stop in the pedestrian crossing area.

Data mining is a promising approach for the identification of some types of norms that can be inferred based on observing the interactions between agents in the society. However, if actions that explicitly signal a sanction or reward are absent or other mechanisms such as reputation are used instead of explicit signals (i.e. reduction in the reputation score of a rogue agent instead of explicit sanctioning that is visible to other agents), then it is difficult to identify norms.

4 Aspects of Active Learning of Agents

Hamada et al. [8] note that *active learning is learning with learners involved in the learning process as active partners: meaning they are "doing", "observing" and "communicating" instead of just "listening" as in the traditional learning style*. An actively learning agent can thus learn about norms in the following three ways.

- **Experiential learning** - This is the ability of an agent *learning by doing*. For example, an agent may litter a park. It may be sanctioned by some other agent(s).

Through the sanction experienced as a result of the littering action, the agent can learn about the norm. Thus, an agent can learn from its personal experience based on sanctions and rewards.

- **Observational learning** - This is the ability of an agent *learning by observing*. For example, an agent may observe littering agents being sanctioned in a society. Through the observation of the sanction on others, an agent can learn about the norm.
- **Communication-based learning** - This is the ability of an agent *learning by communicating* with other agents. For example, an agent may ask another agent in the park what the norms of the park are and that agent may communicate the norm to the agent. Norm communication can happen at a peer-to-peer level or from leaders to follower agents.

Table 1 shows the aspects of learning used by different research works investigating norms. It can be noticed that not all the three types of learning have been investigated by many research works. Only some of the recent research works have considered all the three types of learning [6, 18].

Table 1. Comparison of the types of learning employed by different research works (Yes - considered, No - not considered)

Model	Experiential learning	Observational learning	Communication-based learning
Axelrod, 1986 [2]	No	Yes	No
Shoham and Tennenholtz, 1992 [21]	Yes	No	No
Kittock, 1993 [11]	Yes	No	No
Walker and Wooldridge, 1995 [25]	Yes	No	Yes
Verhagen, 2001 [23]	No	No	Yes
Epstein, 2001 [7]	No	Yes	No
Hoffmann, 2003 [9]	No	Yes	No
Pujol, 2006 [14]	Yes	No	No
Sen and Airiau, 2007 [20]	Yes	No	No
Savarimuthu et al., 2010 [17, 18]	Yes	Yes	Yes
EMIL Project, 2006-2010 [6]	Yes	Yes	Yes

4.1 Comparing Different Combinations of Learning

In this section we demonstrate how the three types of learning can be carried out together in the context of an example that uses the machine learning approach. Most research works using machine learning mechanisms have investigated only the experiential learning aspect [14, 20].

Consider the scenario where agents strive to establish a convention of driving either on the left (L) or the right (R) of the road. The payoff matrix for this coordination game is given in Table 2. The goal of the learning task is to enable the agents in a society to drive either on the right or the left. This goal can be achieved through several combinations of three aspects of learning[2]. In this work, we will compare three combinations just to demonstrate that the use of more than one aspect of learning improves the rate of norm emergence[3]. These three combinations are 1) learning by doing, 2) learning by doing and observing and 3) learning by doing, observing and communicating.

Table 2. Payoff matrix

	L	R
L	1, 1	-1, -1
R	-1, -1	1, 1

Assume that there are 100 agents in the system. In each iteration the agents randomly interact with one other agent by choosing an action (L or R). Based on the outcome of the interaction, the agent learns which action to choose for the next iteration. We use a simple Q-Learning algorithm [26] to facilitate learning similar to other works on norm emergence [24]. The Q-value is calculated using the formula

$$Q^t(a) = (1 - \alpha) * Q^{t-1}(a) + \alpha * R \tag{1}$$

where $Q^t(a)$ is the utility of the action a chosen by an agent after t times, R is the reward for performing an action and α is the learning rate. Agents choose actions that yield the highest utility.

