



**HAL**  
open science

## An Agent-Centric Perspective on Norm Enforcement and Sanctions

Elena Yan, Luis Gustavo Nardin, Jomi Hübner, Olivier Boissier

► **To cite this version:**

Elena Yan, Luis Gustavo Nardin, Jomi Hübner, Olivier Boissier. An Agent-Centric Perspective on Norm Enforcement and Sanctions. International Workshop on Coordination, Organizations, Institutions, Norms and Ethics for Governance of Multi-Agent Systems, May 2024, Auckland, New Zealand. 10.48550/arXiv.2403.15128 . hal-04597649

**HAL Id: hal-04597649**

**<https://hal.science/hal-04597649v1>**

Submitted on 3 Jun 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# An Agent-Centric Perspective on Norm Enforcement and Sanctions

Elena Yan<sup>1</sup>[0009-0000-6660-9378], Luis G. Nardin<sup>1</sup>[0000-0002-4506-2745], Jomi F. Hübner<sup>2</sup>[0000-0001-9355-822X], and Olivier Boissier<sup>1</sup>[0000-0002-2956-0533]

<sup>1</sup> Mines Saint-Etienne, Univ Clermont Auvergne, INP Clermont Auvergne, CNRS, UMR 6158 LIMOS, F-42023 Saint-Etienne France

[elena.yan@emse.fr](mailto:elena.yan@emse.fr), [luisgustavo.nardin@emse.fr](mailto:luisgustavo.nardin@emse.fr), [olivier.boissier@emse.fr](mailto:olivier.boissier@emse.fr)

<sup>2</sup> Department of Automation and Systems Engineering, Federal University of Santa Catarina, Florianópolis, Brazil [jomi.hubner@ufsc.br](mailto:jomi.hubner@ufsc.br)

**Abstract.** In increasingly autonomous and highly distributed multi-agent systems, centralized coordination becomes impractical and raises the need for governance and enforcement mechanisms from an agent-centric perspective. In our conceptual view, sanctioning norm enforcement is part of this agent-centric approach and they aim at promoting norm compliance while preserving agents' autonomy. The few works dealing with sanctioning norm enforcement and sanctions from the agent-centric perspective present limitations regarding the representation of sanctions and the comprehensiveness of their norm enforcement process. To address these drawbacks, we propose the NPL(s), an extension of the NPL normative programming language enriched with the representation of norms and sanctions as first-class abstractions. We also propose a BDI normative agent architecture embedding an engine for processing the NPL(s) language and a set of capabilities for approaching more comprehensively the sanctioning norm enforcement process. We apply our contributions in a case study for improving the robustness of agents' decision-making in a production automation system.

**Keywords:** Sanctions · Normative Programming Language · Normative Agent Architecture · Norm Enforcement.

## 1 Introduction

The concepts used for conceiving a multi-agent system (MAS) are often inspired by human societies, e.g., abstractions of agent, environment, interaction, and organization [3]. Agents are autonomous entities that make independent decisions and act to achieve their goals [30]. In Normative MAS (NMAS), the agents' actions are regulated by norms. A set of norms comprises an explicit and formal specification of the behaviors that agents ought to comply with for the system to achieve its overall objectives. Agent's autonomy, however, allows agents to behave nonconforming to the prescribed norms [7], which may lead the system into undesirable states.

Enforcement mechanisms are introduced to nudge agents to act conforming to the norms aiming to steer the system towards its overall objectives [13]. The institutional, or organization-centric, norm enforcement approach relies on a central authority to control and enforce the norms. This approach has two drawbacks. First, this approach becomes impractical in highly distributed large-scale systems. Second, it usually employs regimented enforcement mechanisms that prevent agents from violating the norms, limiting the agents' autonomy. A solution to overcome these drawbacks and balance agents' autonomy and system regulation is to employ a social, or agent-centric, norm enforcement approach, in which agents supervise and regulate the behavior of others.

In our conceptual view, sanctioning is part of this agent-centric norm enforcement approach. Sanctioning is a common enforcement mechanism that uses sanctions to encourage agents to comply with norms. Sanctions are reactions to a behavior, yet not all reactions can count as sanctions [12].

The NMAS community has focused on incorporating normative capabilities into agent architectures [23]. In these architectures, agents should represent the norms and be able to (i) recognize new norms, (ii) recognize actions affected by the norms, (iii) reason about the norms, and (iv) decide whether or not to comply with them. Assuming that a sanctioning enforcement mechanism is adopted, agents should also represent sanctions and be able to (i) identify norm compliance and violation, (ii) decide for the most appropriate sanctions depending on the situation, and (iii) apply the decided sanctions.

Despite the representation of norms and agents' capabilities for handling them have been well investigated in the context of NMAS, few works have proposed a comprehensive perspective on sanctioning norm enforcement from an agent-centric perspective. To reduce this disparity, we advance in norm enforcement by (i) enriching a normative programming language with the ability to represent norms and sanctions as first-class abstractions, and (ii) embedding sanctioning norm enforcement capabilities into a BDI agent architecture to enable agents to deliberate on their normative state when deciding to act.

The remainder of the paper is organized as follows. In Section 2, we introduce the concepts of norms and sanctions, and discuss normative agent architectures and normative programming languages from the perspective of their support to norm enforcement. We then extend a normative programming language by adding sanction as a first-class abstraction (Section 3) and extend a BDI agent architecture with sanctioning capabilities (Section 4). Next, we illustrate the use of our contributions in an industrial case study (Section 5). Finally, we conclude with some discussions and potential future research directions (Section 6).

