Identifying Conditional Norms in Multi-agent Societies

Bastin Tony Roy Savarimuthu, Stephen Cranefield, Maryam A. Purvis, and Martin K. Purvis

University of Otago, Dunedin, P.O. Box 56, Dunedin, New Zealand
(tonyr,scranefield,tehrany,mpurvis)@infoscience.otago.ac.nz

Abstract. Most works on norms have investigated how norms are regulated using institutional mechanisms which assume that agents know the norms of the society they are situated in. Few research works have focused on how an agent may infer the norms of a society without the norm being explicitly given to the agent. These works do not address how an agent can identify conditional norms. In this paper we describe a mechanism that an agent can use to identify conditional norms which makes use of our previously proposed norm identification framework. Using park littering as an example, we show how conditional norms can be identified. In addition, we discuss the experimental results on the dynamic addition, modification and deletion of conditional norms.

1 Introduction

Most works on norms in normative multi-agent systems have concentrated on how norms regulate behaviour (e.g. [14, 18]). These works assume that the agent somehow knows (*a priori*) what the norms of a society are. For example, an agent may have obtained the norm from a leader [7] or through an institution that prescribes what the norms of the society should be [1, 31].

Only a few researchers have dealt with how an agent may infer what the norms of a newly joined society are [2, 23]. Recognizing the norms of a society is beneficial to an agent. This process enables the agent to know what the *normative expectation* of a society is. As the agent joins and leaves different agent societies, this capability is essential for the agent to modify its expectations of behaviour, depending upon the society of which it is a part. As the environment changes, the capability of recognizing a new norm helps an agent to derive new ways of achieving its intended goals. Such a norm identification mechanism can be useful for software agents that need to adapt to a changing environment. In open agent systems, instead of possessing predetermined notions of what the norms are, agents can infer and identify norms through observing patterns of interactions and their consequences. For example, a new agent joining a virtual environment such as Second Life [21] may have to infer norms when joining a society as each society may have different norms. It has been noted that having social norms centrally imposed by the land owners in Second Life is ineffective and there is a need for the establishment of community driven norms [29]. When a community of agents determines what the norm should be, the norm can evolve over time. So, a new agent joining the society should have the ability to recognize the changes to the norms. In our previous work we have proposed and experimented with a norm identification

M. De Vos et al. (Eds.): COIN 2010 International Workshops, LNAI 6541, pp. 285–302, 2011.

framework which can be used to identify norms in the society [23–25]. The norm identification framework takes into account the social learning theory [5] that suggests that new behaviour can be learnt through the observation of punishment and rewards. This work aims to answer the question of how agents infer conditional norms in a multi-agent society. Conditional norms are defined as norms with conditions. We distinguish norms that are not associated with conditions from the ones that have conditions. An example of a norm without a condition is the norm that prohibits anyone from littering a public park, i.e. *prohibit(litter)*. An example of a norm with condition is a norm that prohibits one from littering as long as there is a rubbish bin within x metres from the agent (e.g. *if (distanceFromBin* $<$ *10) then prohibit(litter)*). Software agents should not only have the ability to identify norms but also the conditions under which these norms hold.

Identifying conditional norms is important because an agent that has inferred that another agent gets punished when that agent littered when it was 25 metres away from the bin may infer that the condition associated with the norm is the distance of 25 metres. But the actual norm could be that no one should litter within 50 metres from the bin. The utility of the agent can be negatively impacted through a sanction if it litters 30 metres away from a bin. In this case, the agent does not know the correct condition associated with the norm. Another example of a conditional norm is the tipping norm. In one society an agent may tip 10% of the bill while in another society an agent might be obliged to tip 20% of the bill. In this work we are interested in experimenting with the formation, modification and the removal of conditional norms in the minds of the agents and the impact of the normative conditions on the utility of the agents.

The paper is organized as follows. Section 2 provides a background on norms and how the concept of norms is investigated in the field of normative multi-agent systems (NorMAS). Section 3 provides an overview of our previous work on the norm identification framework. Section 4 describes a mechanism for identifying conditional norms. Section 5 describes the experiments that we have conducted and the results obtained. Section 6 provides a discussion on the work that has been achieved and the issues that can be addressed in the future. Concluding remarks are presented in Section 7.

2 Background

Due to multi-disciplinary interest in norms, several definitions for norms exist [2]. Elster notes the following about social norms [11]. *"For norms to be social, they must be shared by other people and partly sustained by their approval and disapproval. They are sustained by the feelings of embarrassment,anxiety, guilt and shame that a person suffers at the prospect of violating them. A person obeying a norm may also be propelled by positive emotions like anger and indignation ... social norms have a grip on the mind that is due to the strong emotions they can trigger"*.

Based on the definitions provided by various researchers, we note that the social practices surrounding the notion of a social norm are the following:

- *The normative expectation of a behavioural regularity*: There is a general agreement within the society that a behaviour is *expected* on the part of an agent (or actor) by others in a society, in a given circumstance.

- *A norm enforcement mechanism*: When an agent does not follow a norm, it could be subjected to a sanction. The sanction could include monetary or physical punishment in the real world which can trigger emotions (embarrassment, guilt, etc.) or direct loss of utility (e.g. decrease of its reputation score).
- *A norm spreading mechanism*: Examples of norm spreading mechanisms include the notion of advice from powerful leaders, imitation and learning on the part of an agent.

2.1 Normative Multi-agent Systems

The definition of normative multi-agent systems given by the researchers involved in the NorMAS 2007 workshop is as follows [6]. *A normative multi-agent system is a multi-agent system organized by means of mechanisms to represent, communicate, distribute, detect, create, modify and enforce norms, and mechanisms to deliberate about norms and detect norm violation and fulfillment.*

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 [28, 33]. Epstein has shown that norms reduce the amount of computation required to make a decision [12]. 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 and adapt in response to environmental changes.