The pseudocode described in Algorithm 1 provides the operational details of an agent in learning a norm. An agent may learn based on a combination of different aspects. If an agent learns by doing, then, when it interacts with another agent (in the context of playing the coordination game), it chooses the action for which it has the highest Q-value. It then updates the Q-value of the chosen action based on the reward obtained. If an agent chooses to learn by observing, it learns by observing a randomly chosen agent X interacting with another agent Y. It then updates the Q-value for the action that was chosen by X, based on the reward obtained by X. If an agent learns through communication, the agent asks another agent for the action that it considers to be the norm. It then updates the Q-value of the action recommended as the norm with a reward of one.

[2] Seven combinations are possible without considering repetition. These are *doing, observing, communicating, doing-observing, observing-communicating, communicating-doing, doing-observing-communicating*.

[3] Conventions and norms are broadly considered under the same umbrella of norms in this paper, similar to some other works in this field [20] on norm emergence and the terminologies have been used interchangeably in this paper. We recognize that the distinction between norms and conventions exist in multi-agent systems by the fact that norms have explicit sanctions associated with them.

Algorithm 1.1. Psuedocode for agent learning

```
while iterationNumber ≤ totalNoOfIterations do
    if learningFromDoingEnabled then
        Each agent interacts with one other agent by choosing an action that has the
        highest Q-value;
        Each agent updates the Q-value of the chosen action based on the reward
        obtained;
    if learningFromObservingEnabled then
        Each agent observes the action chosen by one other agent (X) interacting with
        another agent (Y);
        Each agent updates the Q-value of the action chosen by X based on the reward
        obtained by X;
    if learningFromCommunicationEnabled then
        Each agent asks another agent for the action that is considered to be the norm;
        Each agent updates the Q-value of the action recommended as the norm with a
        reward value of one;
```

We have conducted three experiments by fixing the value of α to 0.3. In the first experiment, an agent learns only through its experience (i.e. based on the result of their interaction with other agents). In the second experiment, in addition to experiential learning they also observe one other agent's action and learn from the result of that action (experiential + observational learning). In the third experiment, in addition to the set-up of the second experiment, an agent also learns from the experience of one other agent (i.e. by asking about the action performed by the agent and the reward it obtained). It should be noted that the third experiment involves all the three aspects of learning.

Figure 1 shows the results of the three experiments as three lines. The results are based on the average of 500 runs per experiment. It demonstrates that experiment three that uses experiential, observational and communication based learning results in the fastest convergence of norms. It is intuitive that an agent that makes use of the three aspects of learning performs better since more information is available to the agent from all the three learning channels. However, it is interesting to note that not all the three aspects have been considered by many research works as shown in Table 1. Even though communication is a fundamental aspect of agent systems, many research works have not considered the possibility of agents learning from others. The reason for this may include the lack of trust on other agents (i.e. agents may lie when asked for advice about a norm).

4.2 The Need for Integrating the Three Aspects of Learning

We note that the future research works on norm learning should consider integrating these three aspects where ever possible. The reasons are outlined below.

1. One of the drawbacks on the experiential learning of norms in an agent society is that an agent cannot perform all possible actions in order to find out what the norms

