

Norm Monitoring under Partial Action Observability

Natalia Criado and Jose M. Such, *Member, IEEE*

Abstract—In the context of using norms for controlling multi-agent systems, a vitally important question that has not yet been addressed in the literature is the development of mechanisms for monitoring norm compliance under partial action observability. This paper proposes the reconstruction of unobserved actions to tackle this problem. In particular, we formalise the problem of reconstructing unobserved actions, and propose an information model and algorithms for monitoring norms under partial action observability using two different processes for reconstructing unobserved actions. Our evaluation shows that reconstructing unobserved actions increases significantly the number of norm violations and fulfilments detected.

Index Terms—Norm Monitoring, Action Observability.

I. INTRODUCTION

Within the Multi-agent System (MAS) area, norms are understood as means to coordinate and regulate the activity of autonomous agents interacting in a given social context [22]. The existence of autonomous agents that are capable of violating norms entails the development of control mechanisms that implement norms in agent societies.

In the existing literature, several authors have proposed infrastructures to observe agent actions and detect norm violations upon them [2], [8], [11], [12], [16], [23], [26]. The majority of these proposals have focused on providing efficient and scalable methods to monitor norms in dynamic agent societies, but they assume that all actions of agents are observable. However, this assumption is too strong because it is not necessarily true that all actions to be controlled can always be observed. One reason for this is that observing actions usually entails high costs. For example, the costs of setting, maintaining, and managing traffic radars to detect car speeds are very high, so traffic authorities usually decide to install a few of them in specific and critical locations. Another reason is that illegal actions may take place outside the institution controlled by the monitor; however, the effects of these actions can still be detected within the institution. For example, black market transactions cannot be directly observed by legal authorities, yet the corresponding money laundering transactions can be detected and sanctioned by these authorities.

Very recent work on norm monitoring under partial action observability proposes solutions to ensure complete action observability by increasing the actions that are observed, either by adding more monitors [7] or by adapting the norms to what can be observed [3]. However, these solutions are not always appropriate or feasible. For instance, in e-markets, such

as eBay¹ or Amazon², it is not possible to change trading laws to what can be observed. This paper goes beyond these approaches by also considering actions that were not observed but that can be reconstructed from what was observed.

The main contributions of this paper are: (i) a formalisation of the problem of reconstructing unobserved actions from observed actions for the purpose of norm monitoring; (ii) an exhaustive and an approximation solution to this problem; and (iii) an information model and algorithms used to monitor norms under partial action observability. Through an extensive empirical evaluation, we show that reconstructing unobserved actions increases noticeably the number of norm violations and fulfilments detected.

This paper is organised as follows: Section II contains the preliminary definitions used in this paper. Section III describes the information model of norm monitor proposed in this paper. Section IV contains the algorithms executed by norm monitors. Our proposal is evaluated in Section V. Related work is discussed in Section VI. Finally, conclusions are contained in Section VII.

II. PRELIMINARY DEFINITIONS

\mathcal{L} is a first-order language containing a finite set of predicate and constant symbols, the logical connective \neg , the equality (inequality) symbol $=$ (\neq), the true \top and false propositions \perp , and an infinite set of variables. The predicate and constant symbols are written as any sequence of alphanumeric characters beginning with a lower case letter. Variables are written as any sequence of alphanumeric characters beginning with a capital letter. We will relate our formulae via logical entailment \vdash ($\not\vdash$). We also assume the standard notion of substitution of variables; i.e., a substitution σ is a finite and possibly empty set of pairs Y/y where Y is a variable and y is a term.

The set of grounded atomic formulas of \mathcal{L} is built of a finite set of predicates and objects that characterise the properties of the world relevant to norm monitoring. By a situation, we mean the properties that are true at a particular moment. Some of these properties are static and not altered by action execution, whereas other properties are dynamic and changed due to agent actions. Specifically, we represent static properties as a set³ of atomic grounded formulas of \mathcal{L} , denoted by g . A state s is a set of grounded atomic formulas of \mathcal{L} , describing dynamic properties which hold on state s . Thus, a situation is built on a “closed assumption” and defined by a set of static properties g and a state s . Moreover, there is a set of inference rules (∇) representing domain knowledge.

EXAMPLE 1. *In this paper we will use a running example in which there are three robots that should attend requests at*

N. Criado is with the School of Computing and Mathematical Sciences, Liverpool John Moores University, UK, e-mail: n.criado@ljmu.ac.uk.

J. M. Such is with the School of Computing and Communications Info-lab21, Lancaster University, UK, email:j.such@lancaster.ac.uk.

¹<http://www.ebay.com>

²<http://www.amazon.com>

³In this paper sets are to be interpreted as the conjunction of their elements.




Fig. 1: Example Scenario. Offices are represented by squares, agents are represented by circles and the corridors are represented by arrows. Black arrows correspond to corridors observed by the Norm Monitor (NM) and grey arrows correspond to corridors not observed by the NM.

six offices in a building. The goal of the robots is to attend these requests as soon as possible. Figure 1a depicts our initial scenario. In our example, the language \mathcal{L} contains: 4 predicate symbols (*robot*, *office*, *in*, *corridor*), used to represent the robots and offices, the positions of the robots and the connections between offices in the building; 3 constant symbols to represent the robots ($r1, r2, r3$); and 6 constant symbols to represent the offices (a, b, c, d, e, f). The information about the robots, offices and corridors between offices is static and represented as follows:

$$g = \{robot(r1), robot(r2), robot(r3), office(a), \dots, office(f), \\ corridor(a, b), corridor(b, a), \dots, corridor(e, a)\}$$

The information about the location of the robots is dynamic. Specifically, the initial state s_0 is defined as follows:

$$s_0 = \{in(r1, a), in(r2, d), in(r3, e)\}$$

In this domain there is an inference rule (∇) representing that a robot cannot be in two different offices at the same time:

$$\nabla = \{\{in(R1, OA), in(R1, OB), OA \neq OB\} \vdash \perp\}$$

A. Action Definitions

D is a finite set of action descriptions that induce state transitions. An action description d is represented using preconditions and postconditions. If a situation does not satisfy the preconditions, then the action cannot be applied in this situation. In contrast, if the preconditions are satisfied, then the action can be applied transforming the current state into a new state in which all negative literals appearing in the postconditions are deleted and all positive literals in the postconditions are added. Moreover, actions are executed in a MAS and, as a consequence, we need to be able to represent concurrent actions with interacting effects. For the sake of simplicity, we will represent concurrent actions without an explicit representation of time⁴ as proposed in [6]. The main idea behind this representation is that individual agent actions do interact (i.e., one action might only achieve the intended effect if another action is executed concurrently). Specifically, each

⁴An explicit representation of time may play a role on other problems like scheduling concurrent actions, but is not strictly necessary for monitoring the effects of interaction.

action is also represented by a (possibly empty) concurrent condition that describes the actions that must (or cannot) be executed concurrently⁵.

Definition 1. An action description d is a tuple $\langle name, pre, con, post \rangle$ where:

- *name* is the action name;
- *pre* is the precondition, i.e., a set of positive and negative literals of \mathcal{L} (containing both dynamic and static properties) as well as equality and inequality constraints on the variables;
- *con* is the concurrent condition; i.e., a set of positive and negative action schemata⁶, some of which can be partially instantiated or constrained;
- *post* is the postcondition; i.e., a set of positive and negative literals of \mathcal{L} (containing dynamic properties only).

Given an action description d , we denote by $pre(d)$, $con(d)$, $post(d)$ the action precondition, concurrent condition and postcondition.

EXAMPLE 2. In our example, there is only one action that can be executed by robots:

$$\left\langle \begin{array}{l} move, \{robot(R), office(O1), office(O2), in(R, O1), \\ corridor(O1, O2)\}, \{\}, \{-in(R, O1), in(R, O2)\} \end{array} \right\rangle$$

This action represents the movement of a robot from one office to another. The parameters of this action are the robot (R), the source office ($O1$), the destination office ($O2$). To execute this action, the robot should be located at the source office and the two offices should be connected. Once the operation has been applied, the robot is no longer at the source office and it is at the destination office.

Finally, note that it is possible to define a NOP action that allows agents to remain still. For the sake of clarity and simplicity, the running example does not include the NOP action. However, our proposal is agnostic with respect to the existence of a NOP action and we have carried out experiments that include the NOP action (see Section V-B).

Definition 2. Given a situation represented by the state s and a set of static properties g , and an action description $d = \langle name, pre, con, post \rangle$; an action instance (or action) is a tuple $\langle name, pre', con', post' \rangle$ such that:

- There is a substitution σ of variables in pre , such that the precondition is satisfied (i.e., entailed by) the situation; i.e., $s, g \vdash \sigma \cdot pre$;
- $\sigma \cdot pre, \sigma \cdot post$ are grounded;
- pre' is a set of grounded literals in $\sigma \cdot pre$ containing dynamic properties only;
- $post' = \sigma \cdot post$ and $con' = \sigma \cdot con$.

⁵A more sophisticated definition of the concurrent condition would allow actions to have conditional effects according to the actions that are executed concurrently. Without loss of expressiveness, we will not consider conditional effects in action descriptions (note that any action with conditional effects can be represented by a set of actions with non conditional effects).

⁶An action schema contains an action name and the parameters of this action. Note that positive action schemata are implicitly existentially quantified –i.e., one instance of each positive schema must occur concurrently– and negative schemata are implicitly universally quantified.

Given an action a , we denote by $actor(a)$ the agent performing the action, and by $pre(a), con(a), post(a)$ the precondition, concurrent condition and postcondition.