Research in normative multi-agent systems can be categorized into two branches. 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 [19] for an overview). Researchers have used deontic logic to define and represent norms [16, 35]. Several researchers have worked on mechanisms for norm compliance and enforcement [1, 18].

The second branch of research is related to emergence of norms [15, 27, 28]. Researchers have worked on both prescriptive (top-down) and emergent (bottom-up) approaches to norms. In a top-down approach an authoritative leader or a normative advisor prescribes what a norm of the society should be [32]. In the bottom-up approach, the agents come up with a norm through learning mechanisms [27, 28]. Researchers have used sanctioning mechanisms [4] and reputation mechanisms [10] for enforcing norms.

Many research works assume that norms exist in the society and the focus is on how the norms can be regulated in an institutional setting such as electronic institutions[3]. Very few have investigated how an agent comes to know the norms of the society [2, 23]. We have previously proposed an architecture for norm identification [23, 25]. In this work, we extend our earlier work by incorporating the mechanism for identifying conditional norms. Identifying conditional norms is important because the agent can confidently apply the norm if the conditions associated with the norm are known. This will help the agent not to lose utility by preventing it from applying the norm under wrong conditions.

3 Overview of the Norm Identification Framework

In this section, we provide an brief overview of the norm identification framework that we have proposed and experimented with in earlier works [23–25]. An agent employing this architecture follows a four-step process.

- Step 1: An agent actively perceives the events in the environment in which it is situated.
- Step 2: When an agent perceives an event, it stores the event in its belief base.
- Step 3: Based on recognizing signals (i.e. events that are either rewards or a sanctions), the agent stores them in a "special events" base.
- Step 4: If the perceived event is a special event an agent checks if there exists a norm in its *personal norm* (*p-norm*) base or the *group norm* (*g-norm*) base. An agent may possess some p-norms. A *p-norm* is the personal value of an agent. For example an agent may consider that littering is an action that should be prohibited in a society based on its past experience or preference. This personal value may not be shared by the agents in a society. A *p-norm* may vary across agents, since a society may be made up of agents with different backgrounds and experiences. A *g-norm* is a norm which an agent infers, based on its personnel interactions as well as the interactions it observes in the society. An agent infers g-norms using the norm inference component. The norm inference component of the framework [25] makes use of Candidate Norm Inference (CNI) algorithm. The CNI algorithm uses association rule mining approach to identify sequences of events as candidate norms. The CNI algorithm has two sub-modules to identify prohibition norms [25] and obligation norms [24] respectively.

When a special event occurs an agent may decide to invoke its norm inference component to identify whether a previously unknown norm may have resulted in the occurrence of the special event. In the context of the park-littering scenario, an agent observing a sanctioning event may invoke its norm inference component to find out what events that had happened in the past (or that had not happened in the past) may have triggered the occurrence of the special event. Prohibition norms can be identified by inferring the relevant events that happened in the past [25]. For example an agent may notice that a sanctioning event is always preceded by a littering event. Hence the agent might infer that littering is prohibited in the society. For identifying obligation norms the agent may have to reason about what events that did not happen in the past are the likely reason for a sanction (i.e. not fulfilling an obligation) [24]. For example, an agent may be sanctioned for not tipping a customer in a restaurant. An agent observing the events may infer that the absence of the tipping action is the reason for a sanction. In this work we focus on identifying conditional norms associated with prohibition norms.

The invocation of the norm inference component may result in the identification of a *g-norm*, in which case it is added to the *g-norm* base. An agent, being an autonomous entity, can also decide not to invoke its norm inference component for every occurrence of a special event but may decide to invoke it periodically. When it invokes the norm inference component, it may find a new *g-norm* which it adds to its *g-norm* base. If it does not find a *g-norm*, the agent may change some of its norm inference parameters and repeat the process again in order to find a *g-norm* or may wait to collect more information.

At regular intervals of time an agent re-evaluates the g-norms it currently has, to check whether those norms hold. When it finds that a *g-norm* does not apply (e.g. if it does not find any evidence of sanctions), it deletes the norm from the *g-norm* base. It could be that there are no punishers in the society or all the agents have internalized the norm and are following the norm. Hence, there might be no sanctions in the society. In the case where all agents have internalized the norm and are following the norm, norm deletion on the part of the observer agent may have negative consequence for that agent (i.e. the agent can be sanctioned) in which case it can add the norm again through norm inference.

The next section describes how conditional norms are inferred by an agent. The mechanism for identifying conditional norms is built on top of the norm inference framework.

4 Identifying Conditional Norms

In our framework, when a new agent enters a society it will try to identify the norms that currently hold in that society. Once an agent has identified a norm it may want to identify the context and the exact conditions under which the norm holds. For example, the norm in a public park could be not to litter, i.e. *prohibit (litter)*. It could be that the norm prohibits people from dropping litter in the park as long as a rubbish bin is visible to them (or the rubbish bin is 50 metres away from them). The context here is the rubbish bin and the condition is the distance from the rubbish bin. When an agent identifies the norm in the first instance through observation, it may not know the exact conditions associated with the norm.

Let us assume that an agent upon identifying the norm knows the context of the norm[1]. For example on identifying that littering is prohibited, the agent identifies the presence of the bin as the context. The condition associated with the norm is the distance between the agent and the bin[2]. We call this a contextual condition.

Note that the condition associated with a norm will be specific to the domain under consideration. In the park littering example, the condition can be either one or two-dimensional. For example, the distance between a littering agent and bin is a single dimensional entity. The littering zone can be modelled as a two dimensional entity if it is defined using x and y coordinates (i.e. an agent should not litter within 5 metres from bin's x position and 10 metres from bin's y position). Some researchers have used a two dimensional representation for normative conditions [17, 30]. In this work we have used the distance metric which we call the radius of the non-littering zone.

[1] We assume that an agent knows the context based on the past experience or based on the domain knowledge. For example, an agent may know about littering from its past experience.