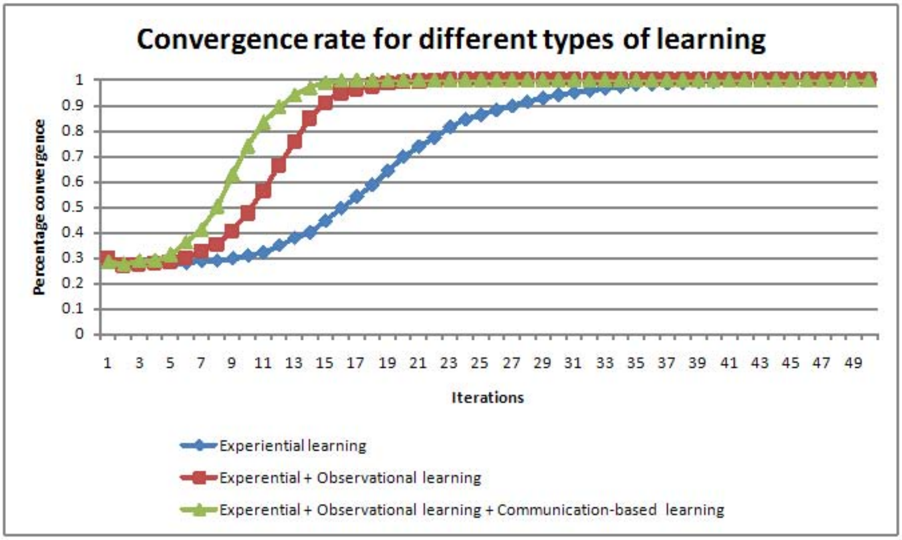

Fig. 1. Comparison of convergence rate in three experiments (varying different aspects of learning)

of the society are. For example, a new agent in a society may not know what the norms are and it may not be desirable to perform all actions to see whether any of those actions result in a sanction by performing it. The state space of actions can be large. Hence, this approach can be computationally expensive. However, if an agent does not actively search for an action that might be sanctioned, but only learns based on receiving a sanction for an action that it performed accidentally, it can use that sanction as a starting point to infer a norm. For example, when it is sanctioned for littering, it can flag the littering action as the potential norm and then check to see in its future interactions with other agents whether littering causes a sanction.

2. Using just the observational learning for learning norms might also cause problems. Assume that agent Z observes agent X punishing Y. Only if Z observes both a) the action responsible for the sanction and b) the sanction itself, it can learn from the observation. However, if the observer (agent Z) does not know which of the actions agent Y had done in the immediate past had caused this sanction, this approach will not be useful. In this case, it has to learn by asking about a norm from an agent in the society (i.e. communication-based learning).
3. Using just the communication learning may be sufficient in regimented societies where norms are prescribed by the organization and in societies where there is no lying. However, in open agent societies it may not be possible to rule out lying. In this case an agent may have to engage in observational learning and/or experiential learning.

5 Lying in Agent Societies

One of the issues with communication-based learning is the ability of agents to lie. Agents being autonomous entities may not be truthful when communicating. For example, an agent that is anti-social may not want a norm to emerge in a society. Therefore, it may try to thwart the process by spreading false information about the norm when asked for advice from other agents.

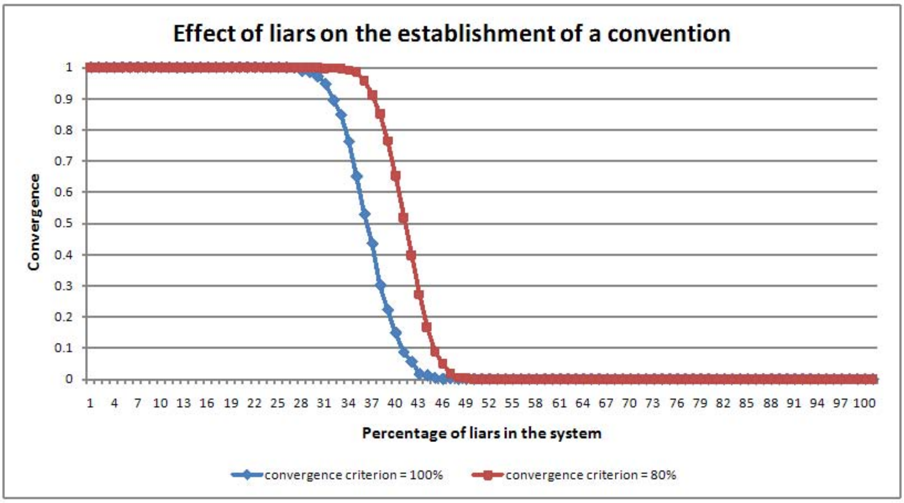

Fig. 2. Effect of liars on convention emergence (observation and communication-based learning)

In this section, we discuss our investigation of the effect of lying in communication-based learning. In Section 5.1, we discuss the effect of lying when communication-based learning is used in conjunction with observation based learning. In Section 5.2 we discuss the effect of lying when communication-based learning is used in conjunction with observation and experiential learning. Our aim is to determine how much lying (i.e. the percentage of liars in a society) would thwart the process of norm emergence. In other words, we aim at investigating how much lying can still be allowed in the society that would not disrupt the norm emergence process.

