

Simulating Normative Behaviour in Multi-Agent Environments using Monitoring Artefacts

Stephan Chang and Felipe Meneguzzi

School of Computer Science, Pontifical Catholic University of Rio Grande do Sul –
Porto Alegre, Brazil

`stephan.chang@acad.pucrs.br`, `felipe.meneguzzi@pucrs.br`

Abstract. Norms are an efficient way of controlling the behaviour of agents while still allowing agent autonomy. While there are tools for programming Multi-Agent Systems, few provide an explicit mechanism for simulating norm-based behaviour using a variety of normative representations. In this paper, we develop an artefact-based mechanism for norm processing, monitoring and enforcement and show its implementation as a framework built with CArtAgO. Our framework is then empirically demonstrated using a variety of enforcement settings.

1 Introduction

Multi-Agent Systems are often used as a tool for simulating interactions between intelligent entities within societies, organisations or other communities. This Agent-based Simulation is useful for studying social behaviour in hypothetical situations or situations that may not be easily reproduced in the real world. The entities being simulated, human or otherwise, are represented by programmable intelligent agents, which must present reactive, pro-active and social behaviour [1].

When working with social simulations, we must consider that agents should be free to act in their own best interest, even though their actions might produce negative effects to other agents. For this reason, rules are established to ensure that certain actions, which would otherwise harm the society’s performance, are prohibited. These rules, referred to as “norms” in multi-agent environments, allow agents to reason and act freely, while still being subject to punishment in the event that a norm is violated [2]. Agents are able to reason whether following a norm brings more positive results by avoiding the penalties associated to its violation. Some mechanisms [3][4][5][6][7][8][9][10] exist that makes reasoning about norms possible. In order to simulate norm-based behaviour, a structure must be defined that allows the specifying of norms in the form of prohibitions and obligations. Once these norms are active, agent interactions shall be observed by a monitoring mechanism and analysed by a norm-enforcing agent, which will then punish agents caught violating norms.

Although there are multiple frameworks for simulating agent-based behaviour, such as MASSim [11] or the agent programming languages Jason [12] and JADE

[13], relatively less attention has been focused on frameworks for norm-based behaviour simulation [14, Chapter 1]. In this paper, we aim to bridge this gap by developing a scalable norm processing mechanism that performs monitoring and enforcement in multi-agent environments. Our contributions are a mechanism to monitor agents's actions in an environment, described in Section 4.4 and a mechanism for norm maintenance and enforcement, described in Section 4.5. In Section 5 we show empirical results of applying our mechanism to a Multi-Agent System.

2 Multi-Agent Simulation

When intelligent agents [1] share an environment, competition between them becomes inevitable [15]. This idea becomes clear when we think of multi-agent systems as societies. Each person in a society has their own goals and plans to achieve them, and it is in their best interest to do so by spending as little effort as possible. Take for an example a person interested in eating an apple and another interested in selling one. For the buying person, its goal is to acquire the apple from the seller for the lowest cost possible, preferably with no cost at all. For the seller, the goal is to sell the apple for as high a price as affordable by the buyer, maybe even higher than that. Now, considering that in this hypothetical world no notion of ethics is known yet, the buyer soon realizes that instead of paying for the apple he wants to eat, he could simply grab it and eat it on the spot.

Competition between agents is often intended when working with agent-based simulations, as we desire to see how agents perform under such circumstances. However, to prevent the system as a whole from descending into chaos, rules must be established in order to control agent interactions while still allowing them to be autonomous. Nevertheless these rules must be limited to directing agents, rather than restraining them, otherwise, much of the benefit from autonomous agents is lost. When rules are set, agents that disregard them are subject to punishment for potentially harming the environment. In our buyer/seller system, we could establish a rule that guarantees items sold at shops must be paid for. If one is caught stealing, it will need to pay for the seller's injury. By doing so, we allow the buyer to reason about the advantages and disadvantages of obeying rules, letting it decide on an appropriate action plan. In multi-agent systems, we refer to these rules as norms.

Usual mechanisms for controlling agent interactions include interaction models, used by simulators such as NetLogo [16], MASON [17] and Repast [18]; strategies, commonly used in Game Theory; and regimented normative systems, such as *Moise* [19]. The disadvantage of these methodologies is that agents are constrained to the rules of their environment. They are not allowed to break rules because the system is rule compliant by design, also known as the regimentation approach [20]. However, unlike environmental constraints, perfectly enforcement (regimentation) for social norms is both undesired, because it prevents agents from occasional violations for the greater good, and unrealistic, as it is not achievable in the real world.

3 Normative Scenario - Immigration Agents

To facilitate explanation and exemplification of our approach, as well as to highlight its capabilities, we present the scenario that was used to test our solution. This scenario helps understand what norms are and how they control interactions in an environment. First, we present a short story that connects the environment to its agents, then we outline the norms that constrain them.