[2] Proximity or the distance of interaction is a contextual condition in many social norms. For example, two people talking tend to speak in a low voice when they walk past others. Another example is the interpersonal distance norm (i.e. how close you can get to someone while holding a conversation without making him/her uncomfortable). Agents may be aware of the distance based contextual condition from their previous experience.

Algorithm 1. Pseudocode of an agent to identify conditional norm associated with the prohibition of littering

```
   Input: Contextual Condition = distance from nearest rubbish bin
1  maxDistanceFromBin = 0, tempDistance = 0 ; /* maxDistanceFromBin
   stores the value of the contextual condition */
2  conditionalNormReferralConsidered = true;
3  conditionalNormRecommenders = ∅;
4  foreach norm inference cycle do
5      Obtain Norms Set (NS) ; /* By invoking Candidate Norm
       Identification algorithm */
6      if NS ≠ ∅ then
7          foreach norm in NS do
8              foreach punished agent with the visibility threshold do
9                  tempDistance = getDistanceFromNearestBin;
10                 if tempDistance > maxDistanceFromBin then
11                     maxDistanceFromBin = tempDistance;
12                 end
13             end
14             if conditionalNormReferralConsidered then
15                 conditionalNormRecommenders =
                   getAgentsFromVicinity;
16                 foreach agent ∈ conditionalNormRecommenders do
17                     if agent.maxDistanceFromBin > maxDistanceFromBin then
18                         maxDistanceFromBin = agent.maxDistanceFromBin;
19                     end
20                 end
21             end
22         end
23     end
24 end
```

Algorithm 1 shows how an agent identifies the conditional norm of the park. In each norm inference cycle an agent will first identify a set of norms using the norm identification framework [25]. Let us assume that the agent has identified *prohibit(litter)* as the norm which is stored in its Norms Set (NS). For each of the littering agents that were observed to be punished, an agent calculates the distance from the nearest bin to the punished agent using Chebyshev's distance metric [36][3]. The agent finds the radius of the non-littering zone (lines 10-12) and stores it in *maxDistanceFromBin*. The agent can

[3] Chebyshev's distance also known as the Chessboard distance is the minimum number of steps required for a King to move from one square of the chessboard to another. In our implementation Chebyshev distance represents the minimum distance between an agent and the nearest bin. Chebyshev distance of length one corresponds to the Moore neighbourhood [34] of size one where an agent in one cell can see all the 8 cells surrounding it.

choose to ask for referral from one or more agents from its vicinity threshold regarding the zone in which littering is prohibited (i.e. *maxDistanceFromBin*). If the referrer's recommended distance is greater than distance observed by the agent the agent increases the distance (lines 14-21).

While Algorithm 1 is specific to the park littering scenario, the generic process of an agent to identify the conditional norm is given in Algorithm 2. Once the agent infers a norm, it will identify the contextual condition. The contextual condition can contain multi-dimensional attributes. For each norm in the norm set (NS), it calculates the value of the contextual condition (line 8). An agent calculates the value for contextual condition based on observing all the punished agents within its visibility threshold.

Algorithm 2. Pseudocode of an agent to identify a conditional norm

```
   Input: Contextual Conditions
1  valueOfContextualCondition [] = ∅;
2  conditionalNormReferralConsidered = true;
3  conditionalNormRecommenders = ∅;
4  foreach norm inference cycle do
5      Obtain Norms Set (NS) ; /* By invoking Candidate Norm
       Identification algorithm                                  */
6      if NS ≠ ∅ then
7          foreach norm n in NS do
8              valueOfContextualCondition [n] =
               calculateContextualConditionalValue; /* This
               is calculated based on the available data on
               all punished agents within the visibility
               threshold                                         */
9              if conditionalNormReferralConsidered then
10                 conditionalNormRecommenders =
                   getAgentsFromVicinity;
11                 foreach agent ∈ conditionalNormRecommenders do
12                     if agent.valueOfContextualCondition is better than
                       valueOfContextualCondition then
13                         valueOfContextualCondition =
                           agent.valueOfContextualCondition;
14                     end
15                 end
16             end
17         end
18     end
19 end
```

The observer agent can optionally ask recommendation from other agents (through referral), on the contextual condition that they have observed (lines 9 and 10). Then, based on the recommendation of other agents it can choose the best value[4] as its value for the contextual condition (lines 11-15).

5 Experiments

In this section we firstly describe the experimental set-up in sub-section 5.1. In the rest of the sub-sections we describe the experiments that were conducted and the results obtained.

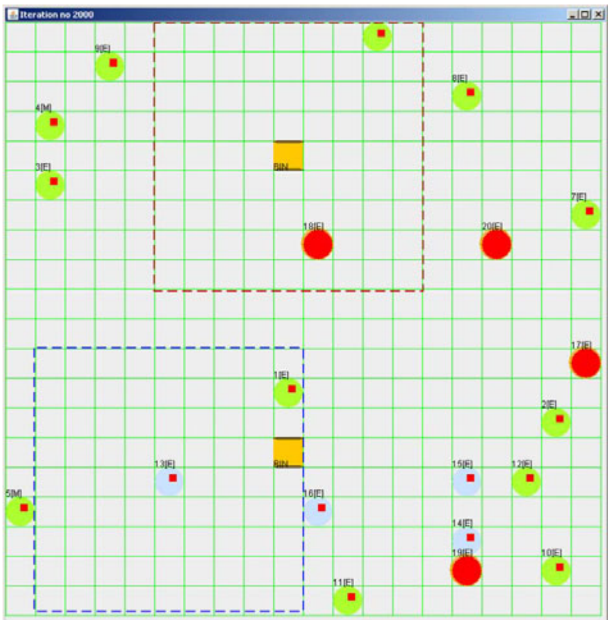

Fig. 1. Snapshot of the simulation of conditional norms

5.1 Experimental Set-Up

We model agents in our virtual society in a two-dimensional space. This virtual environment can be considered as a communal region such as a park shown in Figure 1. The agents explore and enjoy the park by moving around. There are three types of agents in the simulation. They are learning litterers (LL), non-litterers (NL) and non-littering punishers (NLP). There are four possible types of actions defined in the simulation system: *move, eat, litter* and *punish*. The LL agents can *move, eat* and *litter*. The NL agents can *move* and *eat* while the NLP agents can *move, eat* and *punish*. The agents'

[4] The logic for choosing the best value of the condition is domain specific. In a domain the best value may correspond to the least numeric value and in another domain it may correspond to the highest value.

movement can be in one of the four directions: up, down, left or right. The agents that are at the edge of the two dimensional space can again re-appear in the opposite side (i.e. a toroidal grid is implemented). The agents are represented as circles using different colours. The LLs are green, the NLs are blue and the NLPs are red. The id and action that an agent currently does appear above the circle.