5.1 Impact of Lying When Observational and Communication-Based Learning Is Used

In a society of 100 agents, by keeping all the other parameters constant, we varied the number of liars in a society from 1 to 100. We conducted 1000 runs of each experiment, with each run spanning 200 iterations. At the end of the runs we calculated the percentage of convergence of the society to one of the conventions (either left or right). Figure 2 shows the effect of liars on norm emergence. The two lines correspond to different

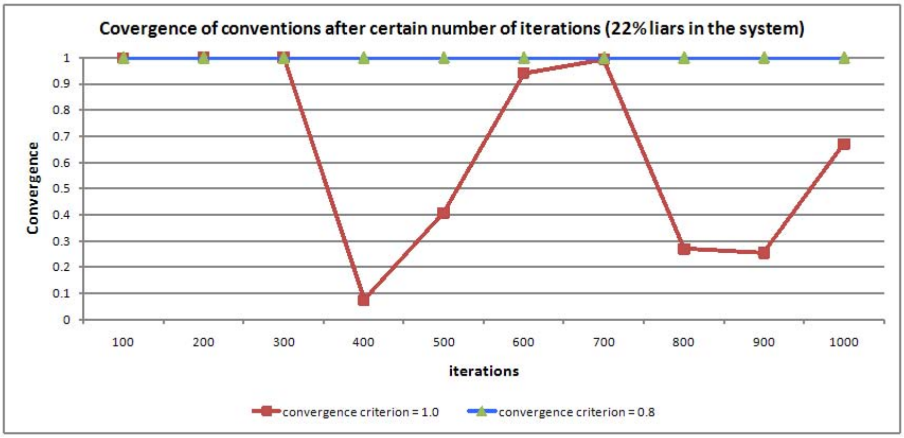

Fig. 3. Effect of liars on convention emergence by varying the number of iterations before convergence is checked

convergence criteria, 80% and 100% in the society respectively[4]. It can be inferred that the system always converges when there are less than or equal to 25% liars in the system when the convergence criterion is 100% and it always converges when there are less than or equal to 29% liars when the convergence criterion is 80%. These points can be considered as tipping points after which there is a rapid drop in the convergence to a norm. Once the system has 46% and 49% liars respectively for convergence criteria of 100% and 80%, the system never converges to a stable norm (i.e. probability of convergence to a norm is 0%).

Since we had only investigated convention emergence after certain number of iterations (i.e. 200 iterations), we investigated convention emergence further systematically at the end of every 100 iterations to a maximum of 1000 iterations (i.e. by conducting 10 experiments). Each experiment was repeated 1000 times. The number of liars was kept to 22. Since the system converged to a convention at 200 iterations, we expected all the experiments after 200 iterations would result in 100% convergence. However, the results that we obtained were surprising (see Figure 3). It can be seen that even though with 22 liars the system had converged in iterations 100, 200 and 300, it did not converge from iterations 400 to 600. It converged in iteration 700 and then it again it did not converge after that. So, there appears to be *cycles* in convergence which we call as *cycle effect*. The cycle is caused by all agents converging to a convention because of learning by observation initially, but once the system has converged to a norm, non-liars are conned by liars and this leads to some of the non-liars moving away from the convention. Soon those non-liar agents that deviated from the convention realize that they have been conned (i.e. through their reduction in utility) and they choose the

[4] 100% convergence means that in all the 1000 runs of the experiment, the system converged to a norm at the end of certain number of iterations (e.g. 200 iterations). 80% convergence criteria means that in all the 1000 runs of the experiment, at least 80% of the agents have converged to a norm.