EXAMPLE 3. In state s_0 , the robot $r1$ moves from office a to office b . This is formalised as follows:

$$\left\langle \begin{array}{l} move, \{robot(r1), office(a), office(b), in(r1, a), \\ corridor(a, b)\}, \{\}, \{-in(r1, a), in(r1, b)\} \end{array} \right\rangle$$

In a MAS, concurrent actions⁷ define state transitions. More formally, a concurrent action $A = \{a_1, \dots, a_n\}$ is a set of individual actions. Given a set of actions $A = \{a_1, \dots, a_n\}$, we define $pre(A) = \bigcup pre(a_i)$, $post(A) = \bigcup post(a_i)$ and $actor(A) = \bigcup actor(a_i)$.

Given a concurrent action $A = \{a_1, \dots, a_n\}$ we say that the concurrent condition of an individual action a_i of A is *satisfied* when for all positive schema in the concurrent condition exists an action a_j ($i \neq j$) in A , such that a_j is an instance of the schema; and for all negative schema none of the elements in A is an instance of the schema. For the sake of simplicity, we assume that each agent performs one action at a time⁸.

Definition 3. (Consistency [6]) Given a concurrent action $A = \{a_1, \dots, a_n\}$ it is consistent if:

- $pre(A)$ is consistent (i.e., $pre(A) \not\vdash \perp$);
- $post(A)$ is consistent (i.e., $post(A) \not\vdash \perp$);
- the concurrent condition of each action is satisfied;
- the concurrent action is complete; i.e., each agent performs one action (possibly the **NOP** action) in A .

EXAMPLE 4. The concurrent action $A = \{move(r1, a, b), move(r2, d, a), move(r3, e, a)\}$ ⁹ is consistent since:

- $pre(A) = \{in(r1, a), in(r2, d), in(r3, e)\}$ which is consistent;
- $post(A) = \{in(r1, b), in(r2, a), in(r3, a), \neg in(r1, a), \neg in(r2, d), \neg in(r3, e)\}$ which is consistent;
- the concurrent conditions of both actions are satisfied;
- each robot performs one action.

A concurrent action $A = \{a_1, \dots, a_n\}$ is applicable in a situation if A is consistent¹⁰ and each individual action $a_i \in A$ is applicable in this situation.

Given a consistent action, we define its effects as the postconditions of its individual actions and the preconditions not invalidated by the postconditions. More formally, given a concurrent action $A = \{a_1, \dots, a_n\}$ its effects are a set of grounded literals as follows:

$$eff(A) = \left(\bigcup_{\substack{\forall pre \in pre(A): \\ pre, post(A) \not\vdash \perp}} pre \right) \bigcup \left(\bigcup_{\forall post \in post(A)} post \right)$$

⁷Concurrent action means actions that occur at the same time and does not necessarily imply agent cooperation or coordination.

⁸This limitation can be relaxed by decomposing agents into groups of agents corresponding to agents' actuators [6].

⁹For simplicity, we represent actions by their schemata.

¹⁰Note that we assume that actions executed by agents at a time are consistent, which entails that observed actions are consistent. Relaxing this assumption may not be desirable, as this would entail issues such as that mutually exclusive actions can be executed simultaneously. For example, the world cannot be in an inconsistent situation in which two robots have moved a piece simultaneously and the piece is in two different locations at the same time.

B. Norm Definitions

We consider *norms* as formal statements that define patterns of behaviour by means of *deontic modalities* (i.e., *obligations* and *prohibitions*). Specifically, our proposal is based on the notion of norm as a conditional rule of behaviour that defines under which circumstances a pattern of behaviour becomes relevant and must be fulfilled [23], [22].

Definition 4. A norm is defined as a tuple $\langle deontic, condition, action \rangle$, where:

- $deontic \in \{\mathcal{O}, \mathcal{P}\}$ is the deontic modality of the norm, determining if the norm is an obligation (\mathcal{O}) or prohibition (\mathcal{P});
- $condition$ is a set of literals of \mathcal{L} as well as equality and inequality constraints that represents the norm condition, i.e., it denotes the situations in which the norm is relevant;
- $action$ is a positive action schema that represents the action controlled by the norm.

EXAMPLE 5. In our example, there is a norm that avoids collisions by forbidding any robot to move into an office when the office is occupied by another robot:

$$\langle \mathcal{P}, in(R1, L2), move(R2, L1, L2) \rangle$$

This norm states that when a robot $R1$ is located in office $O2$ other robots are forbidden to move from any office $L1$ to $L2$.

In line with related literature [2], [3], we consider a *closed legal system*, where everything is considered permitted by default, and obligation and prohibition norms define exceptions to this default permission rule. We also define that a norm is relevant to a specific situation if the norm condition is satisfied in the situation. Besides, we define that a norm condition is satisfied in a given situation when there is a substitution of the variables in the norm condition such that the constraints in the norm condition are satisfied and the positive (vs. negative) literals in the norm condition are true (vs. false) in the situation.

Definition 5. Given a specific situation denoted by a state s and a set of static properties g , and a norm $\langle deontic, condition, action \rangle$; a norm instance is a tuple $\langle deontic, action' \rangle$ such as:

- There is a substitution σ such that the condition is satisfied in the situation; i.e., $s, g \vdash \sigma \cdot condition$;
- $action' = \sigma \cdot action$.

EXAMPLE 6. In state s_0 the norm that forbids robots to move into occupied offices is instantiated as follows:

$$\langle \mathcal{P}, move(R2, L1, d) \rangle \text{ where } \sigma = \{L2/d\}$$

$$\langle \mathcal{P}, move(R2, L1, a) \rangle \text{ where } \sigma = \{L2/a\}$$

$$\langle \mathcal{P}, move(R2, L1, e) \rangle \text{ where } \sigma = \{L2/e\}$$

The semantics of instances (and norms in general) depends on their deontic modality. An obligation instance is fulfilled when the mandatory action is performed and violated otherwise, while a prohibition instance is violated when the forbidden action is performed and fulfilled otherwise. We

classify detected violations (vs. fulfilments) into: *identified* violation (vs. fulfilment), which refers to when the monitor knows the specific action that an agent has executed and violates (vs. fulfils) an instance; and *discovered* violation (vs. fulfilment), which refers to when the monitor knows that an agent has violated (vs. fulfilled) some instance but does not know the forbidden (vs. mandatory) action executed by the agent.

III. NM INFORMATION MODEL

Let us assume a set of agents Ag to be monitored, a set of norms N that regulate the actions of agents, and a set D of action descriptions that represent the actions that can be performed by agents. For the sake of simplicity, we assume that there is a single Norm Monitor (NM) that observes the actions performed by agents and monitors norm compliance¹¹. We also assume that actions are deterministic and that the current state evolves due to action execution only¹². The goal of the NM is to analyse a partial sequence of action observations to detect norm violations. The enforcement of norms is out of the scope of this work and we assume that once the NM detects a norm violation (vs. fulfilment), it applies the corresponding sanction (vs. reward).

A. State Representation

As the NM only observes a subset of the actions performed by agents, it has partial information about the state of the world. The NM represents each partial state of the world, denoted by p , using an “open world assumption” as a set of grounded literals that are known in the state. Thus, a partial state contains positive (vs. negative) grounded literals representing dynamic properties known to be true (vs. false) in the state. The rest of dynamic properties are unknown.

To begin with, assume that the NM monitor has complete knowledge of the initial state (this will be relaxed later). Thus, at t_0 the NM knows which grounded atomic formulas are true or false in the initial state ($p_0 \equiv s_0$). From that moment on, the NM monitors the actions performed by agents at each point in time. At time t_0 the NM carries out a monitoring activity and observes some of the actions performed by agents (Act_0). These actions have evolved s_0 into a new state s_1 . As previously mentioned, the NM has limited capabilities for observing the actions performed by agents. Thus, it is possible that the NM observes a subset of the actions performed by agents. Specifically, if all actions have been observed ($|Act_0| = |Ag|$), then the resulting partial state p_1 can be constructed by considering the effects of actions in Act_0 on p_0 so $p_1 \equiv s_1$. A different case arises when the NM observes a subset of the actions performed by the agents ($|Act_0| < |Ag|$). In this case, the agent cannot be sure about the effects of unobserved actions. Thus, the new partial state p_1 is constructed by assuming that the postconditions of the observed actions must hold on state s_1 (i.e., positive postconditions are positive literals in p_1 and negative postconditions are negative literals

in p_1) and the rest of dynamic propositions are unknown. If the NM takes into account the next sequence of actions that it observes at time t_1 (Act_1), then the NM can also infer that the preconditions of these actions must hold on state s_1 , and, as a consequence, new propositions can be taken for sure in the partial state p_1 , retrospectively. Partial states in the general case are defined as:

Definition 6. *Given a partial state description p_t corresponding to time t , and two consecutive sequences of observed actions Act_t and Act_{t+1} executed by agents at times t and $t + 1$, respectively; the new partial state p_{t+1} resulting from executing actions Act_t in p_t and actions Act_{t+1} in p_{t+1} is obtained as follows:*

$$p_{t+1} = \begin{cases} post(Act_t) \cup pre(Act_{t+1}) & \text{if } |Act_t| < |Ag| \\ p_t^* \cup eff(Act_t) \cup pre(Act_{t+1}) & \text{otherwise} \end{cases}$$

where p_t^* is the set of invariant literals; i.e., literals of p_t that have not been modified by the actions in Act_t and it is defined as follows:

$$p_t^* = \bigcup_{\substack{\forall l \in p_t: \\ l, eff(Act_t) \not\vdash \perp}} l$$