Each agent has a visibility threshold. The visibility threshold of the agent is governed by a Chebyshev distance [36] of a certain length. An agent can observe actions of agents and the interactions that happen between two agents within its visibility threshold. The dashed square that appears at the bottom of Figure 1 shows the visibility range of agent 13 which is at the centre of the dashed square with a Chebyshev distance of four. All the agents make use of the norm inference component [25] to infer norms. The red squares that appear within the circles represent the identification of a norm. Rubbish bins in the simulation environment appear in orange. The non-littering zone with reference to the bin at the top is given by the dashed square that appears at the top of Figure 1. The radius of the non-littering zone in this case is four.

The simulation parameters that were kept constant for all the experiments are given in Table 1. A sample simulation can be viewed from this link[5].

Table 1. Simulation parameters for identifying conditional norms

Parameters	Values
Grid size	20*20
Total number of agents	20
Number of litterers	12
Number of non-litterers	4
Number of non-littering punishers	4
Visibility threshold	5
Number of rubbish bins	2
radius of non-littering zone (maxDistanceFromBin)	10
Number of referrals (when used)	1

5.2 Experiment 1 - Conditional Norm Identification

The objective of the first experiment is to show that agents in a society infer conditional norms using the proposed mechanism. We also compare the rate of norm establishment in the society with the rate of conditional norm establishment in the society.

Figure 2 shows two lines that represent the proportion of agents with norms and the proportion of agents with conditional norms in a society respectively. It can be seen from Figure 2 that even though the norm has been established in the society[6] in iteration 270, the conditional norm (i.e. the agent should not litter when it is within 10 metres from the bin), is not inferred in the society till iteration number 410. This is because the conditional norm identification process is invoked by an agent after it has found a

[5] http://unitube.otago.ac.nz/view?m=iWs217vmy6H

[6] We assume that a norm is established in a society if all the agents (100%) have inferred the norm. Researchers have used different criteria ranging from 35% to 100%[22].

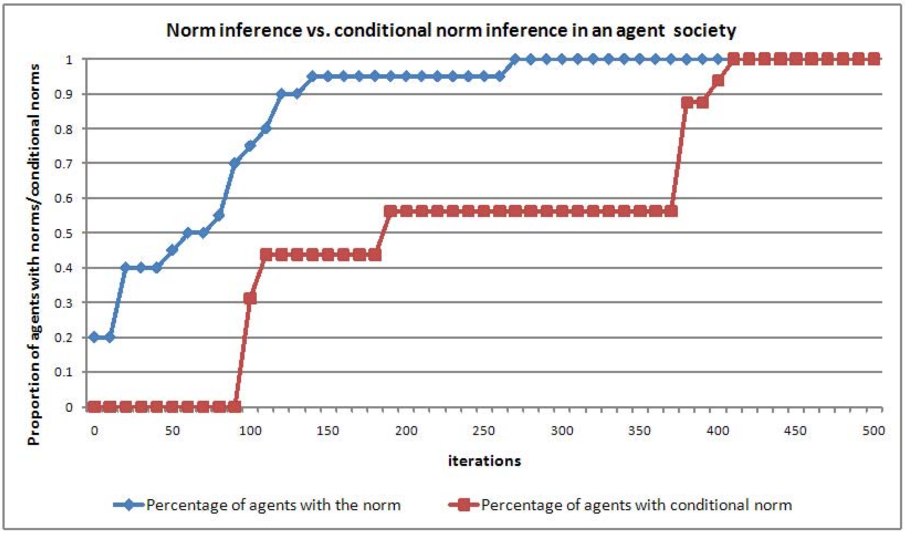

Fig. 2. Conditional norm identification

norm. As the agents interact more and more in the society, they gather more evidence regarding the condition associated with the norm. If the norm does not change, then the correct condition associated with the norm is inferred eventually. When an agent does not infer a norm for certain amount of time or when the norm changes it will remove the norm and its associated conditions from its norms base.

5.3 Experiment 2 - Conditional Norm Identification with and without Referral

An agent can expedite the process of identifying a conditional norm if it asks another agent for its evidence of the normative condition. We call this as the conditional norm referral process. It can be observed from Figure 3 that when the referral is used, the rate of establishment of the conditional norm increases. The agents ask for referral from one other agent in the society. When the number of referees increases, the rate of conditional norm establishment increases. This has also been reported many other works in multi-agent systems[8, 37].

Figure 4 shows the progression of two agents towards the identification of the correct conditional norm (non-littering zone of radius 10) with and without referrals. The progression rates of the two agents are different because of their different paths of travel. If an agent observes more agents on its path, then it has a higher probability of inferring both the norm and the condition associated with the norm. It should be noted that the conditional norm establishment for these two agents improve when the referrals are used.

The two dashed lines in Figure 4 show the radius of the non-littering zone identified by the agents during the simulation. The agent which found the norm first (agent 1, iteration 90) was not the one to find the correct conditional norm first[7]. When agent

[7] The correct conditional norm is the non-littering zone of 10 metres.