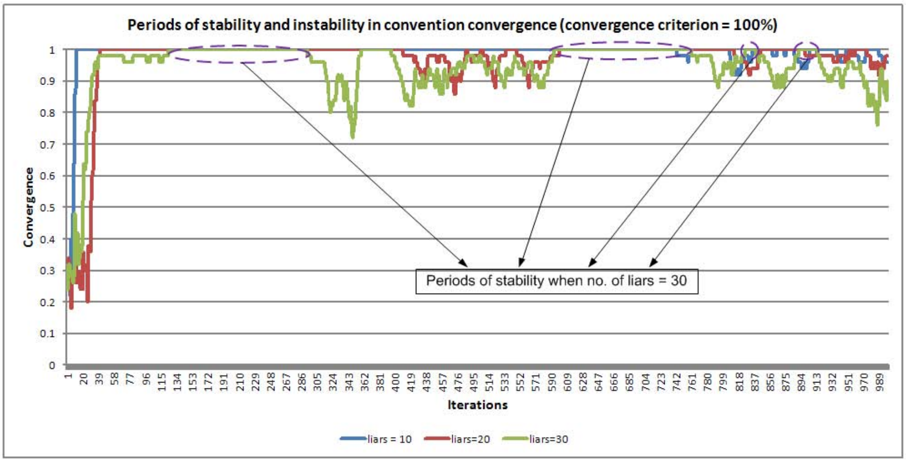

Fig. 4. Periods of stability and instability when the number of liars were varied (convergence criterion = 100%)

appropriate convention. This cycle repeats again. It should be noted that the cycle effect is seen only for 100% convergence and not for 80% convergence which called for further investigation.

In order to investigate the cycle effect further, we examined sample runs (one sample each) from the system where the number of liars were 10%, 20%, and 30% respectively. The convergence criterion was set to 100%. The experiments were conducted for 1000 iterations. It can be observed from Figure 4 that there are periods of stability and instability in convergence (i.e. 100% convergence) in the system. The periods of instability for low liar numbers was smaller than the periods of instability for larger number of liars. These lines also show the cycle associated with convergence.

Figure 5 shows the periods of stability and instability in convergence emergence when the convergence criterion was set to 80% in the system. It is intuitive that there are fewer periods of instability than when the convergence criterion was set to 100% for the same number of liars in the system. Note that there were no periods of instability when the number of liars were 20. This is in agreement with the result shown for approximately same number of liars in Figure 2.

In the light of the result shown in Figure 4 we note that the result shown in Figure 2 still holds since we had investigated convergence in 1000 runs after iteration 200 in each of the runs (in Figure 2). In all 1000 runs the system always converged since iteration 200 falls in the period of stability for 22 liars. However, since the experiment was not run for large number of iterations, it did not capture the dynamics that could have ensued. So, what can be inferred from this result is that we need to measure convergence throughout the entire simulation period instead of just measuring convergence at certain points in time such as iteration 200. Additionally the mechanism employed should check whether a convention is sustained once it has converged. In other words, the investigations should examine whether the system upon reaching the desired convergence level (80% or 100%) becomes unstable again. If it becomes unstable then the convention is said to be unsustained.