## 2 Related Work

Norm is an important concept in human societies that has been investigated in a wide range of disciplines. Norms prescribe how members of a group ought to behave in a given situation [29]. They are expected standards of behavior socially shared and enforced by group members [13]. In MAS, norms are typically

represented as obligations, permissions, or prohibitions; they are used to enable a degree of control over the autonomous agents' actions.

The agents' decision-making processes depend and intervene through their decisions, on the state of the norms (e.g., recognized, adopted, or complied). The transitions between these norm states define what we call the norm's life cycle.

## 2.1 Norm's life cycle

**Global overview.** There are various proposals regarding the life cycle of norms in the literature. They often share a similar structure and process but differ in specific details. Hollander and Wu [13] present the norm life cycle as an overall end-to-end evolutionary process. The norm's life cycle begins with the creation of norms by system designers at design time or by autonomous normative agents at runtime (*creation* phase). In the latter, norms emerge as a consequence of the agents' interactions. Once the norms exist, they are spread in the society from one agent to another through active or passive transmission (*transmission* phase). Next, enforcement mechanisms (e.g., sanctioning) or norm compliers are used to encourage agents to adopt new norms and to comply with the adopted ones (*enforcement* phase). The norms adoption leads to their use in the agents' decision-making (*internalization* phase). As conditions change, norms may become invalid or lose their effectiveness, which cause them to be forgotten (*forgotten* phase).

Next, we describe the *enforcement* phase in more detail since this paper focuses on norm enforcement.

**Norm Enforcement.** There are three traditional approaches to norm enforcement: (i) the *institutional approach*, in which there is a single authority that enforces agents' behaviors; (ii) the *social approach*, in which a subset of agents are equipped with enforcement mechanisms (e.g., police agents) used to enforce other agents' behaviors; and (iii) the *self-enforcement approach*, in which agents enforce the norms on themselves.

Cross-cutting these three approaches, the norm enforcement can be *regimented* or *regulated*. In the regimented norm enforcement, agents are prevented from violating the norms, possibly limiting agents' autonomy. In the regulated norm enforcement, agents can act in violation to the norms, but mechanisms, such as sanctioning, can be used to inflict consequences to their misbehavior in order to encourage their future compliance with the norms.

In the regulated norm enforcement, sanctions can be used as a means to steer agent compliance to the norms. Sanction is a negative or a positive reaction to a violation of or a compliance with a norm. A comprehensive set of dimensions to classify sanctions and an accompanying sanctioning norm enforcement process model has been proposed in [26]. The proposed typology of sanctions enables identifying the types of sanctions that are more or less effective in reducing violation or increasing compliance. The sanctioning process model combines the institutional, the social, and the self-enforcement approaches offering opportunities for agents to reason about and decide for the most appropriate sanctions

to apply in each situation. This model has two normative resources: the *De Jure* that stores the representation of norms and sanctions and the *De Facto* that stores the historical information about sanction decisions, executions, and outcomes. The model is enacted by five capabilities: (i) the *Detector* perceives the environment and detects any action regulated by the norms in the *De Jure*; (ii) the *Evaluator* obtains the norms and sanctions regulating the detected action from the *De Jure* to determine whether the action violates or complies with the norms. The *Evaluator* also combines this information with the historical information from the *De Facto* to decide the sanctions that should be applied, if any; (iii) the *Executor* applies the sanctions decided by the *Evaluator*; (iv) the *Controller* watches the regulated actions previously sanctioned to assess the efficacy of the sanctions; (v) the *Legislator* assesses the *De Jure* in the light of the *De Facto* information and updates the norms and sanctions in the *De Jure*.

Next, we review normative agent architectures (Section 2.2) and normative programming languages (Section 2.3) from the perspective of their support to norm enforcement.

## 2.2 Normative Agent Architectures

Focusing on the BDI (Belief-Desire-Intention) agent model, the literature proposes normative architectures using various approaches to enforce norms.

Most architectures regiment norms by preventing agents from performing actions that violate them. The BOID (Beliefs-Obligations-Intentions-Desires) [6] architecture adds the concept of *obligations* as a mental state that may conflict with the agent’s intentions or desires. BOID agents use a predefined static priority function to resolve conflicts among these concepts. The priority ordering allows classifying agents in different categories, e.g., selfish (i.e., prioritize desires) or socially responsible (i.e., prioritize obligations). Although BOID does not explicitly implement a norm enforcement mechanism, social responsible agents regiment norms since they always choose actions that fulfill their obligations.

Other architectures, like NoA [18] and NBDI [28], regiment norms by filtering out actions or desires that violate norms. In the NoA architecture, agents use filtering functions to simply exclude actions that violate adopted norms. In the NBDI architecture, agents update their beliefs, desires, intentions and set of adopted norms independently. When selecting desires, NBDI agents only select those that are compliant with the adopted norms. Using yet a different approach, the Normative AgentSpeak(L) [24] architecture enables agents to change their plans at runtime to conform to new norms. The approach consists in modifying the agent’s plans library, i.e., create new plans when new obligation norms are adopted and suppress plans containing actions forbidden by prohibition norms. The N-Jason [21] extends the Normative AgentSpeak(L) capabilities by scheduling intentions considering deadline and priority for runtime norm compliance.