EXAMPLE 7. *In our example, the NM knows which grounded atomic formulas are true or false in the initial state:*

$$p_0 = \{in(r1, a), \neg in(r1, b), \neg in(r1, c), \neg in(r1, d), \neg in(r1, e), \\ \neg in(r1, f), in(r2, d), \neg in(r2, a), \neg in(r2, b), \neg in(r2, c), \\ \neg in(r2, e), \neg in(r2, f), in(r3, e), \neg in(r3, a), \neg in(r3, b), \\ \neg in(r3, c), \neg in(r3, d), \neg in(r3, f)\}$$

The NM has some surveillance cameras to monitor the movement of robots in the building. Specifically, the corridors that are monitored are the ones between offices: a and b; b and c; and b and f. These corridors are represented by black arrows in Figure 1, whereas non-monitored corridors are represented by grey arrows. In the initial state (s_0) depicted in Figure 1a, the robots execute the actions $move(r1, a, b)$, $move(r2, d, a)$ and $move(r3, e, a)$ resulting in a new state (s_1) depicted in Figure 1b. However, the NM only observes the action of robot r1, because this action takes place in a monitored corridor; i.e., $Act_0 = \{move(r1, a, b)\}$. In the next state s_1 , the robots execute actions $move(r1, b, c)$, $move(r2, a, e)$ and $move(r3, a, b)$ resulting in a new state (s_2) depicted in Figure 1c. In this case the NM observes two actions; i.e., $Act_1 = \{move(r1, b, c), move(r3, a, b)\}$. Considering these two sets of observed actions¹³ the NM is able to infer the dynamic propositions that are known in s_1 as follows:

$$p_1 = \{in(r1, b), \neg in(r1, a), in(r3, a)\}$$

If the NM uses the information about the states and the observed actions, then no violation of the norm is detected and

¹¹However, our model can be used by a team of monitors as well.

¹²This assumption could be relaxed if NMs have capabilities for observing both state changes and actions.

¹³Inference rules might also be used together with observed actions to infer more propositions. However, we propose to conduct this inference process once unobserved actions have been reconstructed to increase the number of propositions that are inferred, as described in Section III-B.

no robot is sanctioned. However, $r2^{14}$ and $r3$ have violated the norm, since they have moved into an occupied office through non-monitored corridors.

B. Action Reconstruction

NMs use Definition 6 to generate partial state descriptions based on the observed actions. Additionally, we propose that NMs reconstruct the actions that have not been observed. This reconstruction process entails: (i) searching for the actions that have been performed by unobserved agents; and (ii) using the actions found to increase the knowledge about the state of the world. The reconstruction process must be sound, e.g., it cannot indicate that a violation has occurred when it has not in fact occurred. In the following, we introduce full and approximate methods for reconstructing unobserved actions.

1) *Full Reconstruction*: Full reconstruction tries to find exhaustively the actions performed by all the agents that have not been observed. To this aim, the full reconstruction performs a search to identify all solutions to the reconstruction problem.

Definition 7. Given a partial state description p_t corresponding to time t (named initial state), a set of observed actions Act_t at time t , and an partial resulting state p_{t+1} corresponding to time $t + 1$ (named final state); we define search as a function that computes sets of solutions $\mathcal{S} = \{S_1, \dots, S_k\}$ such that each solution S_i in \mathcal{S} is a set of actions such that:

- the concurrent action $S_i \cup Act_t$ is consistent;
- the initial state induced by the concurrent action $S_i \cup Act_t$ is consistent (i.e., $g, p_t, pre(S_i \cup Act_t), \nabla \not\vdash \perp$);
- the final state induced by the concurrent action $S_i \cup Act_t$ is consistent (i.e., $g, p_{t+1}, post(S_i \cup Act_t), \nabla \not\vdash \perp$).

Thus, a solution is a set of actions performed by the agents that have not been observed¹⁵ that are consistent with the states of the world before and after the execution of the actions. Given that the NM has a partial knowledge of the states, we do not require that the preconditions (vs. postconditions) of actions in a solution are met in the initial (vs. final) state, since it is possible that the preconditions (vs. postconditions) are true, but the NM is unaware of it.

EXAMPLE 8. Given the partial state description p_0 (initial state), the set of observed actions Act_0 , and the partial resulting state p_1 (final state), the search function looks for actions of agents $r2$ and $r3$ (since they are the agents that have not been observed). According to the initial position of $r2$, the NM can infer that $r2$ may have performed two different actions $move(r2, d, a)$ and $move(r2, d, e)$ —these two actions are the only ones consistent with p_0 . Similarly, the NM can infer that $r3$ may have performed three different actions $move(r3, e, a)$, $move(r3, e, d)$ and $move(r3, e, f)$ —these three actions are the only ones consistent with p_0 . However, the actions $move(r3, e, d)$ and $move(r3, e, f)$ are

not consistent with the final state —recall that these two actions have as postcondition the fact that $r3$ is in offices d and f , respectively; that p_1 states that $r3$ is in office a (i.e., p_1 was defined as $\{in(r1, b), \neg in(r1, a), in(r3, a)\}$ in Example 7); and that ∇ defines as inconsistent states where any robot is at two different locations. As a result, the solution set for this problem is defined as:

$$\mathcal{S} = \{\{move(r2, d, a), move(r3, e, a)\}, \\ \{move(r2, d, e), move(r3, e, a)\}\}$$

Once all solutions are found, the NM uses this information to extend the information about the actions performed by unobserved agents and the state of the world. To ensure that the reconstruction is sound, the NM calculates the intersection of actions in the solutions to select actions it is completely sure about (i.e., actions belonging to all solutions). Given a set of search solutions $\mathcal{S} = \{S_1, \dots, S_k\}$ for some initial and final states, we define the *reconstruction* action set as follows:

$$R = \bigcap_{\forall S_i \in \mathcal{S}} S_i$$

If $R \neq \emptyset$, then the NM expands its knowledge about the actions performed by agents and it uses this information to increase the knowledge about the initial and final states. More formally, the set of actions observed in t is updated as:

$$Act_t = Act_t \cup R$$

The initial state p_t is updated as follows:

$$p_t = p_t \bigcup pre(R)$$

Finally, the final state is updated as follows:

$$p_{t+1} = \begin{cases} p_{t+1} \bigcup post(R) \bigcup p_t^\bullet & \text{if } |Act_t| < |Ag| \\ p_{t+1} \bigcup p_t^\bullet \bigcup eff(Act_t) & \text{otherwise} \end{cases}$$

where p_t^\bullet is defined as before and p_t^\bullet is the set of extended invariant literals; i.e., literals in p_t that have not been modified since there is not a solution $S_i \in \mathcal{S}$ such that the concurrent action $S_i \cup Act_t$ changes any of these literals:

$$p_t^\bullet = \bigcup_{\substack{\forall l \in p_t: \\ \nexists S_i \in \mathcal{S}: l, post(Act_t \cup S_i) \not\vdash \perp}} l$$

EXAMPLE 9. The reconstruction set for the example is:

$$R = \{move(r3, e, a)\}$$

This action belongs to all solutions, so the NM can be absolutely sure about the performance of this action, even when the NM has not observed it. As a consequence, the NM extends its information as follows:

$$Act_0 = \{move(r1, a, b), move(r3, e, a)\}$$

and p_0 remains unchanged and p_1 is updated as follows:

$$p_1 = \{\neg in(r1, b), \neg in(r1, a), \neg in(r1, c), \neg in(r1, d), \neg in(r1, e), \\ \neg in(r1, f), \neg in(r2, b), \neg in(r2, c), \neg in(r2, f), in(r3, a), \\ \neg in(r3, b), \neg in(r3, c), \neg in(r3, d), \neg in(r3, e), \neg in(r3, f)\}$$

¹⁴Note that proposition $\neg in(r2, d)$ cannot be inferred since our mechanism can also be used in scenarios with the NOP action and, as a result, it cannot assume that agents must move in each step.

¹⁵If all actions were observed, no reconstruction would be needed.

The main disadvantage of full reconstruction is that, for many real-world problems, the number of candidate solutions that needs to be explored is prohibitively large, as shown later in Section IV. In response to this problem, we provide a polynomial approximation below.

2) *Approximate Reconstruction*: Approximate reconstruction includes an approximate search that finds the actions performed by unobserved agents that are consistent with the states of the world before and after action execution. Specifically, approximate reconstruction identifies actions that do not necessarily include the specific actions performed by unobserved agents but that allow the NM to control norms. The main intuition behind approximate reconstruction is as follows: imagine that at a given initial state an agent can perform just one action and that this action is forbidden (vs. mandatory). In this case, the NM *identifies* that the agent has violated (vs. fulfilled) a norm. Besides that, if an agent can perform n different actions and all these actions are forbidden (vs. mandatory), the NM does not need to know which action has been executed to conclude that a norm has been violated (vs. fulfilled)¹⁶. Hence, we say that a violation has been *discovered* (instead of *identified*). Given a set of prohibition instances P and an action a , we define that the action a is forbidden (denoted by $forbidden(P, a)$) when $\exists p \in P : \exists \sigma : \sigma \cdot action(p) = a$. Similarly, given a set of obligation instances O and an action a , we define that the action a is mandatory (denoted by $mandatory(O, a)$) when $\exists o \in O : \exists \sigma : \sigma \cdot action(o) = a$.