In a fictional emerging nation¹, an immigration program was started by the government to accelerate development through the hiring of foreigners. Besides landed immigrants, visitors are also welcome to the country, since money from tourism greatly boosts local economy. At the border, immigration officers are tasked with the inspection of immigrant's passports. The foreigner acceptance policy is quite straightforward, and immigrants with valid passports and no criminal records are to be accepted immediately, while John Does and refugees are to be outright rejected. It is believed that the more immigrants accepted, the better. Each officer's responsibility is to accept as many immigrants as possible, while still following the guidelines that were passed to them. Each accepted able worker nets the officer 5 credits, which eventually turn into a bonus to the officer's salary. There are no rewards for rejecting immigrants. It becomes clear that the bonus each officer accumulates depends entirely on chance, and some officers may accumulate more than others, if at all. As such, some officers might feel inclined to accept immigrants they should not, only to add to their personal gain.

To ensure officers act on the best interests of the nation only, an enforcement system is introduced to the offices at the borders. Among the officers working in the immigration office, one is responsible for observing and recording the behaviour of those working in booths. This officer is known as the "monitor". His job is to write reports about what the officers do and send these reports to another officer, known as the "enforcer". The enforcer then reads the reports that are passed to him and look for any inconsistencies, such as the approval of an illegal immigrant. As this represents a violation of a rule, or norm, the enforcer then carries out an action to sanction the offending officer. The penalties for approving an illegal immigrant are the immediate loss of 10 credits and suspension of work activities for up to 10 seconds. Considering that immigrants arrive at a rate of 1 per 2 seconds, in a 10-second timespan 5 immigrants would have arrived at a given booth, meaning that a violating officer potentially loses 25 credits. Added to the other portion of the sanction, the potential loss rises up to 35 credits.

The enforcement system, however, is not cost free. Each monitor and enforcer has an associated cost and it is within the interests of the nation to spend as little as possible with such a system. Therefore, the government wants to know how intensive the system must be to cover enough cases of disobedience so that officers will know violating norms is a disadvantage rather than an advantage.

¹ Inspired by the game "Papers, please": <http://papersplea.se>

There are two norms that can be extracted from this scenario. These defined in Examples 1 and 2, which are detailed in Section 4.2, where we describe the mathematical representation of norms in our system. These norms concern the stability of the immigration program by assuring valid immigrants are accepted and discouraging corrupt officers to accept those who should not be.

Example 1. “All immigrants holding valid passports must be accepted. Failure to comply may result in the loss of 5 credits.”

Example 2. “All immigrants holding passports that are not valid must not be accepted. Failure to comply may result in the loss of 10 credits and suspension from work activities for up to 10 seconds.”

4 NormMAS Framework

In this section, we develop our normative monitoring and enforcement framework for agent simulation. We start by reviewing the agent approaches that underpin our framework in Section 4.1. We follow with the formalisation of the norms processed in our system in Section 4.2, as well as the way actions are represented in the environment in Section 4.3. With the formalisation in place, we proceed to explaining the monitoring and enforcement systems in Sections 4.4 and 4.5, respectively.

4.1 Jason and CArTAgO

In order to show the feasibility of the mechanism proposed in this paper, we use two programming approaches: agent-oriented programming and environment-oriented programming. The former is provided by the Jason interpreter [12], while the latter is achieved with the Common Artifact infrastructure for Agents Open environments (CArTAgO) [21].

Jason provides us with a means to program agents using the *AgentSpeak* language [22] in a Java environment. Agents are built with the BDI [23] architecture, and so their behaviour is directed by beliefs, goals and plans. Beliefs are logical predicates that represent an agent’s considerations towards its environment. Predicates such as `valid(Passport)` and `wallet(50,dollars)` indicate that the agent believes the given passport variable is valid and that his wallet currently contains 50 dollars. In *AgentSpeak* variables start with an upper-case letter, while constants start with lower-case.

Goals are states which the agent desires to fulfil, and these can be either *achievement goals* or *test goals*. Achievement goals are objectives or milestones that agents pursue when carrying out their duties. To represent these in *AgentSpeak*, the goal’s name is preceded by the ‘!’ character. Test goals, on the other hand, are questions an agent may ask about the current state of the environment. These can be identified by a ‘?’ preceding the goal’s name.

To achieve these goals, agents need to perform sequences of actions that modify the environment towards the desired states. This sequence of actions

is referred to as a *plan*. A plan is not necessarily composed solely of actions, however, it can also contain sub-plans. This allows complex behaviours to be built, creating flows of actions that vary and are influenced by agent beliefs and perceptions.

As with any other programmed system, multi-agent systems must be tested before being effectively deployed to their end environments. To do so, test environments can be programmed for agents to be observed and any faulty behaviour addressed before release. Jason allows the programming of test environments in Java language, by providing an interface between agents and the programmed environment. These environments, however, are centralised, and so they are meant for small systems or specific test scenarios. This hinders scalability, which is an important aspect to consider when working with complex, more realistic scenarios or simply more robust structures. To address this limitation, we use the CArtAgO framework for environment programming.