A NoA extension [19] and the Nu-BDI [25] architectures relax the regimentation constraint allowing the agents to violate the norm, but diminishing the negative impact that the violation may have. In the extended NoA architecture, agents have deliberation mechanisms to annotate actions that violate norms;

these annotations can be later used by the agents to decide whether or not to perform actions leading to the violation of norms. In the Nu-BDI architecture, agents annotate each step of the plans within the scope of a norm with constraints stemming from that norm, and rank the plans using a utility function based on these annotations. When selecting plans, Nu-BDI agents choose the plan that violates the norms less (i.e., highest ranked plan).

López y López et al. [22] implicitly incorporate regulated norm enforcement mechanisms into their agent framework. They define two levels of norms, in which the secondary norms are activated as a result of the fulfillment or violation of the primary norms. So, the secondary norms can be seen as sanctions inflicted on agents fulfilling or violating norms.

Few normative agent architectures deal explicitly with regulated norm enforcement. EMIL-A (EMergence in the Loop Architecture) [2] is a normative BDI agent architecture developed for norm emergence, innovation, and spread. EMIL-A is endowed with modules that allow (i) norm recognition (i.e., recognition and representation of norms as mental states), (ii) norm adoption (i.e., detection and dynamic updating of the norms salience that corresponds to the perceived prominence of a norm within a relevant reference group), (iii) norm compliance (i.e., decision to comply with the adopted norms), and (iv) norm enforcement (i.e., detection of norm compliance and violation, and the application of sanctions). The n-BDI architecture [8] also incorporates a norm enforcement mechanism that always rewards or punishes the agents as a consequence to, respectively, the compliance with or violation of a norm. AORTA [16] proposes a normative agent module that provides normative reasoning capabilities to the agents. These agents are able to reason about norms, automatically detect violations as soon as they become aware of the violation state, and produce a new belief about the violation that may trigger another norm or plan.

In addition to those BDI-based normative agent architectures, Fagundes et al. [10] extend the Markov Decision Process (MDP) framework for creating self-interested agents able to reason about norms. The framework uses an enforcement mechanism based on the imperfect observation of the system state and agents, and inflicts a cost on agents if a violation is detected [11]. Agents make decisions based on the expected utility of conforming with or violating a norm, which takes into account the cost of violating the norm and the probability that the norm violation will be detected.

### 2.3 Normative Programming Languages

Normative programming languages provide abstractions that enable agents to deliberate about their actions bounded by the normative regulations. These languages complement the normative agent architectures as they provide an explicit representation of and enable the handling of normative concepts (e.g., obligation, permission, prohibition, sanction). There has been extensive work on normative programming languages with an organization-centric perspective [9,15]; however, there has been limited interest in languages with an agent-centric perspective.

One of the first normative programming languages with an agent-centric perspective is the NoA language [17]. The language, whose semantics are implemented by the NoA architecture, contains constructs for the specification of beliefs, goals, plans and norms. Norm specifications can regulate states (i.e., the achievement of a particular state of the world) or actions (i.e., the performance of explicit actions). In this language, normative statements are expressed as obligations, permissions, prohibitions, and sanctions. Although sanction exists as an independent construct, it is a syntactic sugar that sets an obligation for an agent to pursue certain activities that represent such sanctions.

The Normative Programming Language (NPL) [14], albeit used mainly within the organization-centric perspective [15], is a general language dedicated to the development of normative programs based on two primitive constructs: obligation and regimentation. Other constructs like prohibition, permission, and sanction are represented using these primitive constructs. Prohibitions are accomplished by regimentation or by delegating to someone an obligation to deal with the situation. Permissions are defined by omission, i.e., NPL adopts the open-world assumption. Likewise the NoA language, the NPL also represents sanctions as obligations.

The N-2APL [1] is an extension to the agent programming language 2APL [9]. The N-2APL provides support for representing beliefs, goals, plans, norms (i.e., obligations or prohibitions), sanctions, deadlines, and durations. A sanction is a consequence to the violation of an obligation or prohibition, i.e., specify the updates to be applied on the environment due to the violation of a norm. In N-2APL, sanction is part of the norm specification, thus tightly associated to it.

## 2.4 Remarks

In our conceptual view, sanctions are part of social mechanisms (i.e., agent-centric perspective) for norm enforcement as they are consequences of norms violation or compliance. Although sanctioning mechanisms are built on top of norms and their state, sanctions are not simply triggered depending on norms state, but also depending on the context in which the compliance or violation take place. Tackling these requirements call for a normative programming language that disentangles sanctions representation from, yet letting it associated to, norms, and a normative agent architecture that triggers sanctions from a normative state contingent upon the agent’s contextual situation in order to fulfill the agent-centric perspective.

In the next sections, we propose an extension to the NPL [14] normative language and to the JaCaMo agent [4] to tackle these requirements. We chose to extend NPL because, to the best of our knowledge, it is the only general purpose normative programming language available, and the JaCaMo agent because NPL has already been integrated into it in an organization-centric perspective [15].

### 3 NPL(s): Extending NPL with Sanctions

Here we extend the NPL [14] language to introduce the concept of sanctions as a first-class abstraction. Our approach consists of having an enriched and separated description in formal language of sanctions and the conditions that allow them to be triggered. This language is interpreted at run-time based on the current state of the normative system. In addition, the NPL interpreter should be extended to compute new sanctions. The new language, named NPL(s), is thus presented in the sequel by its syntax and (informal) semantics.