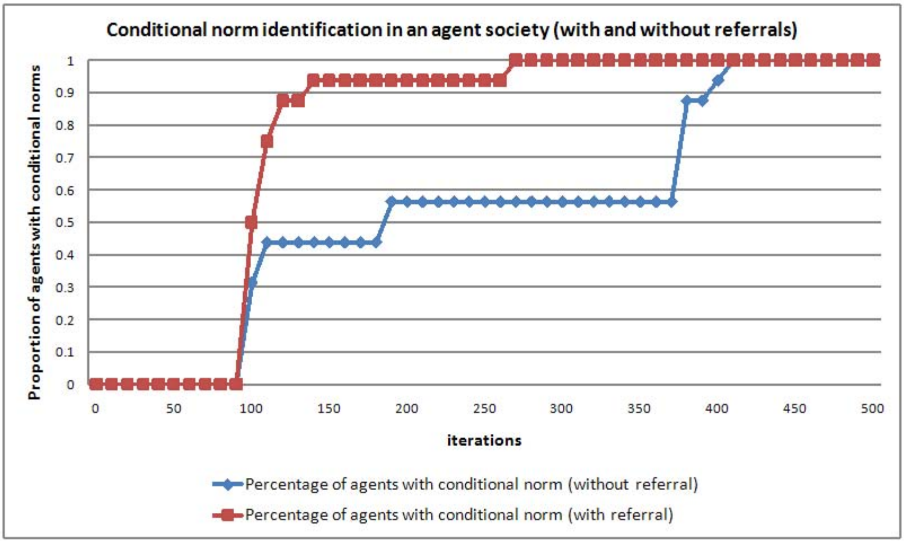

Fig. 3. Rate of norm and conditional norm establishment in an agent society

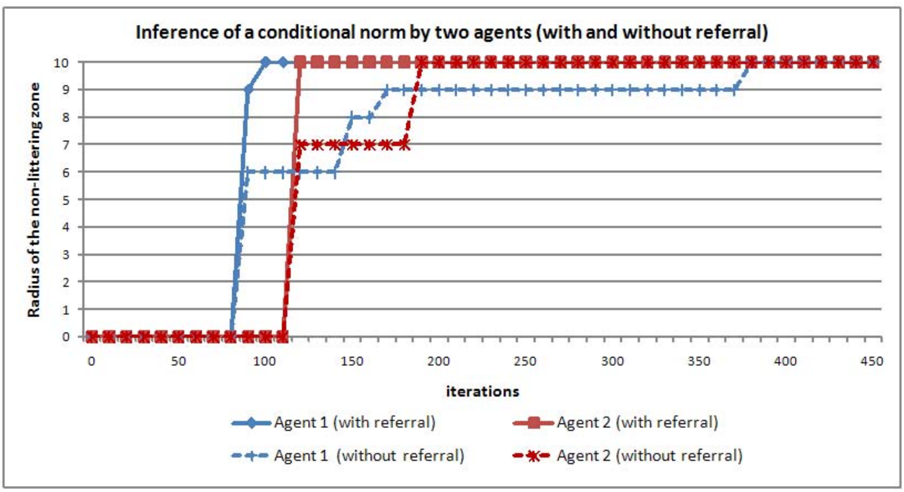

Fig. 4. Rate of conditional norm establishment in two agents with and without referrals

1 found the norm in iteration 90, the non-littering zone identified by the agent was 6 metres (shown using an arrow in the Figure). It found the correct conditional norm in iteration 380. Agent 2, albeit finding the norm second (iteration 110, non-littering zone of radius 7 metres), found the correct conditional norm faster (iteration 190). This again is governed by the number of agents an agent gets to observe (i.e. the path of travel).

The two solid lines show the radius of the non-littering zone identified by the agents when referrals are used. It is interesting to note that when the referral mechanism is used, the agent which found the norm first was also the one that found the normative condition first. This is because once the agent finds the norm it can ask the agents in the vicinity for referral instead of waiting for a long amount of time to find out the maximum distance from the bin from which a violation that was punished occurred.

5.4 Experiment 3 - Dynamic Conditional Norm Identification

An agent should have the ability to dynamically add newly identified norms and remove norms that do not hold. This experiment demonstrates that conditional norms can be added, removed and modified by an agent dynamically depending upon the environmental conditions. The ability to change norms is important for an adaptive agent so that it can flexibly adopt norms. An agent, on identifying a norm, evaluates whether the norm holds at regular intervals of time. If the norm does not hold, it removes the norm from its norm base. When it removes the norm it also removes the condition associated with the norm[8].

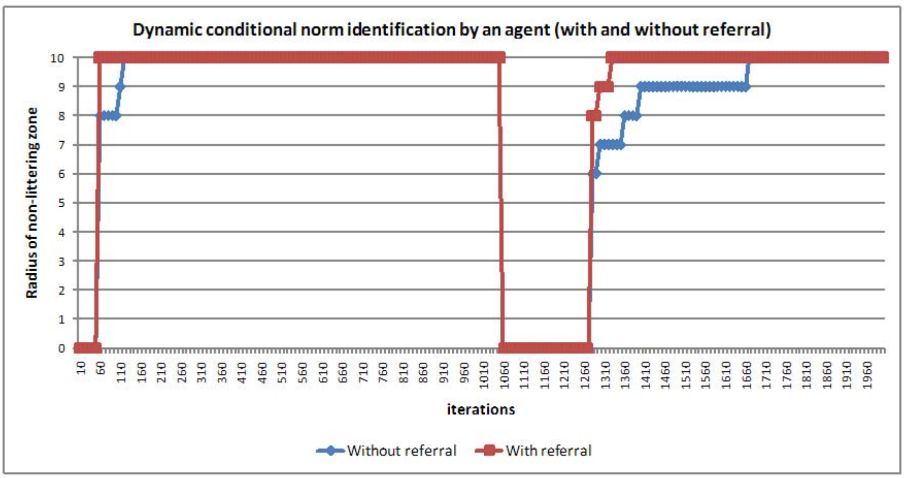

Fig. 5. Dynamic conditional norm identification by an agent

Figure 5 demonstrates that an agent is able to add and remove norms and normative conditions dynamically. Figure 6 demonstrates that agents in our system are able dynamically modify the normative condition. In these experiments, the punishers do not punish from iterations 1000 to 1250. This is to simulate the change in the environment which triggers a norm change. Additionally, having identified a norm, an agent checks for the validity of the norm once again after 5 norm inference cycles (norm inference happens once every 10 iterations). If the norm is found again, then the agent does not delete the norm. If the norm is not found, it removes the norm and the conditions from its norm base.