In CArtAgO, environments are seen as composed by different artefacts. These artefacts represent objects in the environment through which agents interact with one another indirectly. *E.g.* a table in an office, which an agent may put reports on and from where another agent may pick these reports up to read them. The environment then becomes an abstraction, composed of different artefacts, which may be introduced to or removed from the environment whenever convenient. In our work, this allows us to create artefacts specifically for monitoring and enforcement tasks. These normative artefacts are then shared between normative agents so that more monitors and enforcers may be added to the system as it scales up.

4.2 Norms

In order to keep competition between agents manageable, norms are established to direct agent behaviour so that an environment's stability is maintained. This is achieved by specifying obligations and prohibitions [6]. Here, obligations are behaviours that agents must follow in a given context to comply with the norm, and prohibitions behaviours that jeopardise the environment's stability, and so must be avoided. Violating prohibitions is just as harmful as violating obligations, hence both cases must be addressed when detected. We expect that, when agents are punished for transgression, they are able to learn not to misbehave. Examples 1 and 2, in Section 3, correspond to an obligation and a prohibition, respectively.

While norms in the real world are expressed in natural language, they must be translated to a multi-agent environment so that agents are able to reason about them. This requires the extraction of necessary information related to a norm and composition of a mathematical representation. Agents should not have to reason how or why a certain norm came to be, but rather what the norm is about and what are the consequences of violating it. The format can also be extended to include other important information, such as the sanction function associated with a norm's violation, or the conditions for automatic activation

and expiration of the norm [6]. In this paper, norms are specified according to the tuple of Definition 1.

Definition 1. *A norm is represented by the tuple $\mathcal{N} = \langle \mu, \kappa, \chi, \tau, \rho \rangle$, where:*

- $\mu \in \{\textit{obligation}, \textit{prohibition}\}$ represents the norm’s modality.
- $\kappa \in \{\textit{action}, \textit{state}\}$ represents the type of trigger condition enclosed.
- χ represents the set of states (context) to which a norm applies.
- τ represents the norm’s trigger condition.
- ρ represents the sanction to be applied to violating agents.

Using Definition 1, we can proceed to formalising the norms from our example. We can formalize the first norm of our scenario from Example 1, as shown in Example 3.

Example 3. $\langle \textit{obligation}, \textit{action}, \textit{valid}(\textit{Passport}), \textit{accept}(\textit{Passport}), \textit{loss}(5) \rangle$

The process can be repeated for Example 2. By identifying the context of a norm, it is possible to define it solely with predicates and atoms, as shown in Example 4, below.

Example 4. $\langle \textit{prohibition}, \textit{action}, \textit{not valid}(\textit{Passport}), \textit{accept}(\textit{Passport}), \textit{loss}(10) \rangle$

4.3 Action Records

Similarly to norms, actions must also be stored as tuples containing essential information. Actions captured by monitors must only be accessed by agents of the enforcer type, and therefore only the pieces of information that can be associated with norms are deemed essential. These are: what was done; who did it; and under what context it was done. Example 5 shows how a monitor reports its observations to an enforcer:

Example 5. “Officer John Doe approved Passport #3225. The passport was known to be valid.”

From this report, we can extract the following details:

Example 6. $\langle \textit{johndoe}, \textit{approve}(\textit{Passport}), \textit{valid}(\textit{Passport}) \rangle$

In this example, an officer approves the entry of an immigrant holding a valid passport. The next report reads:

Example 7. “Officer John Smith approved Passport #2134. The passport was not known to be valid.”

From this report, we can extract the following details:

Example 8. $\langle \textit{johnsmith}, \textit{approve}(\textit{Passport}), \textit{notvalid}(\textit{Passport}) \rangle$

As such, we define Action Records:

Definition 2. *An Action Record, stored within the Action History, is represented by the tuple: $\mathcal{R} = \langle \gamma, \alpha, \beta \rangle$, where:*

- γ represents the agent executing the action;
- α represents the action being executed by the agent γ ; and
- β represents the agent γ ’s internal state at the moment of execution.

4.4 Monitoring System

The monitoring system has two roles: capturing agents’s actions and forwarding reports, employing a producer/consumer mechanism. An action is captured whenever any agent successfully executes an action. In CArtAgO, this means that each operation executed successfully is saved to the Action History. This function is agent independent and thus it is implemented directly in the simulation engine’s architecture. In our framework, we use an adapted agent architecture for Jason agents acting in CArtAgO environments and extend it so all successful actions are stored in a separate data structure referred to as the *Action History*. Should an action fail for any reason, it is ignored by the capturing system.