A normative program  $np$  in NPL(s) is composed of a set of: (i) facts and inference rules (based on a syntax similar to the one used in Jason [4]), (ii) norms, (iii) sanction rules. While the two first components come from NPL, the last one is the extension introduced by NPL(s). We briefly introduce NPL norms in this paper, but the formal semantics and more details are available at [14]. An NPL norm has the general form:

$$\text{norm } id : \varphi \rightarrow \psi .$$

where  $id$  is a unique *identifier* of the norm;  $\varphi$  is a formula that determines the *activation condition* for the norm; and  $\psi$  is the *consequence* of the activation of the norm. Among the possible consequences, here we focus on obligations (other possible consequences are permissions, prohibitions and failures).

$\psi = \text{obligation}(a, m, g, d)$ : represents the case where an obligation for some agent  $a$  is created. Argument  $m$  is the maintenance condition for the obligation;  $g$  is the formula that represents the obligation itself (a state of the world that the agent must try to bring about, i.e., a goal it has to achieve); and  $d$  is the deadline condition to fulfill the obligation.

The informal semantics defines that whenever  $\varphi$  is true, an obligation for  $a$  is created and its initial state is **active**. If  $m$  becomes false, the obligation state changes to **inactive**. If  $g$  becomes true before  $d$  is true, the obligation state changes to **fulfilled**. If  $d$  becomes true before  $g$  is true, the obligation state changes to **unfulfilled**. The final state of an obligation may thus be **unfulfilled**, **fulfilled**, or **inactive**.

In NPL(s), a norm can be followed by sanction rules to be evaluated according to the final state of an obligation (see *tsr* in Table 1):

$$\begin{aligned} \text{norm } id : \varphi \rightarrow \psi \\ \text{if fulfilled: } sr_1(args), sr_2(args), \dots, sr_n . \end{aligned}$$

A sanction rule has the following general form (see *srule* in Table 1):

$$\text{sanction-rule } sr_i(args) : \rho \rightarrow \text{sanction}(a, s) .$$

where  $sr_i$  is a unique *identifier* of the sanction rule;  $args$  are terms passed as parameters (optional);  $\rho$  is a formula that determines the *activation condition* for the sanction (optional, assumed true if not specified);  $a$  is the target agent of



---

```

np      ::= "np" atom "{" ( rule | norm | srule ) * "}"
rule   ::= atom [ ":" "-" formula ] "."
norm   ::= "norm" id ":" formula "->" ( fail | obl | per | pro ) (tsr)* "."

fail   ::= "fail(" atom ")"
obl    ::= "obligation" dargs
per    ::= "permission" dargs
pro    ::= "prohibition" dargs
dargs  ::= "(" (var | id) "," formula "," formula "," (time | formula) ")"

tsr    ::= "if" ("fulfilled" | "unfulfilled" | "inactive") ":" ( atom ) *

srule  ::= "sanction-rule" atom ":" formula "->" "sanction(" (var | id) "," atom ")"

formula ::= atom | "not" formula | atom ( "&" | "|" ) formula
time    ::= "'" number ( "second" | "minute" | ... ) "'"

```

---

**Table 1.** EBNF of the Normative Programming Language for Sanctions — NPL(s). Non-terminals *atom*, *id*, *var*, and *number* correspond, respectively, to predicates, identifiers, variables, and numbers as used in Prolog.

the sanction; and *s* is the sanction content. Sanction rules can be read as “In the case that a norm has triggered the sanction rule  $sr_i$  because it is (un)fulfilled or inactive, if  $\rho$  is true, then the sanction *s* is created for agent *a*.” A simple example to illustrate the language is given below. Two norms are defined, each associated with a sanction-rule. Norm *n1* specifies that *alice* is obligated to achieve *b(0)* within 3 seconds. Unfulfillment to do so activates sanction-rule *sr1*, resulting in the imposition of a *fine* to *alice*. If the sanction is created and the condition *extra(C)* is met, norm *n2* is activated. This delegates *bob* to apply the *fine* to *alice* in 2 seconds. In the event of non-compliance, a sanction is issued to remove *bob* from the system.

```

// * NORMS *
// alice has 3 seconds to achieve b(0), or else evaluate sr1
norm n1: vl(X) & X > 5
    -> obligation(alice,true, b(0), '3 seconds')
    if unfulfilled: sr1(alice,X) .

// bob is obliged to apply fines in 2 seconds
norm n2: sanction(A,fine(X)) & extra(C)
    -> obligation(bob,true, apply_fine(A,X*C), '2 seconds')
    if unfulfilled: sr2 .

// * SANCTION RULES *
// if A is not in an emergency, create fine sanction for it
sanction-rule sr1(A,V) : not emergency(A) -> sanction(A,fine(V)) .

sanction-rule sr2 -> sanction(bob,remove_from_systems) .

```

Based on this NPL(s) example, we can have the following example story line:

1. A fact `v(20)` is produced and norm `n1` is triggered.
2. `alice` is thus obliged to achieve `b(0)` in 3 seconds.
3. After 3 seconds, `alice` does not fulfill the obligation, triggering the sanction rule `sr1(alice,20)`.
4. Since `alice` has no fact like `emergency(alice)`, the sanction rule `sr1` produces the `sanction(alice,fine(20))` fact.
5. This new sanction fact together with the fact `extra(10)` triggers the norm `n2` that obliges agent `bob` to `apply_fine(alice,200)` in 2 seconds.
6. Supposing that `bob` fulfills its obligation, the story ends here.
7. However, if `bob` does not fulfill the obligation, the sanction rule `sr2` is triggered and the `sanction(bob,remove_from_systems)` fact is produced. Here, we assume the environment is capable of handling and applying the sanction.