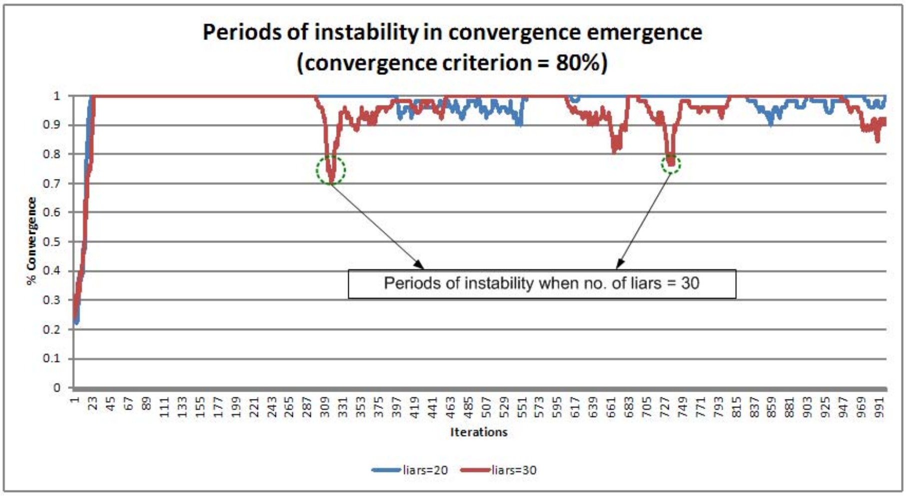

Fig. 5. Periods of stability and instability when the number of liars were varied (convergence criterion = 80%)

The experimental results discussed so far checked for convergence after certain number of iterations. We modified this process by making a simple change where the stability of convention convergence is checked once the initial convergence is reached (i.e. when all the agents converge to a norm). We allowed the agents to interact for certain number of iterations initially before we started measuring convergence. This value was set to 500 in our experiments[5]. Figure 6 shows the results on sustaining norms in an agent society. We have found that in order to sustain convergence there cannot be more than two liars in the society for 100% convergence criterion, and there cannot be more than 11 liars in the society for 80% convergence criterion. These are the tipping points for the number of liars a system can have above which the probability of the system sustaining the emerged norm is less than one. The results reported in this experiment are based on running the experiment 10,000 times (i.e. for each number between 1 to 100 that represents the number of liars in the system, the experiment was repeated 10,000 times).

It can be observed from Figure 6 that there is a considerably long period where the convergence is closer to 100% (percentage of liars from 12 to 22) for 100% convergence criterion. Once the number of liars in the system reach 5 and 33 respectively for 100% and 80% convergence criteria, the probability of system converging to a norm is zero.

5.2 Impact of Lying When All the Three Aspects of Learning Are Considered

We also conducted experiments to investigate the impact of lying when all the three aspects of learning are used in an agent society. For the same parameter settings of the previous experiment (results shown in Figure 6), the tipping points obtained for

[5] We note this value can be changed.

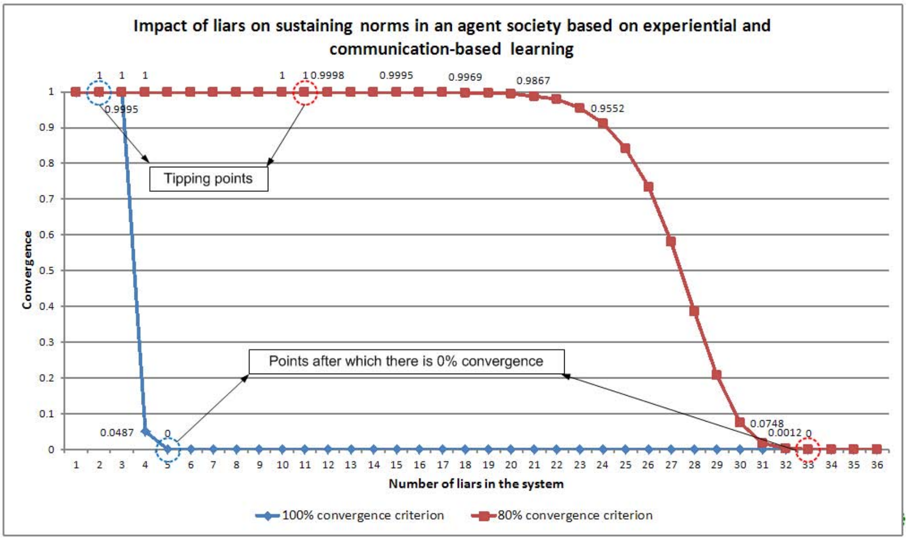

Fig. 6. Tipping points of norm sustainability on varying convergence criteria (80% and 100%)