Definition 8. Given a partial state p_t , a set of observed actions Act_t at time t , and a partial resulting state p_{t+1} ; we define approximate search as a function that calculates the set of all unobserved applicable actions $\tilde{S} = \{a_i, \dots, a_n\}$ such that:

- the preconditions of each action in \tilde{S} are consistent with the initial state (i.e., $\forall a_i \in \tilde{S} : g, p_t, pre(a_i), \nabla \not\vdash \perp$);
- the postconditions of each action in \tilde{S} are consistent with the final state (i.e., $\forall a_i \in \tilde{S} : g, p_{t+1}, post(a_i), \nabla \not\vdash \perp$);
- actions in \tilde{S} are performed by unobserved agents (i.e., $actor(\tilde{S}) \cap actor(Act_t) = \emptyset$);
- all unobserved agents perform at least one action in \tilde{S} .

EXAMPLE 10. Given the partial state description p_0 (initial state), the set of observed actions Act_0 , and the partial resulting state p_1 (final state), the approximate search function looks for actions of agents $r2$ and $r3$ (since they are the agents that have not been observed). According to the initial position of $r2$, the NM can infer that $r2$ may have performed two different actions $move(r2, d, a)$ and $move(r2, d, e)$. Again, $r3$ may have performed action $move(r3, e, a)$ —recall that actions $move(r3, e, d)$ and $move(r3, e, f)$ are inconsistent with the final state p_1 since it

¹⁶Note that the purpose of this paper is to monitor norms, not to determine whether agents are responsible for norm violations/fulfilments. Monitoring situations where agents can only execute forbidden/obligatory actions can help to detect norm-design problems. Additionally, the fact that an agent can only execute forbidden actions may be explained by the agent putting itself into these illegal situations (e.g., I am allowed to overtake but overtaking may put me in a situation where I can only exceed the speed limit).

was defined as $\{in(r1, b), -in(r1, a), in(r3, a)\}$ in Example 7. The approximate solution for this problem is defined as:

$$\tilde{S} = \{move(r2, d, a), move(r2, d, e), move(r3, e, a)\}$$

As in full reconstruction, the NM uses approximate search solutions (\tilde{S}) to expand its knowledge about the actions performed by unobserved agents and to increase the knowledge about the initial and final states. When an unobserved agent may have executed only one action, then the NM knows for sure that this action was executed. More formally, the *reconstruction* action set is defined as follows:

$$R = \bigcup_{\substack{\forall a \in \tilde{S}: \exists a' \in \tilde{S}: \\ a \neq a' \wedge actor(a) = actor(a')}} a$$

The set of actions observed in t is updated as:

$$Act_t = Act_t \cup R$$

Then the initial state p_t is updated as follows:

$$p_t = p_t \bigcup pre(R)$$

The final state is updated as follows:

$$p_{t+1} = \begin{cases} p_{t+1} \bigcup post(R) \bigcup p_t^o & \text{if } |Act_t| < |Ag| \\ p_{t+1} \bigcup p_t^* \bigcup eff(Act_t) & \text{otherwise} \end{cases}$$

where p_t^* is defined as before and p_t^o is the set of extended invariant literals in p_t ; i.e., literals that have not been modified since there is not an observed action or an applicable action that changes them:

$$p_t^o = \left(\bigcup_{\substack{\forall l \in p_t: \\ l, post(Act_t) \not\vdash \perp}} l \right) \cap \left(\bigcup_{\substack{\forall l \in p_t: \\ l, post(\tilde{S}) \not\vdash \perp}} l \right)$$

Finally, the set of discovered violations and fulfilments is a set of actions:

$$D = \{a_1, \dots, a_j\}$$

where for each action a in D : a is in \tilde{S} and the set of unobserved applicable actions for the agent that performs a (i.e., $\{a' : a' \in \tilde{S} \wedge actor(a') = actor(a)\}$) is formed by two or more forbidden (vs. mandatory) actions; and a is one of these forbidden (vs. mandatory) actions.

When an agent is only able to perform forbidden (vs. mandatory) actions, an action among these can be selected according to various criteria. For example, in a normative system where the presumption of innocence principle holds, the NM should assume that the agent has violated (vs. fulfilled) the least (vs. most) important norm and the action that violates (vs. fulfils) this norm is selected. Note that discovering violations is very useful in many practical applications, in which it would allow the NM to ban offender agents (e.g., Intrusion Detection/Prevention Systems [4]), to stop the execution of any offender agent (e.g., Business Process Compliance monitoring [29]), or to put offender agents under close surveillance (e.g.,

Model-Based Diagnosis Systems [24]), even when the specific action performed is not known.

EXAMPLE 11. *In case of the approximate reconstruction, r3 is only able to perform one action, which entails that the NM can be absolute sure about the performance of this action and the reconstruction set is defined as:*

$$R = \{\text{move}(r3, e, a)\}$$

As a consequence, the NM extends its information as follows:

$$\text{Act}_0 = \{\text{move}(r1, a, b), \text{move}(r3, e, a)\}$$

p_0 remains unchanged and p_1 is updated as follows:

$$p_1 = \{\text{in}(r1, b), \neg\text{in}(r1, a), \neg\text{in}(r1, c), \neg\text{in}(r1, d), \neg\text{in}(r1, e), \\ \neg\text{in}(r1, f), \neg\text{in}(r2, b), \neg\text{in}(r2, c), \neg\text{in}(r2, f), \text{in}(r3, a), \\ \neg\text{in}(r3, b), \neg\text{in}(r3, c), \neg\text{in}(r3, d), \neg\text{in}(r3, e), \neg\text{in}(r3, f)\}$$

In this situation, r2 is only able to execute forbidden actions —recall that the instances $\langle \mathcal{P}, \text{move}(R2, L1, a) \rangle$ and $\langle \mathcal{P}, \text{move}(R2, L1, e) \rangle$ forbid any robot to move into offices a and e and that r2 may have executed actions $\text{move}(r2, d, a)$ and $\text{move}(r2, d, e)$ ¹⁷. Thus, the set of discovered violations and fulfilments can be defined as either containing $\text{move}(r2, d, a)$ or $\text{move}(r2, d, e)$. Note that the discovered violation does not correspond to the action executed by r2, however, it allows the NM to determine that r2 must have violated an instance.

C. Norm Monitoring

Once all the information about the actions performed by the agents and the partial states has been reconstructed, the NM checks the actions of agents to determine which instances have been violated or fulfilled. Recall that norms in our model are defined as conditional rules that state which actions are obligatory or forbidden. Given that the NM has partial knowledge about the state of the world, the NM should control norms only when it is completely sure that the norms are relevant to ensure that the norm monitoring process is sound. In particular, we define that a norm is relevant to a partial situation when the norm condition is satisfied by the partial situation —i.e., a norm $\langle \text{deontic}, \text{condition}, \text{action} \rangle$ is relevant to a partial situation represented by a partial state p , the static properties g and the domain knowledge ∇ if $\exists \sigma$ such that $p, g, \nabla \vdash \sigma \cdot \text{condition}$.

EXAMPLE 12. *In state p_0 the norm that forbids robots to move into occupied offices is instantiated three times as follows:*

$$\langle \mathcal{P}, \text{move}(R2, L1, d) \rangle \text{ where } \sigma = \{L2/d\} \\ \langle \mathcal{P}, \text{move}(R2, L1, a) \rangle \text{ where } \sigma = \{L2/a\} \\ \langle \mathcal{P}, \text{move}(R2, L1, e) \rangle \text{ where } \sigma = \{L2/e\}$$

Once the NM has determined which norm instances hold in a given situation, it has to check the actions of agents to

¹⁷Note forbidden actions do not necessary lead to a collision since neither robot r1 nor r3 have moved to office d.

determine which instances have been violated and which ones have been fulfilled.

Obligation Instance. In presence of partial knowledge about the actions performed by agents, the NM can only determine that an obligation instance has been fulfilled. If the NM knows all the actions performed by agents, then it can determine whether an obligation has been fulfilled or violated.

Definition 9. *Given an obligation instance $\langle \mathcal{O}, \text{action}' \rangle$ and a set of observed actions Act , then the obligation is defined as:*

$$\begin{cases} \text{fulfilled} & \text{iff } \exists \sigma : \sigma \cdot \text{action}' \in \text{Act} \\ \text{violated} & \text{iff } (\nexists \sigma : \sigma \cdot \text{action}' \in \text{Act}) \wedge |\text{Act}| = |\text{Ag}| \\ \text{unknown} & \text{otherwise} \end{cases}$$

Prohibition Instance. In presence of partial knowledge about the actions performed by agents, the NM can only determine that a prohibition instance has been violated. If the NM knows all the actions performed by agents then it can determine whether a prohibition has been fulfilled or violated.

Definition 10. *Given a prohibition instance $\langle \mathcal{P}, \text{action}' \rangle$ and a set of observed actions Act , then the prohibition is defined as:*

$$\begin{cases} \text{violated} & \text{iff } \exists \sigma : \sigma \cdot \text{action}' \in \text{Act} \\ \text{fulfilled} & \text{iff } (\nexists \sigma : \sigma \cdot \text{action}' \in \text{Act}) \wedge |\text{Act}| = |\text{Ag}| \\ \text{unknown} & \text{otherwise} \end{cases}$$

Finally, the set of discovered violations and fulfilments is used to identify those agents that have violated or fulfilled an instance.

EXAMPLE 13. *Taking into account the set of actions Act_0 , the NM can identify that robot r3 has violated the instance $\langle \mathcal{P}, \text{move}(R2, L1, a) \rangle$, even though this forbidden action has not been observed by the NM. Specifically, there is $\sigma = \{R2/r3, L1/e\}$ such that $\sigma(\text{move}(R2, L1, a)) \in \text{Act}_0$. Besides that, the approximate reconstruction discovers that robot r2 has violated a prohibition instance though it does not know the exact action performed —recall that either $\text{move}(r2, d, a)$ or $\text{move}(r2, d, e)$ are in D . Had the NM not performed the proposed reconstruction processes, none of these violations would have been detected.*

IV. NM ALGORITHMS

Algorithm 1 contains the NM pseudocode. In each step, the NM observes the actions of agents and uses this information to update the current and the previous partial states (lines 4-9). If all the actions have not been observed in the previous state, then the NM executes the *reconstruction* algorithm to reconstruct unobserved actions (lines 11-14). Then, the *checkNorms* function is executed to determine which norms have been violated and fulfilled in the previous state (line 15) according to Definitions 9 and 10.