[8] In our previous work [25], we have demonstrated how an agent adds and removes norms dynamically. In this experiment, we show how conditions associated with norms are dynamically added and removed.

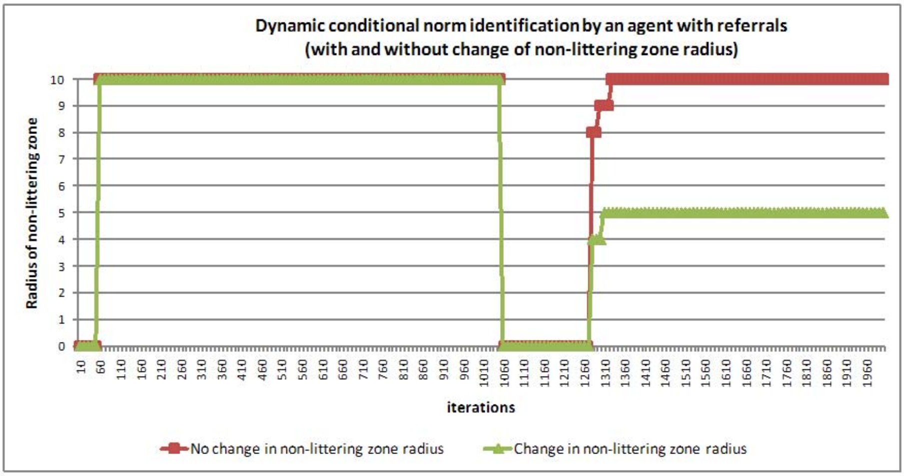

Fig. 6. Dynamic conditional norm identification by an agent

Figure 5 shows two lines that represent an agent adding and removing norms based on changing environmental conditions. The line is red represents the agent using the referral mechanism and the line in blue represents the agent without using the referral mechanism. It can be observed that the agent without using referral identifies a conditional norm in iteration 60, and the correct conditional norm in iteration 120 while it infers the norm faster when it uses referral. In this experiment, when the punishers do not punish, the norm is not inferred for 50 iterations (5 norm inference cycles from iteration 1010 to 1050). So, the agent removes the norm and the conditions associated with the norm (with and without referral) in iteration 1050. The agent that does not use the referral finds a conditional norm again in iteration 1280 and the correct conditional norm in iteration 1670. It can be observed that when the referral is used by the agent it identifies the correct conditional norm earlier (iteration 1330).

Figure 6 shows two lines that represent the identification of different normative conditions under changing environmental conditions (with and without change of non-littering zone) for the same agent. By keeping all the other parameters the same, we varied the radius of the non-littering zone (i.e. the punishment zone for littering). This is to demonstrate that when the radius of the littering zone varies, the agent infers the change. After iteration 1250 all NLP agents punished only those agents that littered within 5 metres from the bin (as opposed to 10 metres which was used in iterations 1 to 1000). It can be observed from the green line in Figure 6 that the agent inferred the new normative condition (radius = 5)[9]. Note that the agent has made use of referral in this experiment. The red line which converges to the littering zone of radius 10 that appears at the top of Figure 6 is the same as the red line shown at the top of Figure 5 which represents the normative behaviour of an agent that uses the referral process.

[9] The simulation video can be found at http://unitube.otago.ac.nz/view?m=nQ6y17frCcJ

5.5 Experiment 4 - Comparison of Utility of Agents with and without Conditional Norm Identification

The objective of this experiment is to compare the utility benefits of an agent when it identifies norms with and without conditions. In this experiment, an agent has a utility value which we call the satisfaction level (S) which varies from 0 to 100.

An agent's satisfaction level (S) decreases in the following situations:

- When a litterer is punished, its utility decreases (-1).
- For all agents, littering activity results in the decrease of the utility. This is because each littering activity ruins the "commons" area (-1/number of agents in the society).

An agent's satisfaction level (S) increases (i.e. it gains utility) in the following situation:

- When a litterer litters, it gains utility in a society (+1).

We have experimented with the utility of the agent with and without conditional norm identification. An LL agent is better off by using conditional norm (CN) identification. Once identifying a norm an LL agent may choose to abstain from the action that is being prohibited. In the case of a conditional norm, it learns the exact condition under which it should not violate the norm. By this process, it can improve its utility. It can be observed from Figure 7 that an LL agent's utility increases (greater than 50) when it has identified the conditional norm than just identifying the norm without conditions (less than 50). This is because when the littering agent finds the norm without the normative condition, it abstains from the littering activity which does not lead to an increase in its utility. But, when it identifies the normative condition, it now can litter outside the non-littering zone which results in the increase in its utility.

For an NL agent, when it identified a norm without identifying the condition, the utility initially decreases but then stabilizes to a constant value because when all the agents inferred the norm, there aren't any littering actions in the society. When the NL agent identifies the conditional norm, its utility continues to decrease because whenever an LL agent litters outside the not-to-litter zone, its utility decreases[10]. Similarly, for an LL agent, its utility decreases because of the littering action. However, its net utility increases as it gains from the littering action (i.e. it can litter outside the non-littering zone).