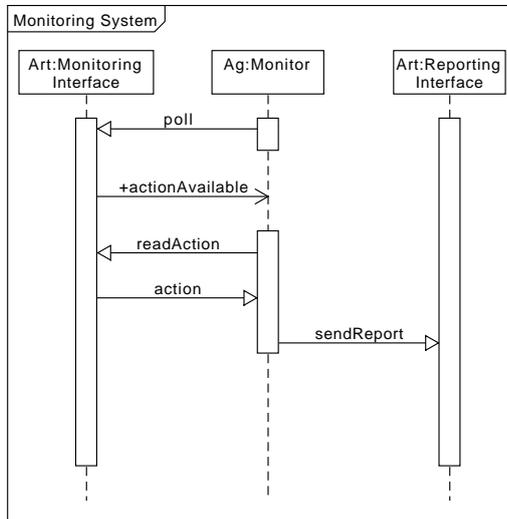


Fig. 1. Monitoring System Sequence Flow.

For these actions to be analysed, they must be sent to an enforcer agent in the form of a report. To achieve that, we use producer/consumer dynamic, in which an agent is tasked with continuously providing information through a channel, while another agent consumes this information. With this in mind, we can identify four components that are necessary for this set-up: a Producer, a Consumer, a channel for communications and the information itself. In our context, the role of Producer is given to the Monitor Agent; the role of Consumer is given to the Enforcer Agent; the communications channel is an interface called “Reporting Interface”; and the information that transits through this channel are reports containing the actions executed by agents. This process is illustrated in Figure 1.

Monitoring tasks are not cost-free, and monitoring costs grow with its intensity [24]. For this reason, it must be possible to adjust monitoring intensity so that enforcement can be performed effectively at a cost considered affordable by the society. Adjustments can be made either by configuring the monitor’s enforcement intensity (inducing some desirable probability of reading actions) or by creating monitoring strategies. In this paper, we use a purely probabilistic strategy to study the general behaviour of our simulation.

The **Action History** is a queue-like data structure that stores captured actions, from which monitors gather the information that is sent to the norm enforcers. Actions are stored in the format discussed in Section 4.3 and are removed from the queue as soon as a monitor attempts to read them, regardless of the monitor’s success when doing so. This represents the chance a violation will go unpunished.

4.5 Enforcement System

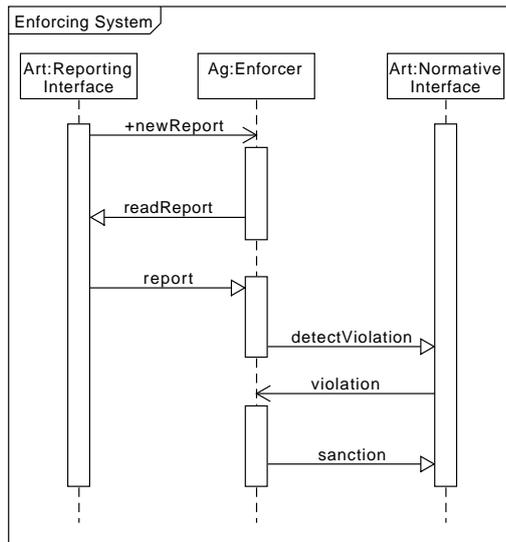


Fig. 2. Enforcement System Sequence Flow.

The enforcement system represents the Consumer entity in the normative mechanism’s Producer/Consumer scheme. An enforcer agent connects to the Reporting Interface and awaits the arrival of new reports to analyse. The arrival of new reports is perceived by the enforcer, and in our implementation this perception is mapped to the `+newReport` signal. Once the report submission is

perceived, the enforcer accesses the Normative Interface in search of currently activated norms and checks for any possible violations by the reported action.

During the violation detection routine, the perception of violations is also mapped to a signal, represented in the sequence diagram of Figure 2 as the `+violation` event. When a violation is perceived, it falls to the enforcer to apply associated sanctions. The sanctioning step is the last in this process, and it starts as soon as detection finishes.

In order to sanction violating agents, the normative mechanism must be able to recognise them. It does not make sense to be told “John has approved an invalid passport. He violated a norm.” if we do not know who John is in the first place. Therefore agents must be registered to the normative system prior to execution of their designed plans, similar to how people are registered for government issued IDs. In CArtAgO, this is accomplished through an operation in the Normative Interface that adds the agent’s ID to a list, so that they may be found when needed. The ID they are registered with should be the same that appears in Action Records.

Normative Base When norms are created, they must be stored within the system so that they may be accessed by an enforcer attempting to detect violations. The Normative Base structure holds all the norms that exist in the system, active or not. Every time a norm is created, it is stored in a list structure with a unique identification. Norms may be activated or deactivated through the Normative Interface. Every time a norm is created, activated, deactivated or destroyed, agents connected to the Normative Interface perceives the event.