From this story, we can notice that the NPL(s) interpreter is limited to compute sanction facts. The impact of these sanctions on the running system depends on something else that reads the sanction fact and implements it. This could be integrated in an *agent-*, *environment-*, or *organization-*centric perspective. For instance, it could be an agent that reads the sanction fact and *decides* to run a procedure to implement it. However, it could be the case that the environment is capable of reading sanction facts and implementing them, as it is the case of automatic fines created when we cross red traffic lights. Finally, the NPL(s) interpreter itself can read sanction facts, as in the example of the activation of norm `n2`. In this paper we focus on the agent-centric perspective.

## 4 Normative Agent Architecture

Here we describe an extension of the JaCaMo BDI agent architecture to handle normative states as illustrated in Fig. 1<sup>3</sup>.

This normative architecture integrates an NPL(s) Engine enabling agents to interpret and compute the current state of norms and sanctions specified in the NPL(s) language (Section 2.3). Additionally, the normative architecture draws inspiration from the sanctioning norm enforcement process model proposed in [26] (Section 2.1) to define the specialized resources and enforcement capabilities needed to implement a normative agent. Although these capabilities are present in the architecture, they are implemented as plans that are triggered by normative facts. This characteristic gives modularity and flexibility to agents implementing this architecture as (i) not all capabilities have to be implemented in each agent and (ii) different agents can implement these capabilities differently, i.e., agents can deal differently with norms and sanctions. An exception is the Detector capability, which, in practice, is performed by the NPL(s) Engine.

<sup>3</sup> The source code of the normative agent architecture is available at <https://github.com/moise-lang/npl/blob/master/src/main/java/npl/NormativeAg.java>

<sup>4</sup> Note that this architecture depicts only norm enforcement capabilities. Others were omitted due to our interest in norm enforcement in this paper.

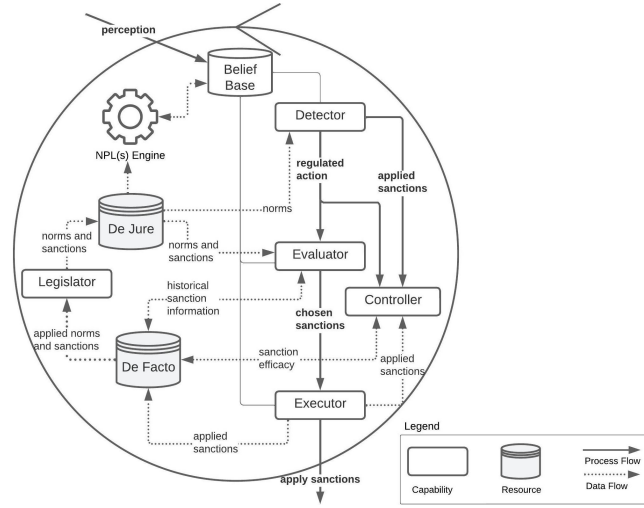


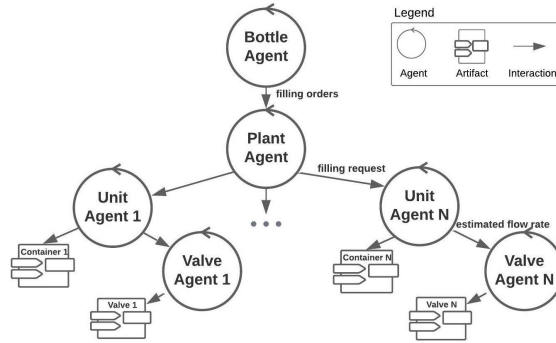
Fig. 1. Normative BDI agent architecture extended to handle normative states.<sup>4</sup>

The De Jure is a normative program, i.e., norms, sanction rules, and sanctions, specified in NPL(s) language. The specification is read once and used by the NPL Engine to together with the belief base information of the agent determine the current state of norms and sanctions, as well as add normative facts in the belief base. Once added into the belief base, these normative facts may trigger specific capabilities (i.e., plans). For instance, once a norm is activated, an obligation is added to the belief base causing the execution of a plan; or once norm is unfulfilled, a sanction rule is triggered producing sanction facts in the agent’s belief base. Contrary to the De Jure, the De Facto is part of the belief base of the agent that contains particular types of beliefs, i.e., beliefs related to the decision and application of sanctions and their efficacy.

## 5 Case Study: Laboratory Plant MyJoghurt

Here we present a case study illustrating the practical potential of the extended normative programming language NPL(s) integrated into the proposed normative agent architecture. The use case is focused on improving the robustness of agents’ decision-making in a production automation system. We took as reference the industrial use case of the laboratory plant *myJoghurt* introduced in [20].

The *myJoghurt* is composed of filling stations responsible for filling recipe-specific liquids (e.g., yogurt, milk) into bottles that are transported by the logistics system. Orders are received and the filling tasks are distributed among the process plant. The orders and filling process are controlled by a multi-agent system (illustrated in Fig. 2), whose agents control the process plant into filling the correct amount of liquid in the bottles according to the order received.