The utilities of NLP agents are not discussed here because we assume these agents have other utility functions for punishing (e.g. a leader who wants to promote a smoother functioning of the society, or an altruistic agent who does not care about its diminishing utility). We note that if non-altruistic punishers are present in the society, then the cost

[10] It should be noted that when the utility of an NL agent goes below a certain threshold, it can leave the society, or can become a litterer or become a punisher. This is explored in another work [26]. Additionally, if the parameters of this experiment are varied (for example if the utility gain of a litterer is changed to 0.5 and the utility loss on receiving a punishment is changed to 0.25) the results obtained will be different. The objective here is to show that the conditional norm identification has an impact on the utility of the agents.

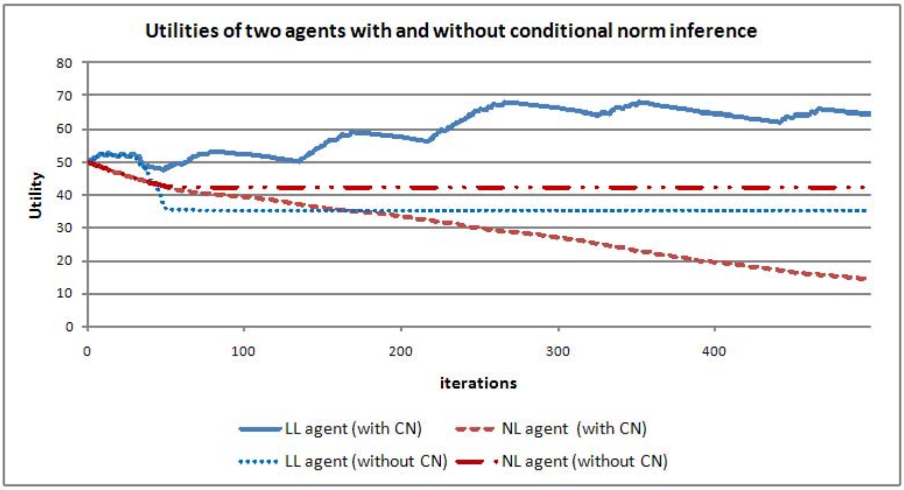

Fig. 7. Utility comparison of two agents

incurred by the non-altruistic punishers will play a role in the establishment of a norm in the society (see [26]). Several other works have investigated the role of punishment costs on norm spreading [13, 20].

6 Discussion

The issue of conditional norm identification has not been dealt with by researchers in the field of normative multi-agent systems. To this end, we have experimented how a conditional norm can be identified by an agent in the context of park-littering. Identifying norms with conditions can be beneficial in several settings. For example, the norm identification architecture can be used to infer norms in Massively Multi-player Online Games (MMOGs). Players involved in massively multi-player games perform actions in an environment to achieve a goal. They may play as individuals or in groups. When playing a cooperation game (e.g. players forming groups to slay a dragon), individual players may be able to observe norms. For example, a dragon can only be slayed if two players are within certain distance from the dragon. An agent that identifies this condition (the distance) will be better-off than an agent that just infers the norm of co-operation (i.e. two players are needed to slay a dragon). The mechanism proposed in this paper can be used to identify norms with conditions. This mechanism can also be used in virtual environments such as Second Life to infer conditional norms.

Another application of identifying conditional norms is in the area of e-commerce. For example, in one society, the norm associated with the deadline for payment (i.e. obligations with deadlines as in [9]) may be set to 120 minutes after winning the item. Depending upon what an agent has observed, agents may have subtly different norms (e.g. one agent may notice that *pay* follows *win* after an average of 80 minutes while another may notice this happens after 100 minutes). Both these agents could still infer the norm but the deadlines they had noticed can be different. This may result in an unstable equilibrium with reference to the norms and hence conflict resolution mechanisms should be used to resolve them [17, 30].

We note that we have modelled and experimented with a simple domain. The number and type of agents can easily be increased and the normative conditions identified can be richer and more complex depending upon the problem domain. However, we believe the main contribution is the mechanism for the identification of conditions associated with norms. We have also shown how an agent can dynamically add, remove and modify conditions associated with the norms.

7 Conclusion

This paper addresses the question of how conditional norms can be identified in an agent society using the norm inference architecture. Identification of conditional norms has been demonstrated in the context of a simple park-littering scenario. The ability of an agent to add, delete and modify a conditional norm has also been demonstrated. It has also been shown that identifying norms with conditions has an impact on the utility of the agents in the society.

References

[1] Aldewereld, H., Dignum, F., García-Camino, A., Noriega, P., Rodríguez-Aguilar, J.A., Sierra, C.: Operationalisation of norms for usage in electronic institutions. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2006), pp. 223–225. ACM Press, New York (2006)

[2] 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. Dagstuhl Seminar Proceedings, vol. 07122, Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany (2007)

[3] Arcos, J.L., Esteva, M., Noriega, P., Rodrguez-aguilar, J.A., Sierra, C.: Environment engineering for multiagent systems. Engineering Applications of Artificial Intelligence 18(2), 191–204 (2005)

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

[5] Bandura, A.: Social Learning Theory. General Learning Press (1977)

[6] Boella, G., van der Torre, L., Verhagen, H.: Introduction to the special issue on normative multiagent systems. Autonomous Agents and Multi-Agent Systems 17(1), 1–10 (2008)

[7] Boman, M.: Norms in artificial decision making. Artificial Intelligence and Law 7(1), 17–35 (1999)

[8] Candale, T., Sen, S.: Effect of referrals on convergence to satisficing distributions. In: Proceedings of the Fourth International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2005), pp. 347–354. ACM, New York (2005)

[9] Cardoso, H.L., Oliveira, E.C.: Directed deadline obligations in agent-based business contracts. In: Padget, J.A., Artikis, A., Vasconcelos, W.W., Stathis, K., da Silva, V.T., Matson, E.T., Polleres, A. (eds.) COIN@AAMAS 2009. LNCS, vol. 6069, pp. 225–240. Springer, Heidelberg (2010)