Detecting Violations The detection operation runs for each action report received by an enforcer agent. Each action read is verified against the normative base, along with the context under which the action was executed. Since it is possible for an action to violate more than one norm, we utilize a list structure to take note of all violations detected so they will be properly addressed at a later time. At first, no norm is seen as violated and thus the list is empty. A norm is only added to the list when all verification steps finish with the variable’s *isViolated* value set to *True*. The procedure for detecting violations can be seen in Algorithm 1 and is explained further.

Detection of violations can be achieved in two steps: context analysis and trigger condition analysis. Context analysis is about making sure that the action’s execution context is the same as the one predicted by a norm. If it is, then there is a possibility of violation and further analysis is required. Otherwise, violation is considered an impossibility and the routine carries on. Formally, we define the norm’s context as χ and the acting agent’s belief-base as β . Hence, the context analysis returns *True* value if $\chi \subseteq \beta$. Algorithm 2 is used for comparing sets of predicates. It checks if all the predicates defined in context χ are present in the agent’s belief-base β , one by one. If a predicate in χ is negated (*e.g not valid(Passport)*), then the algorithm checks for its absence in belief-base β in-

stead. This is to reflect how **not** operator works in Jason. If the trigger condition is satisfied, the routine returns *True* value, and *False* otherwise.

Algorithm 1 Violation detection algorithm.

```

1: function DETECTVIOLATION( $\langle \gamma, \alpha, \beta \rangle$ )
2:    $V \leftarrow []$ 
3:   for each  $n = \langle \mu, \kappa, \chi, \tau, \rho \rangle \in ActiveNorms$  do
4:     if CONTEXTAPPLIES( $\chi, \beta$ ) then
5:       if CONDITIONAPPLIES( $\kappa, \tau, \alpha, \beta$ ) then
6:         if  $\mu = prohibition$  then
7:            $V \leftarrow V \cup \{n\}$   $\triangleright$  Violation detected! Adds to the list of violated norms.
8:         else
9:           if  $\mu = obligation$  then
10:             $V \leftarrow V \cup \{n\}$   $\triangleright$  Violation detected! Adds to the list of violated norms.
11:    for each  $n \in V$  do
12:      SIGNALVIOLATION( $n, \gamma$ )

```

A trigger condition of a norm can be either an action that was executed or a state the agent has reached. This is specified by the norm’s trigger condition type and directs the way in which the detection algorithm executes. If we are working with an action trigger, then we must compare the action that was executed with the one specified by the norm. However, if we are working with a state trigger, then two contexts must be compared: the agent’s belief-base and the norm’s state trigger condition. These are compared using the context analysis algorithm of Algorithm 2. The pseudo-code for the trigger analysis procedure can be seen in Algorithm 3.

Algorithm 2 Context comparison sub-routine.

```

1: function CONTEXTAPPLIES( $\chi = [l_1, \dots, l_n], \beta = [l_1, \dots, l_n]$ )
2:   Require  $count(\chi) \leq count(\beta)$ 
3:   for each  $p \in \chi$  do
4:      $isPresent \leftarrow False$ 
5:      $checkAbsence \leftarrow False$ 
6:     if  $p$  is of the form  $\neg\phi$  then
7:        $p \leftarrow \phi$ 
8:        $checkAbsence \leftarrow True$ 
9:     for each  $l \in \beta$  do
10:      if  $l = p$  then
11:         $isPresent \leftarrow True$ 
12:      break
13:     if  $checkAbsence = isPresent$  then
14:       return False
15:   return True

```

When both context and trigger conditions are satisfied, we need only verify whether the norm is an obligation or prohibition to conclude if it was violated or not. A prohibition means that a certain action or state is undesired under the given context. If all the conditions up to now have been met, we conclude that said undesired state has been reached and the norm was violated. On the other hand, an obligation requires the flow specified by the norm to be followed strictly, and if this is the case, we conclude that the norm was complied with. By negating our conditions, we also negate its results: if in a prohibition context the conditions were not met, then we would be home free; if they are not met while in an obligation context, however, we would have just violated it.

Algorithm 3 Trigger condition analysis sub-routine.

```

1: function CONDITIONAPPLIES( $\kappa, \tau, \alpha, \beta$ )
2:   if  $\kappa = \text{action}$  then
3:     return  $\tau = \alpha$ 
4:   return CONTEXTAPPLIES( $\tau, \beta$ )
    
```

Their modality notwithstanding, every norm that is violated is added to a list that is processed when all norms have been verified. Sanction functions are then executed and agents perceive their punishments. Penalties can be brought directly upon agents through perception or carried out by a third party, while records on agent transgressions can be maintained in a separate structure for greater consistency.

5 Evaluation

In order to test our solution, we programmed agents using Jason and deployed them in a CArtAgO environment following the scenario described in Section 3. To visualise the difference between compliant and non-compliant behaviours, two types of agents were used: the *normal* type and the *corrupt* type. The normal type is programmed to approve only those passports that are truly valid, whereas the corrupt one will approve passports indiscriminately for his own personal gain. By making it so, we can more easily tell the effectiveness of the norm enforcing mechanism. Therefore, the following results were expected:

- Corrupt agents attain more credits when under lower monitoring intensity.
- Standard agents maintain an average quantity of credits through all simulations.
- At some point, corrupt agents should start performing poorly due to higher monitoring intensity. This marks the point at which monitoring can change the environment.