both 100% and 80% convergence criteria were the same. In both cases, any society that has more than two liars resulted in the destabilization of a norm in the society (not shown here). It is an interesting result because one would assume that the tipping point would be higher (i.e. the set-up can sustain more lying than the result presented in Figure 6) because the impacts of experiential learning and observational learning are the same, hence their additive effect should help in norm stabilization (as opposed to an agent using just the experiential learning). However, it is not the case because once the destabilization phase starts (i.e. liars start impacting the non-liars), the system quickly becomes unstable because the rate of destabilization is twice faster than the previous experiment. It should be noted that the rates of both stabilization and destabilization of a norm are faster when experiential and observational learning are used in conjunction with communication-based learning that permits lying.

6 Discussion

We note that there is scope for further investigation on the lying aspect on norm emergence. In the future the following extensions will be considered.

- The impact of assigning certain weights to each of the three types of learning on the time taken to reach convergence with and without the presence of liars can be investigated.
- The role of network topologies on lying can be investigated in the future. If the agents that are hubs are the liars, then the impact of the lie spreading would be more pronounced than the agent being a leaf node in a network topology.

– Mechanisms for preventing lying in an agent society can be investigated. For example, spreading information about liars using decentralized mechanisms such as gossips [19] can be undertaken.

In this work, we have arrived at the tipping points using a simulation-based approach. It would be desirable to mathematically model scenarios such as the approach used in the work of Brooks et al. [4] to accurately estimate when the liars can thwart the emergence of a norm.

7 Conclusion

The objectives of this paper were two fold. First, in the context of discussing the approaches to norm learning in agent societies, it discussed how three aspects of active learning (learning by doing, observing and communicating) can be integrated to facilitate better norm learning in agent societies. It also demonstrated using a simple example how the three aspects can be integrated. Second, it demonstrated what the effect of liars are on convention emergence when communication-based learning mechanism is used. It identified the tipping points where the convention can no longer be sustained in an agent society.

References

1. Andrighetto, G., Conte, R., Turrini, P., Paolucci, M.: Emergence in the loop: Simulating the two way dynamics of norm innovation. In: Boella, G., van der Torre, L., Verhagen, H. (eds.) Normative Multi-agent Systems. No. 07122 in Dagstuhl Seminar Proceedings, Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany (2007)
2. Axelrod, R.: An evolutionary approach to norms. The American Political Science Review 80(4), 1095–1111 (1986)
3. Boman, M.: Norms in artificial decision making. Artificial Intelligence and Law 7(1), 17–35 (1999)
4. Brooks, L., Iba, W., Sen, S.: Modeling the emergence and convergence of norms. In: International Joint Conference on Artificial Intelligence (IJCAI), pp. 97–102 (2011)
5. Castelfranchi, C., Conte, R., Paolucci, M.: Normative reputation and the costs of compliance. Journal of Artificial Societies and Social Simulation 1(3) (1998)
6. EMIL project report: Emil - emergence in the loop: simulating the two way dynamics of norm innovation. Emil-t, deliverable 5.1, final project report, EMIL consortium (2010), `http://cfpm.org/EMIL-D5.1.pdf`
7. Epstein, J.M.: Learning to be thoughtless: Social norms and individual computation. Computational Economics 18(1), 9–24 (2001)
8. Hamada, M.: Web-based Environment for Active Computing Learners. In: Gervasi, O., Murgante, B., Laganà, A., Taniar, D., Mun, Y., Gavrilova, M.L. (eds.) ICCSA 2008, Part I. LNCS, vol. 5072, pp. 516–529. Springer, Heidelberg (2008)
9. Hoffmann, M.J.: Entrepreneurs and Norm Dynamics: An Agent-Based Model of the Norm Life Cycle. Tech. rep., Department of Political Science and International Relations, University of Delaware, USA (2003)
10. Jones, A.J.I., Sergot, M.: On the characterisation of law and computer systems: The normative systems perspective. In: Deontic Logic in Computer Science: Normative System Specification, pp. 275–307. John Wiley and Sons (1993)