[10] Castelfranchi, C., Conte, R., Paolucci, M.: Normative reputation and the costs of compliance. Journal of Artificial Societies and Social Simulation 1(3) (1998)

[11] Elster, J.: Social norms and economic theory. The Journal of Economic Perspectives 3(4), 99–117 (1989)

[12] Epstein, J.M.: Learning to be thoughtless: Social norms and individual computation. Computational Economics 18(1), 9–24 (2001)

[13] Fehr, E., Fischbacher, U.: Third-party punishment and social norms. Evolution and Human Behavior 25(2), 63–87 (2004)

[14] Gaertner, D., García-Camino, A., Noriega, P., Rodrguez-aguilar, J.A., Vasconcelos, W.W.: Distributed norm management in regulated multi-agent systems. In: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007), Honolulu, Hawaii, pp. 624–631. ACM, New York (2007)

[15] Griffiths, N., Luck, M.: Norm diversity and emergence in tag-based cooperation. In: De Vos, M., et al. (eds.) COIN 2010. LNCS (LNAI), vol. 6541, pp. 230–249. Springer, Heidelberg (2011)

[16] 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, Chichester (1993)

[17] Kollingbaum, M.J., Vasconcelos, W.W., García-Camino, A., Norman, T.J.: Managing conflict resolution in norm-regulated environments. In: Artikis, A., O'Hare, G.M.P., Stathis, K., Vouros, G.A. (eds.) ESAW 2007. LNCS (LNAI), vol. 4995, pp. 55–71. Springer, Heidelberg (2008)

[18] López y López, F.: Social Powers and Norms: Impact on Agent Behaviour. PhD thesis, Department of Electronics and Computer Science, University of Southampton, United Kingdom (2003)

[19] Neumann, M.: A classification of normative architectures. In: Proceedings of World Congress on Social Simulation (2008)

[20] Ohtsuki, H., Iwasa, I., Nowak, M.A.: Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. Nature 457(7225), 79–82 (2009)

[21] Rymaszewski, M., Au, W.J., Wallace, M., Winters, C., Ondrejka, C., Batstone-Cunningham, B., Rosedale, P.: Second Life: The Official Guide. SYBEX Inc., Alameda (2006)

[22] Savarimuthu, B.T.R., Cranefield, S.: A categorization of simulation works on norms. In: Boella, G., Noriega, P., Pigozzi, G., Verhagen, H. (eds.) Normative Multi-Agent Systems, Dagstuhl, Germany. Dagstuhl Seminar Proceedings, (09121), pp. 43–62. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany (2009)

[23] Savarimuthu, B.T.R., Cranefield, S., Purvis, M.A., Purvis, M.K.: Internal agent architecture for norm identification. In: Padget, J., Artikis, A., Vasconcelos, W., Stathis, K., da Silva, V.T., Matson, E., Polleres, A. (eds.) COIN@AAMAS 2009. LNCS, vol. 6069, pp. 241–256. Springer, Heidelberg (2010)

[24] Savarimuthu, B.T.R., Cranefield, S., Purvis, M., Purvis, M.: A data mining approach to identify obligation norms in agent societies. In: Cao, L., Bazzan, A., Gorodetsky, V., Mitkas, P., Weiss, G., Yu, P. (eds.) ADMI 2010. LNCS, vol. 5980, pp. 43–58. Springer, Heidelberg (2010)

[25] Savarimuthu, B.T.R., Purvis, M.A., Purvis, M.K., Cranefield, S.: Norm identification in multi-agent societies. Discussion Paper 2010/03, Department of Information Science, University of Otago (2010)

[26] Savarimuthu, B.T.R., Purvis, M., Purvis, M.K., Cranefield, S.: Social norm emergence in virtual agent societies. In: Baldoni, M., Son, T.C., van Riemsdijk, M.B., Winikoff, M. (eds.) DALT 2008. LNCS (LNAI), vol. 5397, pp. 18–28. Springer, Heidelberg (2009)

[27] Sen, S., Airiau, S.: Emergence of norms through social learning. In: Proceedings of Twentieth International Joint Conference on Artificial Intelligence (IJCAI 2007), pp. 1507–1512. AAAI Press, Menlo Park (2007)

[28] Shoham, Y., Tennenholtz, M.: Emergent conventions in multi-agent systems: Initial experimental results and observations. In: Proceedings of Third International Conference on Principles of Knowledge Representation and Reasoning, San Mateo, CA, pp. 225–231. Morgan Kaufmann, San Francisco (1992)

[29] Stoup, P.: Athe development and failure of social norms in second life. Duke Law Journal 58(2), 311–344 (2008)

[30] Vasconcelos, W.W., Kollingbaum, M.J., Norman, T.J.: Normative conflict resolution in multi-agent systems. Autonomous Agents and Multi-Agent Systems 19(2), 124–152 (2009)

[31] Vázquez-Salceda, J.: Thesis: The role of norms and electronic institutions in multi-agent systems applied to complex domains. the harmonia framework. AI Communications 16(3), 209–212 (2003)

[32] Verhagen, H.: Simulation of the Learning of Norms. Social Science Computer Review 19(3), 296–306 (2001)

[33] 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, San Francisco, CA, pp. 384–389. MIT Press, Cambridge (1995)

[34] Weisstein, E.W.: Moore neighbourhood (2010), http://mathworld.wolfram.com/MooreNeighborhood.html

[35] Wieringa, R.J., Meyer, J.-J.C.: Applications of deontic logic in computer science: a concise overview. In: Deontic Logic in Computer Science: Normative System Specification, pp. 17–40. John Wiley & Sons, Inc., New York (1994)

[36] Wikipedia. Chebyshev distance, http://en.wikipedia.org/wiki/Chebyshev_distance (last accessed on September 15, 2010)

[37] Yolum, P., Singh, M.P.: Emergent properties of referral systems. In: Proceedings of the Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003), pp. 592–599. ACM, New York (2003)