Note that the NM code can be executed while actions are performed without delaying agents. Regarding the temporal cost of the algorithm executed by NMs, it is determined by the cost of the *reconstruction* algorithm, the implementations of which (full and approximate) are discussed below.

Algorithm 1 NM Algorithm

Require: Ag, N, D, ∇, g

```

1:  $p_0 = \emptyset$   $\triangleright p_0$  is an empty conjunction of literals
2:  $t \leftarrow 0$ 
3: while true do
4:    $Act_t \leftarrow observeActions()$ 
5:   if  $|Act_t| < |Ag|$  then
6:      $p_{t+1} \leftarrow post(Act_t)$ 
7:   else
8:      $p_{t+1} \leftarrow p_t^* \wedge eff(Act_t)$ 
9:    $p_t \leftarrow p_t \wedge pre(Act_t)$ 
10:  if  $t > 0$  then
11:    if  $|Act_{t-1}| < |Ag|$  then
12:       $TA \leftarrow Ag \setminus actors(Act_{t-1})$   $\triangleright$  Target Agents
13:       $D \leftarrow \emptyset$   $\triangleright$  Discovered violations and fulfilments
14:       $reconstruction(p_{t-1}, p_t, Act_{t-1}, TA, D)$ 
15:       $checkNorms(p_{t-1}, Act_{t-1})$ 
16:     $t \leftarrow t + 1$ 

```

Full Reconstruction (Algorithm 2). This pseudocode corresponds to the full reconstruction algorithm. This algorithm calls the function *search* to search the actions of target agents (line 2). Then, for all the solutions found, the NM checks if they are consistent according to Definition 7 (lines 4-6). Finally, consistent solutions are used to extend the set of observed actions and the knowledge about the initial and final states (lines 7-14).

Algorithm 2 Full Reconstruction Algorithm

```

1: function FULLRECONSTRUCTION( $i, f, Act, TA, D$ )
2:    $S' \leftarrow search(i, f, Act, TA)$   $\triangleright$  Candidate Solutions
3:    $S \leftarrow \emptyset$   $\triangleright$  Consistent Solutions
4:   for all  $S_j \in S'$  do
5:     if  $checkSolutionConsistency(Act, S_j, i, f)$  then
6:        $S \leftarrow S \cup S_j$ 
7:    $R \leftarrow \bigcap_{S_j \in S} S_j$ 
8:   if  $R \neq \emptyset$  then
9:      $Act \leftarrow Act \cup R$ 
10:     $i \leftarrow i \wedge pre(R)$ 
11:    if  $|Act| < |Ag|$  then
12:       $f \leftarrow f \wedge post(R) \wedge i^*$ 
13:    else
14:       $f \leftarrow i^* \wedge f \wedge eff(Act)$ 

```

Algorithm 3 contains the pseudocode of the recursive *search* function that computes all the sequences of consistent actions that may have been executed by the agents that have not been observed. It starts by checking that there is at least one target agent (line 2). If so, it identifies all actions that might have been executed by one target agent (lines 3-4). An action might have been executed if it is consistent according to the static properties, the domain knowledge, and the initial and final states. For each consistent action, it reconstructs the actions of the remaining agents recursively (lines 5-13).

Completeness of the full reconstruction algorithm is achieved as the NM considers all possible combinations of unobserved actions. The termination of the algorithm is guaranteed as agents, actions, predicate and constant symbols are finite. In the worst case, the temporal cost of this algorithm is $O(|Ag|^{|D| \times I_D})$, where Ag is the set of agents, D is the set of action descriptions and I_D is the maximum number of

instantiations per action. This situation arises when no action is observed and all actions are applicable for all agents —note the asymptotic cost of the reconstruction algorithm is determined by the search function. Thus, the full reconstruction algorithm is a factorial algorithm, which entails that the algorithm may not be feasible even for reasonable-sized inputs.

Algorithm 3 Search Function

```

1: function SEARCH( $i, f, Act, TA$ )
2:   if  $TA \neq \emptyset$  then
3:     for all  $d \in D$  do  $\triangleright$  Identify consistent actions
4:       if  $\exists \sigma : checkActionConsistency(\sigma \cdot d, i, f) \wedge$   

 $actor(\sigma \cdot d) \in TA$  then
5:          $\alpha = actor(\sigma \cdot d)$ 
6:          $i' \leftarrow i \wedge pre(\sigma \cdot d)$ 
7:          $f' \leftarrow f \wedge post(\sigma \cdot d)$ 
8:          $Act' \leftarrow Act \cup \sigma \cdot d$ 
9:          $TA' \leftarrow TA \setminus \{\alpha\}$ 
10:         $S \leftarrow search(i', f', Act', TA')$ 
11:        for all  $S_i \in S$  do
12:           $S_i \leftarrow S_i \cup \sigma \cdot d$ 
13:        return  $S$ 
14:   else
15:     return  $\emptyset$ 

```

Approximate Reconstruction (Algorithm 4). This algorithm calls the function *ApproximateSearch* to search the applicable actions per each target agent (line 2). Then, the list of applicable actions per each agent is checked (lines 3-12). Specifically, if an agent may have executed one action only, then the NM knows that this action was executed and it updates the reconstructed action set (lines 3-5). Then, the set of observed actions and the knowledge about the initial and final states is updated (lines 6-12). Finally, discovered violations and fulfilments are calculated (lines 14-19).

Algorithm 5 contains the pseudocode of the *ApproximateSearch* function. It starts by initialising the list of applicable actions per agent (lines 3-4). Then it calculates the set of instances that are relevant to the initial state (line 7). The function calculates per each target agent the list of applicable actions that it may have executed (lines 9-11). Then, the list of applicable actions per each agent is checked (lines 12-18). Specifically, if an agent may have executed one action only, then the NM knows that this action was executed and it updates the list of applicable actions, the initial and final states, and retracts the agent from the target agents (lines 14-17). This process is repeated until there are no more target agents or the initial and final states remain unchanged. Then the set of instances that are relevant to the initial state is calculated (line 13). Finally, the list of applicable actions per agent is updated with actions of remaining target agents (lines 19-20).

Completeness of the approximate reconstruction algorithm is not achieved as the NM does not consider all possible combinations of unobserved actions. However, the termination of the algorithm is guaranteed as agents, actions, predicate and constant symbols are finite. The temporal cost of this algorithm is $O(|Ag|^2 \times |D| \times I_D)$. Thus, the reconstruction of the approximate sequence of actions that matches partial

initial and final states can be computed by a feasible algorithm in polynomial time.

Algorithm 4 Approximate Reconstruction Algorithm

```

1: function APPROXIMATERECONSTRUCTION( $i, f, Act, TA, D$ )
2:    $\tilde{S} \leftarrow \text{approximateSearch}(i, f, Act, TA)$ 
3:   for all  $\alpha \in TA$  do
4:     if  $|\tilde{S}_\alpha| = 1$  then
5:        $R \leftarrow R \cup S_\alpha$ 
6:   if  $R \neq \emptyset$  then
7:      $Act \leftarrow Act \cup R$ 
8:      $i \leftarrow i \wedge \text{pre}(R)$ 
9:     if  $|Act| < |Ag|$  then
10:       $f \leftarrow f \wedge \text{post}(R) \wedge i^\circ$ 
11:     else
12:       $f \leftarrow i^* \wedge f \wedge \text{eff}(Act)$ 
13:    $O, P \leftarrow \text{calculateInstances}(i)$ 
14:   for all  $\alpha \in TA$  do
15:     if  $|\tilde{S}_\alpha| > 1$  then
16:       if  $\exists a \in \tilde{S}_\alpha : \neg \text{forbidden}(P, a)$  then
17:          $D \leftarrow D \cup a$   $\triangleright a$  is an action from  $\tilde{S}_\alpha$ 
18:       else if  $\exists a \in \tilde{S}_\alpha : \neg \text{mandatory}(O, a)$  then
19:          $D \leftarrow D \cup a$   $\triangleright a$  is an action from  $\tilde{S}_\alpha$ 

```

Algorithm 5 Approximate Search Function

```

1: function APPROXIMATESEARCH( $i, f, Act, TA$ )
2:    $continue \leftarrow true$ 
3:   for all  $\alpha \in TA$  do
4:      $\tilde{S}_\alpha \leftarrow \emptyset$   $\triangleright$  List of approximate actions per agent
5:   while  $continue \wedge TA \neq \emptyset$  do
6:      $continue \leftarrow false$ 
7:     for all  $\alpha \in TA$  do
8:        $L_\alpha \leftarrow \emptyset$ 
9:       for all  $d \in D$  do
10:      if  $\exists \sigma : \text{checkActionConsistency}(\sigma \cdot d, i, f) \wedge$   

 $actor(\sigma \cdot d) \in TA$  then
11:         $L_{actor(\sigma \cdot d)} \leftarrow L_{actor(\sigma \cdot d)} \cup \sigma \cdot d$ 
12:      for all  $\alpha \in TA$  do
13:        if  $|L_\alpha| = 1$  then
14:           $\tilde{S}_\alpha \leftarrow L_\alpha$ 
15:           $i \leftarrow i \wedge \text{pre}(L_\alpha)$ 
16:           $f \leftarrow f \wedge \text{post}(L_\alpha)$ 
17:           $TA \leftarrow TA \setminus \{\alpha\}$ 
18:           $continue \leftarrow true$ 
19:   for all  $\alpha \in TA$  do
20:      $\tilde{S}_\alpha \leftarrow L_\alpha$ 
21:   return  $\tilde{S}$ 

```

V. EVALUATION

This section compares the performance of a NM with full reconstruction, a NM with approximate reconstruction and a traditional norm monitor—which is the method used in the majority of previous proposals [8], [25], [16], [26], [11]—that only considers the observed actions to detect violations, with respect to their capabilities to monitor norm compliance. We have evaluated our proposal in a case study, which allows us to contextualise the results and to give a meaningful interpretation to them; and in a series of random experiments, which allow us to evaluate our proposal under a wide range of different situations and parameter values.