We ran 35 simulations for 11 different values of monitoring intensity. Intensity values range from 0 to 100, with a step value of 10. Each simulation was run for 10 minutes. In this timespan, with our set-up, around 1048 immigrants attempt to cross the border. In what follows, we refer to an agent's obtained credits, or their performance measure, as their utility. We use that measure in the graph of Figure 3, which illustrates how the environment's monitoring intensity affects the utilities of corrupt agents 1 and 2. The monitoring intensity is the probability as a percentage of a monitor to read an agent's action. A value of 100 means that all actions are read, while a value of 0 means no actions are read by the monitor. We notice that, as the intensity of the monitoring mechanism increases, the utility of corrupt agents decreases to the point where performing badly and not performing at all yield the same utility, whereas normal agents maintain their average utility. This allows us to conclude that, for a monitoring intensity value of 40 or more, following norms is a better decision than the contrary.

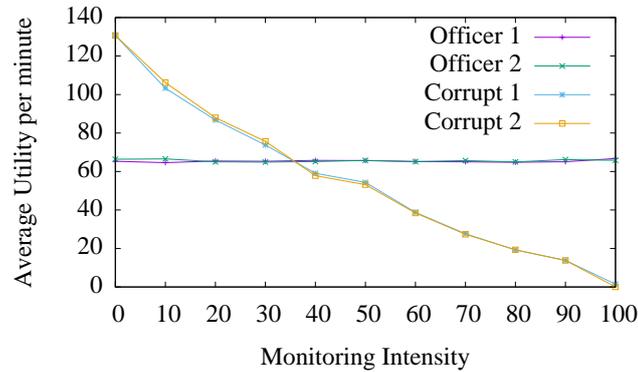


Fig. 3. Utility of corrupt agents is affected by monitoring intensity.

The data used to plot the graph of Figure 3 can be seen in Table 1. Values for μ and σ represent the arithmetic mean and standard deviation, respectively. These were calculated to show that utility values for normal agents are near constant. The μ values for corrupt agents show that, at the end of the simulation, their average performance is worse than those of normal agents, due to their constant violation of norms. A high σ value for these agents shows that their performance suffers between simulations. We can then see that through the analysis of recorded agent actions and successful identification of violation occurrences, violating agents are punished by the enforcement system and have their utilities affected.

Table 1. Agent Utilities \times Monitoring Intensity.

Intensity	<i>officer1</i>	<i>officer2</i>	<i>corrupt officer1</i>	<i>corrupt officer2</i>
0	65,3285	66,3714	130,6571	130,7000
10	64,5871	66,5714	103,3000	106,2285
20	65,4428	65,0142	86,8000	87,9571
30	65,3142	64,8714	73,7571	75,6571
40	65,7857	65,1857	59,0571	57,8142
50	65,6714	65,7714	54,3285	53,1857
60	65,1571	65,1714	38,7714	38,4571
70	65,0142	65,6571	27,6428	27,3714
80	64,7857	64,9571	19,2285	19,3428
90	65,0714	66,1714	13,7857	13,8142
100	66,7571	65,8000	1,4714	0,0285
μ	65.3559	65.5948	55,3454	55,5051
σ	0.5569	0.5705	38,4836	39,1996

6 Related Work

There are multiple tools available for programming multi-agent environments, few of which provide mechanisms for norm specification. These tools range from programming libraries to model-based simulators. To name a few, NetLogo [16] and its distributed version HubNet [25] are of the model-based type and allow users to work with educational projects and, to some extent, professional ones. Other tools include MASON [17] and Repast [18]. MASON is a simulation library developed in Java that provides functions for modelling agents and visualising simulations as they run. As for Repast, it uses interaction models much like NetLogo does, although it is meant for professional use and thus offers more alternatives for agent programming. One final example worth mentioning is MASSim [11], which promotes multi-agent research and is used in the MAS Programming Contest² [26]. This one, however, provides only the tools related to the contests. Although it is possible to develop custom agents for operation within the simulator, the practice is not encouraged by its developers.

Building a full-fledged norm-based behaviour simulation engine is not a trivial task, and the “Emergence in the Loop” (EMIL) [27] project built a set of tools to accomplish this objective. A toolset which includes an extension of the BDI architecture that is capable of simulating the processes referred to as “immersion” and “emergence” of norms [28]; and an integration with multi-agent modelling tools such as NetLogo [16] and Repast [18]. In this way, agents are modelled in one of these environments and then simulated using the EMIL agent architecture. It is a very powerful tool for studying social behaviour in autonomous agents, since agents can reason about norms and, together, create conventions of what kinds of behaviours must be avoided or followed. EMIL’s approach to normative simulation is more focused on agents and their experience

² <https://multiagentcontest.org>

with norms. This contrasts with our approach in that we are more focused on norm monitoring and enforcement tasks, and little is said about these matters in the EMIL literature. We also consider the environmental aspects of Normative Multi-Agent Systems, which is why we employ CArtAgO in our implementation.