**Fig. 2.** *myJoghurt* multi-agent architecture (based on [20]).

The *myJoghurt* control system is composed of five types of components: bottle, plant, unit, container, and valve. They are controlled by four types of agents that interact with each other and two types of artifacts. The *bottle agent* receives and sends filling orders from customers with the type and the amount of liquid to the *plant agent*. The *plant agent* manages the process at the cell level and negotiates with the *unit agent* when orders are received. The *plant agent* then selects and assigns the order to *unit agents* based on their availability. The *unit agent*, which controls the filling station, interacts with the filling station's container and valve artifacts for the filling process. The *container*, represented as an artifact, manages the level of liquids in the container. The *valve agent* estimates the flow rate considering the type of liquid (e.g., a liquid with a higher viscosity leads to a faster flow) and whether the valve is clogged (e.g., residues of liquid could clog the valve over time). Based on the estimated value, the *valve agent* controls the opening window of the valve through the *valve* artifact.

Enforcement mechanisms and sanctions are used to steer the decisions of agents to improve the accuracy of the flow rate estimation in various types of liquid and in the presence of non-deterministic situations. The norms defined in this use case refer to ensuring the correct fulfillment of order requests of each bottle (i.e., bottles filled with the correct liquid and volume within a specified tolerance range). The system considers the deviation and learning factors to determine whether and which sanctions to apply. The deviation factor comprises polarity, determined by whether an action is complied with (positive polarity) or violated (negative polarity), and magnitude that measures the degree of deviation from the expected fill level. The learning factor considers the historical behavior of the *valve agents*, represented by the image of the agent (i.e., assessment based on direct experience with the agent [27]), the frequency of violations, and the efficacy of sanctions applied.

Table 2 illustrates the sanctions defined for the case study. The activation conditions of sanctions are evaluated based on the compliance or violation of the norm, and sanctions are applied accordingly. All sanctions have a negative polarity [26], intended to punish the *valve* or the *unit agent* for their misbe-

havior. S1 aims to guide the *valve agent* by adjusting its estimated flow rate, while all other sanctions aim to reduce (e.g., S3) or prevent future misbehavior temporarily (e.g., S2 and S5) or permanently (e.g., S4).

Id	Sanctioner	Target	Sanction	Pre-condition
S1	unit	valve	Adjust the estimated flow rate	The image is below a threshold
S2	unit	valve	Activate the self-cleaning procedure	The violation occurs three consecutive times
S3	plant	unit	Adjust the unit agent's image	The image is below a threshold
S4	plant	unit	Disregard the unit agent as an option for subsequent filling orders	The violation occurs five consecutive times
S5	plant	unit	Require manual intervention to repair the filling station	The violation occurs five consecutive times

**Table 2.** Sanctions defined for the case study. The sanctioner agent applies the sanction to the target agent if the activation pre-condition holds.

We used JaCaMo [3] with the extended NPL(s) normative language to implement the case study<sup>5</sup>.

### 5.1 Representation of Norms and Sanctions

Here, we present how norms and sanctions are represented in the De Jure repository using the NPL(s) normative language (see Section 3). The listing below illustrates the *Unit De Jure* repository with the representation of the *unit agent's* norms, sanction rules, and sanctions<sup>6</sup>.

```
norm n1: fill_bottle(LQ, X, MN, MX) & .my_name(U)
  -> obligation(U, n1, fill(LQ,X,MN,MX), level(X,L) & (L<MN | L>MX)).

norm n2: level(V,X,L) & .my_name(U)
  -> obligation(U, n2, update_factors(V,X,L),
    deviation_factor(V,X,"negative",_) & learning_factor(V,X,I,_,_,_) &
    threshold(T,_) & I<T)
  if unfulfilled: s1(V,X), s2(V,X).

sanction-rule s1(V,X): learning_factor(V,X,_,_,_,C) & threshold(_,T) & C<T
  -> sanction(V, adjust_flow_rate(X)).

sanction-rule s2(V,X): learning_factor(V,X,_,_,_,C) & threshold(_,T) & C>=T
  -> sanction(V, self_cleaning(X)).
```

The working principle of the multi agent-centric sanctioning mechanism including the NPL(s) engine operates as follows. The NPL(s) activates the norm *n1* whenever the *unit agent* receives a filling request `fill_bottle(LQ, X, MN,`

<sup>5</sup> The source code available at <https://github.com/yan-elena/myjoghurt-case-study>

<sup>6</sup> The *plant agent's* norms, sanction rules, and sanctions are available at [https://github.com/yan-elena/myjoghurt-case-study/blob/main/src/npl/plant\\_de.jure.npl](https://github.com/yan-elena/myjoghurt-case-study/blob/main/src/npl/plant_de.jure.npl)

MX) from the *plant agent*. This norm produces an obligation  $\text{fill}(\text{LQ}, \text{X}, \text{MN}, \text{MX})$  for the *unit agent* to fill the current bottle X with the liquid type LQ within the filling level range MN-MX. Once the bottle is filled, the norm n2 is activated and produces an obligation for the *unit agent* updating the deviation and learning factors based on the bottle filled. If the polarity is negative and the image drops below a threshold, the norm becomes *unfulfilled* and one of two sanction facts is produced, i.e., s1 and s2. If the number of consecutive violations (C) of the *valve agent* (V) is less than a threshold T, then the sanction fact s1 is produced ordering the *valve agent* (V) to adjust its flow rate; otherwise, the sanction fact s2 is produced ordering the *valve agent* (V) to activate the self cleaning procedure in the *valve* artifact.

Next the implementation of this sanctioning enforcement mechanism is shown in the context of the *myJoghurt* MAS architecture.