A. Case Study

We implemented in Java a simulator of the paper example in which robots attend requests in offices connected through corridors. Compliance with the collision avoidance norm is controlled by a monitor that observes surveillance cameras. In each simulation, we generate corridors and cameras randomly. In each step of the simulation, each robot chooses randomly one applicable action to be executed. The simulation is executed 100 steps and repeated 100 times to support the findings. We conducted experiments in which the number of offices O took a random value within the $\llbracket 3, 500 \rrbracket$ interval and the number of robots R took a random value within the $\llbracket 2, 250 \rrbracket$ interval. Besides that, to be able to compare with the full NM, we also considered small scenarios only, in which the number of offices O takes a random value within $\llbracket 3, 10 \rrbracket$ and the number of robots R takes a random value within $\llbracket 2, 5 \rrbracket$, as the full reconstruction has an exponential cost and it is intractable for most of the cases with the default intervals.

1) *Action Observability*: To analyse the performance and scalability of monitors with respect to their capabilities to observe actions, we defined the number of corridors C as a random value within the $\llbracket O, O \times (O - 1) \rrbracket$ interval and varied the ratio of cameras to corridors (action observability). Table I shows the percentage of violations detected per each type of monitor. The higher the ratio of cameras, the more actions are observed and the better the performance of all monitors. Moreover, the approximate NM offers on average a 39% performance improvement over a traditional monitor (i.e., it identifies 16% more violations plus a further 24% of discovered violations). That is, an approximate NM outperforms a traditional monitor with the same capabilities to observe actions. When compared to full NM in small scenarios ($O \in \llbracket 3, 10 \rrbracket$ and $R \in \llbracket 2, 5 \rrbracket$), approximate NM performs similarly. This is explained by the fact that there is a single norm in this scenario, actions have no concurrency conditions, and the preconditions and postconditions of actions are disjoint. In these circumstances, the approximate reconstruction process reconstructs actions similarly to the full reconstruction¹⁸.

Cameras Ratio	Traditional Monitor	Approximate NM Identify+Discover	Cameras Ratio	Traditional Monitor	Full NM	Approximate NM Identify+Discover
0%	0%	0+0%	0%	0%	0%	0+0%
20%	11%	14+9%	20%	16%	32%	32+6%
40%	31%	40+12%	40%	32%	68%	67+5%
60%	55%	71+10%	60%	56%	88%	88+3%
80%	78%	91+4%	80%	76%	99%	99+0%
100%	100%	100+0%	100%	100%	100%	100+0%

$O \in \llbracket 3, 500 \rrbracket$ and $R \in \llbracket 2, 250 \rrbracket$ $O \in \llbracket 3, 10 \rrbracket$ and $R \in \llbracket 2, 5 \rrbracket$

TABLE I: Action Observability Experiment

2) *Action Instantiations*: To analyse the performance and scalability of monitors with respect to agent capabilities to execute actions (i.e., the number of instantiations per action), we varied the ratio of corridors¹⁹ (e.g., a ratio of 0% means $C = O$) and defined the number of cameras as a random value within the $\llbracket 0, C \rrbracket$ interval. Table II shows the results of this experiment. The approximate NM offers on average a 43% performance improvement over a traditional monitor (i.e., it

¹⁸Note that full reconstruction does not guarantee completeness.

¹⁹Recall that C takes values within the $\llbracket O, O \times (O - 1) \rrbracket$ interval.

identifies 29% more violations plus a further 14% of discovered violations). That is, given the same number of possible instantiations per action, an approximate NM outperforms a traditional monitor. Besides, we can see that, as in the previous experiment, the approximate NM performs similarly to the full NM. In particular, when the ratio of corridors is higher than 0%, agents are capable of executing different actions and the reconstruction process becomes more complex, which decreases the performance of full and approximate NMs. However, full and approximate NMs noticeably outperform the traditional monitor regardless of the ratio of corridors.

Corridors Ratio	Traditional Monitor	Approximate NM Identify+Discover	Corridors Ratio	Traditional Monitor	Full NM	Approximate NM Identify+Discover
0%	51%	98+0%	0%	52%	99%	99+0%
20%	48%	55+6%	20%	59%	80%	79+3%
40%	48%	56+7%	40%	55%	74%	74+3%
60%	41%	47+6%	60%	51%	70%	69+4%
80%	49%	56+13%	80%	55%	68%	68+5%
100%	42%	49+7%	100%	57%	70%	69+4%

$O \in \llbracket 3, 500 \rrbracket$ and $R \in \llbracket 2, 250 \rrbracket$ $O \in \llbracket 3, 10 \rrbracket$ and $R \in \llbracket 2, 5 \rrbracket$

TABLE II: Action Instantiations Experiment

B. Random Experiments

We implemented a simulator in Java in which there is a set of agents that perform actions in a monitored environment as defined below. In particular, our simulator does not model a specific scenario; rather it creates a different scenario in each simulation (i.e., generating randomly agent capabilities, the environment properties, actions and norms). As in the previous experiments, we have considered big and small scenarios. In particular, the number of agents G in small scenarios took a random value within the $\llbracket 1, 5 \rrbracket$ interval, whereas in big scenarios G took a random value within the $\llbracket 1, 500 \rrbracket$ interval. The number of actions A took a random value within the $\llbracket 1, 50 \rrbracket$ interval. Again, the simulation is executed 100 steps and repeated 1000 times to ensure that the values of the simulation parameters range over possible values²⁰.

Agent Definition. We modelled different types of agents with different capabilities to perform actions. In particular, the set of actions available to each agent depends on the function/s assumed by each agent in a particular simulation. To model these capabilities, a set of roles is created at the beginning of each simulation. Specifically, the number of roles created took a random value within the $\llbracket 1, A \rrbracket$ interval. For each role a subset of the actions are randomly selected as the role capabilities; i.e., all agents enacting this role are able to perform these actions²¹. To avoid that all roles have similar capabilities, which would lead to simulations populated by homogeneous agents, the number of actions selected as role capabilities took a random value within the $\llbracket 1, \lceil 0.1 * A \rceil \rrbracket$ interval (i.e., at maximum each role is capable of performing a 10% of the actions). At the beginning of each simulation, each agent is defined as enacting a random subset of the roles. In

²⁰Note that in the random experiments there are more simulation parameters than in the case-study simulator and a higher number of repetitions is required to support the findings.

²¹This condition has been formulated in action preconditions as explained below.

each step of the simulation, each agent selects randomly one action among the available actions that it can execute in the current state.

Environment Definition. In the simulator, the environment is described in terms of different situations or states of affairs that can be true or false. Each one of these states of affairs is represented by a grounded proposition. Thus, the state of the environment is defined in terms of a set of propositions. For simplicity, we assumed that these propositions are independent (i.e., propositions are not logically related). In our simulations, the number of propositions P took a random value within the $\llbracket A, 2 * A \rrbracket$ interval (i.e., there is at least one proposition per each action²²). Besides that, there is a set of grounded atomic formulas describing the roles played by agents and the actions that can be performed by each role. The relationship between agents and roles is formally represented by a binary predicate (*play*). Specifically, the expression $play(g, r)$ describes the fact that the agent identified by g enacts the role identified by r . Similarly, relationship between roles and actions is formally represented by a binary predicate (*capable*). Specifically, the expression $capable(a, r)$ describes the fact that agents enacting role r are capable of performing the action identified by a . For simplicity, we assume that the roles enacted by the agents and the role capabilities are static properties of the environment.

Action Definition. Actions allow agents to change the state of the environment. At the beginning of each simulation, a set of actions is randomly generated. For each action $\langle name, pre, con, post \rangle$ the elements are defined as follows: *name* is initialised with a sequential identifier a ; *pre* is defined as $\{play(A, R), capable(R, a), p_1, \dots, p_n\}$ where the elements p_1, \dots, p_n are randomly selected from the proposition set; *con* is defined as $\{a_1(A_1, R_1), \dots, a_m(A_m, R_m)\}$, where each a_i is an action randomly selected from the action set such that $a_i \neq a$ and A_i, R_i are free variables representing the agent performing the action and the role capable of performing this action, respectively; and *post* is defined as $\{p_1, \dots, p_k\}$ where each p_i is a proposition randomly selected from the proposition set. To avoid that actions have too many constraints, which would be unrealistic and make actions to be only executed on few situations, the number of propositions in *pre* and *post* takes a random value within the $\llbracket 1, \lceil 0.1 * P \rceil \rrbracket$ interval. Similarly, the number of actions in *con* takes a random value within the $\llbracket 0, \lceil 0.1 * A \rceil \rrbracket$ interval.