7 Conclusions and Future Work

In this paper, we constructed a mechanism of norm processing and enforcement in a multi-agent environment. We show its feasibility with an implementation using Jason [12] and Cartago [21] technologies. By keeping track of agent activities and analysing actions against a normative base, it is possible to detect violations and enforce norms through the sanctioning of violating agents. With this framework, it is possible to evaluate different implementations [6,29,30,31] of normative behaviour. Statistics collection can also be customised so that results may be compared between simulations.

CArtAgO allows us to build environments in a distributed manner, therefore providing scalability for realistic simulation scenarios or complex multi-agent systems. The philosophy behind CArtAgO, which sees the environment as the composition of artefacts through which agents interact, also aided in the framework's construction. Artefacts are modular, they can be attached or detached to a multi-agent system seamlessly. Meaning that artefacts can be created to suit an agent's or group of agents's specific needs, and agents may connect only to those artefacts that are related to their designs. We took advantage of those features to build the interfaces for the monitoring system to access the Action History and Normative Base structures.

As future work, we aim to build improvements and extensions to the framework, such as: a mechanism to be added to the normative system that allows activation and expiration of norms following predefined conditions; agent architectures that can learn from normative environments, and with that avoid penalties by violation or minimising performance loss when violations are inevitable [6]; enable agents to learn about the enforcing intensity and use that information to their advantage [24]; and the introduction of agent hierarchies to control normative power [32].

References

1. M. Wooldridge, "Intelligent Agents," in *Multi-Agent Systems* (G. Weiss, ed.), pp. 3–50, The MIT Press, 2nd ed., 2013.
2. A. J. I. Jones and M. Sergot, *Deontic Logic in Computer Science: Normative System Specification*, ch. Chapter 12: On the characterisation of law and computer systems: the normative systems perspective, pp. 275–307. Wiley Professional Computing Series, Wiley, 1993.
3. M. Kollingbaum, *Norm-governed Practical Reasoning Agents*. PhD thesis, University of Aberdeen, 2005.

4. J. Broersen, M. Dastani, J. Hulstijn, Z. Huang, and L. van der Torre, "The BOID architecture: conflicts between beliefs, obligations, intentions and desires," in *Proceedings of the Fifth International Conference on Autonomous Agents*, pp. 9–16, 2001.
5. G. Governatori and A. Rotolo, "BIO logical agents: Norms, beliefs, intentions in defeasible logic," *Autonomous Agents and Multi-Agent Systems*, vol. 17, no. 1, pp. 36–69, 2008.
6. F. Meneguzzi and M. Luck, "Norm-based behaviour modification in BDI agents," in *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems*, pp. 177–184, 2009.
7. N. Criado, *Using Norms To Control Open Multi-Agent Systems*. PhD thesis, Universitat Politècnica de València, 2012.
8. N. Alechina, M. Dastani, and B. Logan, "Programming norm-aware agents," in *Autonomous Agents and Multi-Agent Systems* (W. van der Hoek, L. Padgham, V. Conitzer, and M. Winikoff, eds.), pp. 1057–1064, IFAAMAS, 2012.
9. S. Panagiotidi, J. Vázquez-Salceda, and F. Dignum, "Reasoning over norm compliance via planning," in *Coordination, Organizations, Institutions, and Norms in Agent Systems VIII, 14th International Workshop, Held Co-located with the International Conference on Autonomous Agents and Multi-Agent Systems, Valencia, Spain, June 5, 2012, Revised Selected Papers* (H. Aldewereld and J. S. Sichman, eds.), vol. 7756 of *Lecture Notes in Computer Science*, pp. 35–52, Springer, 2012.
10. F. Meneguzzi, S. Mehrotra, J. Tittle, J. Oh, N. Chakraborty, K. Sycara, and M. Lewis, "A cognitive architecture for emergency response," in *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems*, pp. 1161–1162, 2012.
11. T. M. Behrens, M. Dastani, J. Dix, and P. Novák, "MASSim: Multi-Agent Systems Simulation Platform," *In Begehung des Simulationswissenschaftlichen Zentrums*, 2008.
12. R. H. Bordini, J. F. Hübner, and M. Wooldridge, *Programming Multi-Agent Systems in AgentSpeak Using Jason (Wiley Series in Agent Technology)*. John Wiley & Sons, 2007.
13. F. L. Bellifemine, G. Caire, and D. Greenwood, *Developing Multi-Agent Systems with JADE (Wiley Series in Agent Technology)*. John Wiley & Sons, 2007.
14. R. Conte, G. Andrighetto, and M. Campennl, *Minding Norms: Mechanisms and Dynamics of Social Order in Agent Societies*. Oxford Series on Cognitive Models and Architectures, OUP USA, 2013.
15. M. Fagundes, S. Ossowski, and F. Meneguzzi, "Analyzing the tradeoff between efficiency and cost of norm enforcement in stochastic environments populated with self-interested agents," in *Proceedings of the 21st European Conference on Artificial Intelligence*, 2014.
16. U. Wilensky, "NetLogo." Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL, 1999. <http://ccl.northwestern.edu/netlogo/>.
17. S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan, "MASON: A Multiagent Simulation Environment," *Simulation*, vol. 81, pp. 517–527, July 2005.
18. M. North, N. Collier, J. Ozik, E. Tatara, C. Macal, M. Bragen, and P. Sydelko, "Complex adaptive systems modeling with repast symphony," *Complex Adaptive Systems Modeling*, vol. 1, no. 1, 2013.
19. J. F. Hubner, J. S. Sichman, and O. Boissier, "Developing Organised Multiagent Systems Using the MOISE+ Model: Programming Issues at the System and Agent Levels," *Int. J. Agent-Oriented Softw. Eng.*, vol. 1, pp. 370–395, Dec. 2007.