### 5.2 myJoghurt MAS architecture

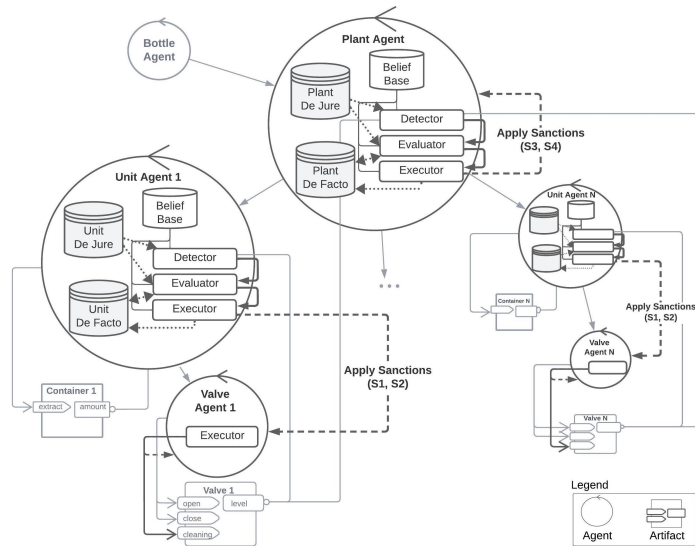


Fig. 3. Overview of integration of enforcement mechanisms into the MAS Architecture.

The normative agent architecture equipped with the NPL(s) engine proposed in Section 4 is used to design agents in the *myJoghurt* MAS architecture (Fig. 3). The *plant agent* and *unit agents* set the norms and sanctions in their De Jure repository and they are endowed with the enforcement. Because the applied agent architecture is modular, we can adapt each agent to specific requirements of the *myJoghurt* architecture. So, the *plant agent* and the *unit agents* implement the capabilities of Detector, Evaluator, and Executor, while the *valve agent* implements only the capability of Executor.

The Detector capability of the *unit agent* and *plant agent* perceives the filling level information of each completed bottle and detects any violation or compliance with the norms.

Once a violation is detected, both the *plant agent* and *unit agents* calculate and update the factors to determine whether and which sanctions to apply using their Evaluator capability. The Evaluator considers the norms in the De Jure and the historical information in the De Facto to make these decisions. Based on the sanctions identified in Table 2, the *unit agent* may apply sanctions S1 and S2 to the *valve agent* responsible for the violation. S2 and S3 are applied by the *plant agent* to itself in order to indirectly update the *unit agents'* image or even remove it from the considerations of future orders. S5 is not represented in the diagram since it is a sanction to be sent to an external actor in the system (e.g., a human operator) responsible for repairing the filling station.

The Executor capabilities of *plant agent*, *unit agents*, and *valve agent* handle the actual enforcement of the sanctions. We assume Executors are obedient and they enforce the sanctions chosen by the Evaluator. However, there may be cases in which the Executor may decide whether or not to enforce the sanctions. The actual sanction enforcement is specified in the agent code as a plan. The following listing shows the *unit agent* internal plan for reacting to the sanction fact of sanction S1 linked to adjusting the *valve agent* flow rate estimation.

```
+sanction(V, adjust_flow_rate(X))
  <- //...
    .send(V, signal, update_estimation(M)).
```

For the sake of simplicity in this paper, we have not considered the Controller and Legislator capabilities.

## 6 Conclusion and Future Work

In this paper, we have presented an agent-centric perspective on norm enforcement and sanctions by proposing a twofold contribution. We have extended the normative programming language NPL to NPL(s) with an explicit representation of sanctions as first-class abstraction in the definition of a normative program. Thanks to it, normative programs gain flexibility and expressivity enabling the association of sanctions with different norms on one hand, and, on the second hand, open the possibility for agents to locally decide on sanctions and how to enforce them.

Based on this enriched normative programming language, we proposed a normative agent architecture able to reason explicitly on norms and sanctions. The use of these new features has been illustrated in a multi-agent based control production system in the context of the industrial case study of the laboratory plant *myJoghurt*. We have shown how this normative agent-centric architecture opens a comprehensive, modular, and flexible perspective on norm enforcement and sanctions in a multi-agent system. Future work is to experiment the case study in the real setting of the *myJoghurt* plant to assess the enforcement approach and

to evaluate the agent’s performance. From a theoretical point of view, our proposal opens the path to explore transparency and explainability of the normative functioning of multi-agent systems taking profit of this agent-centric reasoning on norms and sanctions extending the approach proposed in [31]. Future extension is to consider the distinctions between different types of obligations (i.e., an obligation as an action to be performed and/or to achieve a certain state) and how agents adapt their sanctioning strategy effectively, especially in complex social scenarios. Another extension concerns the connection of our sanctioning mechanisms with the approach developed in Situated Artificial Institution by connecting the sanctioning process directly to the environment [5].

## Acknowledgment

Partially funded by ANR-FAPESP NAIMAN project (ANR-22-CE23-0018-01), GFA through RAMP-UP II project, and CAPES, The Brazilian Agency for Higher Education, under the project PrInt CAPES-UFSC “Automation 4.0”.