Besides these actions, a NOP action, which has no effect on the environment, was created. Our simulator models scenarios where the NOP action can always be observed. This is the case in many real domains such as Intrusion Detection Systems or Autonomous Systems, where it is not always possible to analyse the data (e.g., the packages) sent by agents (e.g., hosts) to infer the actions performed, but it is always possible to know which agents have performed an action (i.e., which agents have sent packages). To maximise the number of actions executed in the simulations, which may entail more violations and fulfilments, we defined that the NOP action can only be executed by agents when none of their available actions can be

²²Note that an action can change the truth value of several propositions.

executed. However, similar results would have been obtained if this condition was relaxed²³.

Norm Definition. Agents’ actions are regulated by a set of norms. At the beginning of each simulation, a set of norms is randomly created. In particular the number of norms took a random value within the $\llbracket 1, A \rrbracket$ (i.e., there is at maximum one norm per each action). For each norm $\langle deontic, condition, action \rangle$ the elements are defined as follows: *deontic* is randomly initialised with a deontic operator; *condition* is defined as $\{p_1, \dots, p_k\}$ where each p_i is a proposition randomly selected from the proposition set; and *action* is randomly initialised with an action. To allow norms to be instantiated, the number of propositions in *condition* takes a random value within the $\llbracket 0, \lceil P * 0.1 \rceil \rrbracket$ interval.

1) *Action Observability:* To analyse the performance and scalability of monitors with respect to their capabilities to observe actions, we varied the observation probability. Tables III and IV show the percentage of detected fulfilments and violations, respectively. Again, the approximate NM offers a significant performance improvement over a traditional monitor; i.e., the approximate NM offers on average a 74% performance improvement over a traditional monitor. When compared to full NM in small scenarios ($A \in \llbracket 1, 50 \rrbracket$ and $G \in \llbracket 1, 5 \rrbracket$), the full NM offers on average a 21% performance improvement over an approximate NM. This is explained by the fact that this experiment is more complex than the case study; i.e., there are several norms (both prohibition and obligation norms), actions have concurrent conditions and actions may have conflicting preconditions and postconditions (i.e., conditions that are defined over the same propositions). Note that the traditional monitor detects violations and fulfilments even when the observation probability is 0%. These detections correspond to situations in which none of the agents can execute any action (i.e., all agents execute the NOP action) which leads to the fulfilment of prohibition instances and the violation of obligation instances. This phenomenon is more frequent in case of small scenarios since the lower the number of agents, the higher the probability that all agents cannot execute any action.

Observ. Prob.	Traditional Monitor	Approximate NM Identify+Discover	Observ. Prob.	Traditional Monitor	Full NM	Approximate NM Identify+Discover
0%	2%	35+7%	0%	20%	46%	34+1%
20%	18%	47+6%	20%	23%	51%	39+1%
40%	35%	62+4%	40%	31%	57%	45+1%
60%	50%	72+3%	60%	43%	69%	56+0%
80%	66%	80+2%	80%	58%	79%	70+0%
100%	100%	100+0%	100%	100%	100%	100+0%

$A \in \llbracket 1, 50 \rrbracket$ and $G \in \llbracket 1, 500 \rrbracket$ $A \in \llbracket 1, 50 \rrbracket$ and $G \in \llbracket 1, 5 \rrbracket$

TABLE III: Fulfilments Detected in the Action Observability Experiment

2) *Action Possibilities:* To analyse the performance and scalability of monitors with respect to agent capabilities to execute actions (i.e., the number of actions available), we defined the observation probability as a random value within the $\llbracket 0\%, 100\% \rrbracket$ interval and we varied the number of actions.

²³Note that when the NOP action is always observable, the capabilities of monitors to detect violations and fulfilments do not depend on the fact that agents are allowed to perform the NOP action in any situation.

Observ. Prob.	Traditional Monitor	Approximate NM Identify+Discover	Observ. Prob.	Traditional Monitor	Full NM	Approximate NM Identify+Discover
0%	2%	36+8%	0%	20%	44%	33+1%
20%	18%	50+6%	20%	23%	50%	37+0%
40%	35%	62+6%	40%	32%	56%	45+0%
60%	49%	70+2%	60%	42%	68%	55+0%
80%	65%	78+2%	80%	56%	78%	69+0%
100%	100%	100+0%	100%	100%	100%	100+0%

$A \in \llbracket 1, 50 \rrbracket$ and $G \in \llbracket 1, 500 \rrbracket$ $A \in \llbracket 1, 50 \rrbracket$ and $G \in \llbracket 1, 5 \rrbracket$

TABLE IV: Violations Detected in the Action Observability Experiment

Tables V and VI show the percentage of detected fulfilments and violations, respectively. In this experiment, the more actions, the more complex the reconstruction problem is. As a consequence, the improvement offered by an approximate NM over a traditional monitor decreases as the number of actions increases. However, the approximate NM still offers on average a 56% performance improvement over a traditional monitor. When the number of actions is very high (e.g., when the number of actions is 128 in small scenarios), then action preconditions become very complex and most of the times the NOP action is executed by all agents, which entails that the all monitors obtain a good performance. We can see that, as in the previous experiment, the approximate NM performs slightly worse than the full NM (i.e., the full NM offers on average a 15% performance improvement over an approximate NM). However, full and approximate NMs noticeably outperform the traditional monitor regardless of the number of actions.

Actions	Traditional Monitor	Approximate NM Identify+Discover	Actions	Traditional Monitor	Full NM	Approximate NM Identify+Discover
2	52%	95+8%	2	69%	99%	95+3%
8	49%	75+1%	8	47%	75%	71+0%
32	32%	48+0%	32	38%	63%	48+0%
128	29%	40+0%	128	53%	83%	57+0%

$G \in \llbracket 1, 500 \rrbracket$ $G \in \llbracket 1, 5 \rrbracket$

TABLE V: Fulfilments Detected in the Action Possibilities Experiment

Actions	Traditional Monitor	Approximate NM Identify+Discover	Actions	Traditional Monitor	Full NM	Approximate NM Identify+Discover
2	53%	93+5%	2	68%	99%	95+2%
8	49%	76+6%	8	48%	76%	72+1%
32	35%	50+1%	32	38%	62%	47+0%
128	31%	39+0%	128	56%	85%	59+0%

$G \in \llbracket 1, 500 \rrbracket$ $G \in \llbracket 1, 5 \rrbracket$

TABLE VI: Violations Detected in the Action Possibilities Experiment

C. Summary

The conclusions of our evaluation are threefold:

- 1) Both approximate and full reconstruction processes are more effective (i.e., detect more norm violations and fulfilments) than traditional monitoring approaches regardless of the scenario complexity (i.e., action possibilities and observability). Both in the case study and in the random experiments our algorithms improved significantly the percentage of violations and fulfilments detected.
- 2) Approximate reconstruction is slightly less effective than full reconstruction. In the case study, where a single

prohibition norm was monitored; the approximate NM obtained almost the same results as the full NM. In our random experiments, where several prohibition and obligation norms were monitored, the full NM offered on average an improvement of a 18% over an approximate NM.

- 3) Approximate reconstruction is scalable with the scenario size (i.e., the number of agents and actions to be monitored). In particular, our experiments demonstrate that the approximate algorithm can be used to monitor a large number of agents (we simulated scenarios with up to 500 agents), actions (we simulated scenarios with up to 128 actions), and norms (we simulated scenarios with up to 128 norms).

D. Applications

The information model and algorithms proposed in this paper will be of interest to a wide range of businesses and sectors within the cybernetic area. For example, monitoring is already used widely in the security sector from network monitoring, to malware detection software, and abnormal and intrusion detection tools. All these areas could benefit from having advanced algorithms for run-time norm monitoring. In particular, the existing solutions [2], [8], [11], [12], [16], [23], [26] for norm monitoring are ineffective for these application domains due to the inherent incompleteness of observations; e.g., intrusion detection and network monitoring tools only check a subset of the network information, cannot monitor all events in the system and can be partially evaded by hackers. Even those solutions that try to overcome the partial observability problem cannot be applied to these domains as the security tools should be able to operate in real-time with limited resources and the security policies cannot be modified, which entails that solutions proposing to increase the number of monitors [7] and solutions proposing the adaptation of norms to what is observable [3] are both unsuitable. Besides that, the algorithms we propose for reconstructing unobserved actions and monitoring norm compliance can be applied into solutions for controlling autonomous systems and robotics. Finally, our proposal could also be applied to other domains such as the finance industry to develop into anti-money laundering software capable of inferring illegal transactions.

VI. RELATED WORK

Recent works have also confronted with the problem of how to check norm compliance [10]. Specifically, this problem has been tackled from a logical and formal perspective; e.g., the proposals in [1], [13] describe logic formalisms and axioms for checking norm compliance. Besides that, there are works that have tackled this issue from a rule-based perspective; e.g., the proposals in [17], [5] describe rules and inference engines for checking norm compliance. Finally, there are proposals on the development of practical mechanisms and algorithms to check norm compliance. Given that our proposal falls into this last category, this section reviews the most relevant control mechanisms that allow norms to have an effective influence on agent behaviours [18]. These *control* mechanisms

are classified into two main categories [18]: *regimentation* mechanisms, which make the violation of norms impossible; and *enforcement* mechanisms, which are applied after the detection of norm violations and fulfilments, reacting upon them.

Regimentation mechanisms prevent agents from performing forbidden actions (vs. force agents to perform obligatory actions) by mediating access to resources and the communication channel, such as Electronic Institutions (EIs) [15]. However, the regimentation of all actions is often difficult or impossible. Furthermore, it is sometimes preferable to allow agents to make flexible decisions about norm compliance [9]. In response to this need, *enforcement* mechanisms were developed. Proposals on the enforcement of norms can be classified according to the entity that monitors whether norms are fulfilled or not. Specifically, norm compliance can be monitored by either agents themselves or the underlying infrastructure may provide monitoring entities.