11. Kittock, J.E.: Emergent conventions and the structure of multi-agent systems. In: Nadel, L., Stein, D.L. (eds.) 1993 Lectures in Complex Systems, Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley (1995)
12. Lotzmann, U., Möhring, M.: Simulating norm formation: an operational approach. In: Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 1323–1324. International Foundation for Autonomous Agents and Multiagent Systems (2009)
13. Neumann, M.: A classification of normative architectures. In: Takadama, K., Cioffi-Revilla, C., Deffuant, G. (eds.) Simulating Interacting Agents and Social Phenomena: The Second World Congress, Agent-Based Social Systems, vol. 7, pp. 3–18. Springer, Heidelberg (2010)
14. Pujol, J.M.: Structure in Artificial Societies. Ph.D. thesis, Departament de Llenguatges i Sistemes Informàtics, Universitat Politènica de Catalunya (2006)
15. Savarimuthu, B.T.R., Cranefield, S.: Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. Multiagent and Grid Systems 7(1), 21–54 (2011)
16. Savarimuthu, B.T.R., Cranefield, S., Purvis, M., Purvis, M.K.: Role model based mechanism for norm emergence in artificial agent societies. In: Coordination, Organizations, Institutions, and Norms in Agent Systems III, pp. 203–217 (2007)
17. Savarimuthu, B.T.R., Cranefield, S., Purvis, M.A., Purvis, M.K.: Norm identification in multi-agent societies. Discussion Paper 2010/03, Department of Information Science, University of Otago (2010), http://eprints.otago.ac.nz/873/
18. Savarimuthu, B.T.R., Cranefield, S., Purvis, M.A., Purvis, M.K.: Obligation norm identification in agent societies. Journal of Artificial Societies and Social Simulation 13(4), 3 (2010), http://jasss.soc.surrey.ac.uk/13/4/3.html
19. Savarimuthu, S., Purvis, M.A., Purvis, M.K., Savarimuthu, B.T.R.: Mechanisms for the Self-Organization of Peer Groups in Agent Societies. In: Bosse, T., Geller, A., Jonker, C.M. (eds.) MABS 2010. LNCS, vol. 6532, pp. 93–107. Springer, Heidelberg (2011)
20. Sen, S., Airiau, S.: Emergence of norms through social learning. In: Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI), pp. 1507–1512. AAAI Press (2007)
21. Shoham, Y., Tennenholtz, M.: Emergent conventions in multi-agent systems: Initial experimental results and observations. In: Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR), pp. 225–231. Morgan Kaufmann, San Mateo (1992)
22. Stane, S.A.B., Zytniewsk, M.: Normative multi-agent enriched data mining to support e-citizens. In: Cao, L. (ed.) Data Mining and Multi-Agent Integration, pp. 291–304. Springer, Heidelberg (2009)
23. Verhagen, H.: Simulation of the Learning of Norms. Social Science Computer Review 19(3), 296–306 (2001)
24. Villatoro, D., Sabater-Mir, J., Sen, S.: Interaction, observance or both? study of the effects on convention emergence. In: Proceeding of the 2009 Conference on Artificial Intelligence Research and Development: Proceedings of the 12th International Conference of the Catalan Association for Artificial Intelligence, pp. 189–196. IOS Press (2009)
25. Walker, A., Wooldridge, M.: Understanding the emergence of conventions in multi-agent systems. In: Lesser, V. (ed.) Proceedings of the First International Conference on Multi-Agent Systems (ICMAS), pp. 384–389. AAAI Press, Menlo Park (1995)
26. Watkins, C.J.C.H., Dayan, P.: Q-learning. Machine Learning 8(3-4), 279–292 (1992)