20. A. J. I. Jones and M. Sergot, "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.
21. A. Ricci, M. Viroli, and A. Omicini, "CARTAgO: A Framework for Prototyping Artifact-based Environments in MAS," in *Proceedings of the 3rd International Conference on Environments for Multi-agent Systems III*, (Berlin, Heidelberg), pp. 67–86, Springer-Verlag, 2006.
22. A. S. Rao, "AgentSpeak(L): BDI Agents Speak Out in a Logical Computable Language," in *Agents Breaking Away, 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Eindhoven, The Netherlands, January 22-25, 1996, Proceedings* (W. V. de Velde and J. W. Perram, eds.), vol. 1038 of *Lecture Notes in Computer Science*, pp. 42–55, Springer, 1996.
23. M. E. Bratman, *Intention, Plans and Practical Reason*. Harvard University Press, 1987.
24. F. Meneguzzi, B. Logan, and M. S. Fagundes, "Norm monitoring with asymmetric information," in Bazzan *et al.* [33], pp. 1523–1524.
25. U. Wilensky and W. Stroup, "HubNet." Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL., 1999. <http://ccl.northwestern.edu/netlogo/hubnet.html>.
26. T. M. Behrens, M. Dastani, J. Dix, J. Hübner, M. Köster, P. Novák, and F. Schlesinger, "The Multi-Agent Programming Contest," *AI Magazine*, vol. 33, no. 4, pp. 111–113, 2012.
27. G. Andrighetto, R. Conte, P. Turrini, and M. Paolucci, "Emergence In the Loop: Simulating the two way dynamics of norm innovation," in *Normative Multi-agent Systems, 18.03. - 23.03.2007* (G. Boella, L. W. N. van der Torre, and H. Verhagen, eds.), vol. 07122 of *Dagstuhl Seminar Proceedings*, Internationales Begegnungs- und Forschungszentrum für Informatik, Schloss Dagstuhl, Germany, 2007.
28. G. Andrighetto, M. Campenni, R. Conte, and M. Paolucci, "On the Immersion of Norms: a Normative Agent Architecture," in *Proceedings of the Association for the Advancement of Artificial Intelligence Symposium, Social and Organizational Aspects of Intelligence*, Forthcoming, 2007.
29. J. Lee, J. Padget, B. Logan, D. Dybalova, and N. Alechina, "Run-time norm compliance in BDI agents," in Bazzan *et al.* [33], pp. 1581–1582.
30. W. W. Vasconcelos, M. J. Kollingbaum, and T. J. Norman, "Normative conflict resolution in multi-agent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 19, no. 2, pp. 124–152, 2009.
31. N. Criado, E. Argente, V. J. Botti, and P. Noriega, "Reasoning about norm compliance," in *10th International Conference on Autonomous Agents and Multiagent Systems, Taipei, Taiwan, May 2-6, 2011, Volume 1-3* (L. Sonenberg, P. Stone, K. Tumer, and P. Yolum, eds.), pp. 1191–1192, IFAAMAS, 2011.
32. N. Oren, M. Luck, and S. Miles, "A model of normative power," in *9th International Conference on Autonomous Agents and Multiagent Systems, Toronto, Canada, May 10-14, 2010, Volume 1-3* (W. van der Hoek, G. A. Kaminka, Y. Lespérance, M. Luck, and S. Sen, eds.), pp. 815–822, IFAAMAS, 2010.
33. A. L. C. Bazzan, M. N. Huhns, A. Lomuscio, and P. Scerri, eds., *International conference on Autonomous Agents and Multi-Agent Systems, Paris, France, May 5-9, 2014*, IFAAMAS/ACM, 2014.