Regarding *agent monitoring*, this approach is characterized by the fact that norm violations and fulfilments are monitored by agents that are involved in an interaction [31], [12], or other agents that observe an interaction in which they are not directly involved [30], [14], [33]. The main drawback of proposals based on agent monitoring is the fact that norm monitoring and enforcement must be implemented by agent programmers.

Regarding *infrastructural monitoring*, several authors proposed developing entities at the infrastructure level that are in charge of both monitoring and enforcing norms. Cardoso & Oliveira [8] proposed an architecture in which the monitoring and enforcement of norms is made by a single institutional entity. This centralized implementation represents a performance limitation when dealing with a considerable number of agents. To this limitation of centralized approaches, distributed mechanisms for an institutional enforcement of norms were proposed in [25], [16], [26], [23], [11].

All of the aforementioned proposals on monitoring assume that monitors have complete observational capabilities. Exception to these approaches is the recent work of Bulling et al. [7] and Alechina et al. [3]. In [7], the partial observability problem is addressed combining different norm monitors to build ideal monitors (i.e., monitors that together are able to detect the violation of a given set of norms). In [3], the authors propose to synthesise an approximate set of norms that can be monitored given the observational capabilities of a monitor. However, there are circumstances in which norms cannot be modified (e.g., contract and law monitoring) or ideal monitors are expensive and/or not feasible. We take a different approach in which norms and monitors' observation capabilities remain unchanged and monitors reconstruct unobserved actions.

Our approach is also related to planning, where methods (e.g., POMDPs [21]) for choosing optimal actions in partially observable environments have been proposed. A major difference between these proposals and our proposal is that NMs do not perform practical reasoning, i.e., they do not try to optimise or achieve a practical goal. Instead, NMs perform both deductive and abductive reasoning [27] to reason from observed actions to reach a conclusion about the state of the world, and to infer unobserved actions.

VII. CONCLUSION

In this paper, we propose information models and algorithms for monitoring norms under partial action observability by reconstructing unobserved actions from observed actions using two different reconstruction processes: full and approximate. Our experiments demonstrate that both reconstruction processes significantly detect more norm violations than traditional monitoring approaches. Approximate reconstruction performs slightly worse than full reconstruction, whereas its computational cost is much cheaper, making it suitable to be applied in practice.

The reconstruction algorithms proposed in this paper can be applied to several domains that require action monitoring; from normative MAS [32], to intrusion detection systems [19], to control systems [28] and to intelligent surveillance systems [20]. As future work, we plan to investigate domain-dependent approximations that could speed up action reconstruction.

REFERENCES

- [1] T. Ågotnes, W. Van der Hoek, and M. Wooldridge. Robust normative systems and a logic of norm compliance. *Logic Journal of IGPL*, page jzp070, 2009.
- [2] N. Alechina, M. Dastani, and B. Logan. Reasoning about normative update. In *Proc. of IJCAI*, 2013.
- [3] N. Alechina, M. Dastani, and B. Logan. Norm approximation for imperfect monitors. In *Proc. of AAMAS*, pages 117–124, 2014.
- [4] T. Bass. Intrusion detection systems and multisensor data fusion. *Communications of the ACM*, 43(4):99–105, 2000.
- [5] G. Boella and L. van der Torre. Norm governed multiagent systems: The delegation of control to autonomous agents. In *Intelligent Agent Technology, 2003. IAT 2003. IEEE/WIC International Conference on*, pages 329–335. IEEE, 2003.
- [6] C. Boutilier and R. I. Brafman. Partial-order planning with concurrent interacting actions. *Journal of Artificial Intelligence Research*, 14(1):105–136, 2001.
- [7] N. Bulling, M. Dastani, and M. Knobbout. Monitoring norm violations in multi-agent systems. In *Proc. of AAMAS*, pages 491–498, 2013.
- [8] H. Cardoso and E. Oliveira. Institutional reality and norms: Specifying and monitoring agent organizations. *International Journal of Cooperative Information Systems*, 16(1):67–95, 2007.
- [9] C. Castelfranchi. Formalising the informal? Dynamic social order, bottom-up social control, and spontaneous normative relations. *Journal of Applied Logic*, 1(1-2):47–92, 2003.
- [10] N. Criado, E. Argente, and V. Botti. Open issues for normative multi-agent systems. *AI Communications*, 24(3):233–264, 2011.
- [11] N. Criado, E. Argente, P. Noriega, and V. Botti. Manea: A distributed architecture for enforcing norms in open mas. *Engineering Applications of Artificial Intelligence*, 26(1):76–95, 2013.
- [12] A. Daskalopulu, T. Dimitrakos, and T. Maibaum. Evidence-based electronic contract performance monitoring. *Group Decision and Negotiation*, 11(6):469–485, 2002.
- [13] M. Dastani, D. Grossi, J.-J. C. Meyer, and N. Tinnemeier. Normative multi-agent programs and their logics. In *Knowledge Representation for Agents and Multi-Agent Systems*, pages 16–31. Springer Berlin Heidelberg, 2009.
- [14] A. P. de Pinninck, C. Sierra, and W. M. Schorlemmer. A multiagent network for peer norm enforcement. *Journal of Autonomous Agents and Multi-Agent Systems*, 21(3):397–424, 2010.
- [15] M. Esteva, B. Rosell, J. A. Rodríguez-Aguilar, and J. L. Arcos. Ameli: An agent-based middleware for electronic institutions. In *Proc. of AAMAS*, pages 236–243, 2004.
- [16] D. Gaertner, A. García-Camino, P. Noriega, J.-A. Rodríguez-Aguilar, and W. Vasconcelos. Distributed norm management in regulated multiagent systems. In *Proc. of AAMAS*, pages 624–631, 2007.
- [17] A. García-Camino, J. A. Rodríguez-Aguilar, C. Sierra, and W. Vasconcelos. A rule-based approach to norm-oriented programming of electronic institutions. *ACM SIGecom Exchanges*, 5(5):33–40, 2006.
- [18] D. Grossi, H. Aldewereld, and F. Dignum. Ubi lex, ibi poena: Designing norm enforcement in e-institutions. In *COIN II*, pages 101–114. Springer, 2007.
- [19] W. Hu, W. Hu, and S. Maybank. Adaboost-based algorithm for network intrusion detection. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 38(2):577–583, 2008.
- [20] C.-M. Huang and L.-C. Fu. Multitarget visual tracking based effective surveillance with cooperation of multiple active cameras. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 41(1):234–247, 2011.
- [21] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Journal of Artificial Intelligence*, 101(1):99–134, 1998.
- [22] F. López y López, M. Luck, and M. d’Inverno. A normative framework for agent-based systems. *Computational & Mathematical Organization Theory*, 12(2-3):227–250, 2006.
- [23] F. Meneguzzi, S. Modgil, N. Oren, S. Miles, M. Luck, and N. Faci. Applying electronic contracting to the aerospace aftercare domain. *Engineering Applications of Artificial Intelligence*, 25(7):1471–1487, 2012.
- [24] R. Micalizio, P. Torasso, and G. Torta. On-line monitoring and diagnosis of multi-agent systems: a model based approach. In *Proc. of ECAI*, volume 16, page 848, 2004.
- [25] N. Minsky and V. Ungureanu. Law-governed interaction: a coordination and control mechanism for heterogeneous distributed systems. *ACM Transactions on Software Engineering and Methodology*, 9(3):273–305, 2000.
- [26] S. Modgil, N. Faci, F. Meneguzzi, N. Oren, S. Miles, and M. Luck. A framework for monitoring agent-based normative systems. In *Proc. of AAMAS*, pages 153–160, 2009.
- [27] G. Paul. Approaches to abductive reasoning: an overview. *Journal of Artificial Intelligence Review*, 7(2):109–152, 1993.
- [28] C. Rieger, D. Scheidt, and W. Smart. Guest editorial: Introduction to the special issue on resilient control architectures and systems. *IEEE Transactions on Cybernetics*, 44(11):1994–1996, Nov 2014.
- [29] S. Sadiq, G. Governatori, and K. Namiri. Modeling control objectives for business process compliance. In *Business process management*, pages 149–164. Springer, 2007.
- [30] S. Sen and S. Airiau. Emergence of norms through social learning. In *Proc. of IJCAI*, pages 1507–1512, 2007.
- [31] M. Venkatraman and M. Singh. Verifying compliance with commitment protocols. *Journal of Autonomous Agents and Multi-Agent Systems*, 2(3):217–236, 1999.
- [32] C. Yu, M. Zhang, and F. Ren. Collective learning for the emergence of social norms in networked multiagent systems. *IEEE Transactions on Cybernetics*, 44(12):2342–2355, Dec 2014.
- [33] W. Zeng and M.-Y. Chow. Resilient distributed control in the presence of misbehaving agents in networked control systems. *IEEE Transactions on Cybernetics*, 44(11):2038–2049, Nov 2014.

Natalia Criado is Senior Lecturer in the School of Computing and Maths at Liverpool John Moores University (UK), which she joined in 2014. Her main research interests are on the intersection between Artificial Intelligence and Normative Systems, and in particular, intelligent/automated approaches to norm reasoning, with a strong focus on Multi-agent Systems and BDI Agents.



Jose M. Such is Lecturer (Assistant Professor) in the School of Computing and Communications at Lancaster University (UK), which he joined in 2012. His main research interests are on the intersection between Artificial Intelligence, Cyber Security and Human Factors.