## References

1. Alechina, N., Dastani, M., Logan, B.: Programming norm-aware agents. In: Conitzer, V., Winikoff, M., Padgham, L., van der Hoek, W. (eds.) Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012). vol. 2, pp. 1057–1064. IFAAMAS, Valencia (2012)
2. Andrighetto, G., Conte, R., Turrini, P., Paolucci, M.: Emergence in the loop: Simulating the two way dynamics of norm innovation. In: Dagstuhl Seminar Proceedings. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2007)
3. Boissier, O., Bordini, R.H., Hubner, J., Ricci, A.: Multi-agent oriented programming: programming multi-agent systems using JaCaMo. Mit Press (2020)
4. Bordini, R.H., Hübner, J.F., Wooldrige, M.: Programming Multi-Agent Systems in AgentSpeak using *Jason*. John Wiley & Sons (2007)
5. de Brito, M., Hübner, J.F., Boissier, O.: Coupling the normative regulation with the constitutive state management in situated artificial institutions. *The Knowledge Engineering Review* **34** (2019)
6. Broersen, J., Dastani, M., Hulstijn, J., Huang, Z., van der Torre, L.: The boid architecture: Conflicts between beliefs, obligations, intentions and desires. In: Proceedings of the Fifth International Conference on Autonomous Agents. pp. 9–16. ACM, New York, NY (2001)
7. Conte, R.: Emergent (info)institutions. *Cognitive Systems Research* **2**(2), 97–110 (2001)
8. Criado, N., Argente, E., Noriega, P., Botti, V.: Reasoning about norms under uncertainty in dynamic environments. *International Journal of Approximate Reasoning* **55**(9), 2049–2070 (2014)
9. Dastani, M., Tinnemeier, N.A.M., Meyer, J.J.C.: A programming language for normative multi-agent systems. In: Dignum, V. (ed.) *Handbook of Research on Multi-Agent Systems: Semantics and Dynamics of Organizational Models*. pp. 397–417. IG Global, Hershey, PA (2009)
10. Fagundes, M.S., Billhardt, H., Ossowski, S.: Normative reasoning with an adaptive self-interested agent model based on markov decision processes. In: Kuri-Morales, A., Simari, G.R. (eds.) *Advances in Artificial Intelligence – IBERAMIA 2010*. pp. 274–283. Springer (2010)
11. Fagundes, M.S., Ossowski, S., Meneguzzi, F.: Imperfect norm enforcement in stochastic environments: An analysis of efficiency and cost tradeoffs. In: Bazzan, A.L., Pichara, K. (eds.) *Advances in Artificial Intelligence – IBERAMIA 2014*. pp. 523–535. Springer International Publishing, Cham (2014)
12. Gibbs, J.P.: Sanctions. *Social Problems* **14**(2), 147–159 (1966)



13. Hollander, C.D., Wu, A.S.: The current state of normative agent-based systems. *Journal of Artificial Societies and Social Simulation* **14**(2), 6 (2011)
14. Hübner, J.F., Boissier, O., Bordini, R.H.: A normative programming language for multi-agent organisations. *Annals of Mathematics and Artificial Intelligence* **62**(1-2), 27–53 (2011)
15. Hübner, J.F., Boissier, O., Kitio, R., Ricci, A.: Instrumenting multi-agent organisations with organisational artifacts and agents: “giving the organisational power back to the agents”. *Journal of Autonomous Agents and Multi-Agent Systems* **20**(3), 369–400 (2010)
16. Jensen, A.S.: The AORTA Reasoning Framework – Adding Organizational Reasoning to Agents. Ph.D. thesis, Danmarks Tekniske Universitet (DTU) (2015)
17. Kollingbaum, M.J., Norman, T.J.: Supervised interaction: Creating a web of trust for contracting agents in electronic environments. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1*. p. 272–279. AAMAS '02, ACM (2002)
18. Kollingbaum, M.J., Norman, T.J.: NoA - A normative agent architecture. In: Gottlob, G., Walsh, T. (eds.) *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*. pp. 1465–1466 (2003)
19. Kollingbaum, M.J., Norman, T.J.: Informed deliberation during norm-governed practical reasoning. In: Boissier, O., Padget, J., Dignum, V., Lindemann, G., Matson, E., Ossowski, S., Sichman, J.S., Vázquez-Salceda, J. (eds.) *Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems*. pp. 183–197. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
20. Land, K., Nardin, L., Vogel-Heuser, B.: Increasing robustness of agents’ decision-making in production automation using sanctioning. In: *2023 IEEE 21st International Conference on Industrial Informatics (INDIN)*. pp. 1–6 (2023)
21. Lee, J., Padget, J., Logan, B., Dybalova, D., Alechina, N.: N-jason: Run-time norm compliance in agentspeak(1). In: Dalpiaz, F., Dix, J., Birna van Riemsdijk, M. (eds.) *Engineering Multi-Agent Systems. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 8758, pp. 367–387. Springer (2014), 2nd International Workshop on Engineering Multi-Agent Systems, EMAS 2014 ; Conference date: 05-05-2014 Through 06-05-2014
22. López y López, F., Luck, M., d’Inverno, M.: A normative framework for agent-based systems. *Computational & Mathematical Organization Theory* **12**, 227–250 (2006)
23. Luck, M., Mahmoud, S., Meneguzzi, F., Kollingbaum, M., Norman, T.J., Criado, N., Fagundes, M.S.: Normative agents. In: Ossowski, S. (ed.) *Agreement Technologies*, pp. 209–220. Springer (2013)
24. Meneguzzi, F., Luck, M.: Norm-based behaviour modification in bdi agents. In: *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*. vol. 1, pp. 177–184 (2009)
25. Meneguzzi, F., Vasconcelos, W., Oren, N., Luck, M.: Nu-bdi: norm-aware bdi agents. In: *Proceedings of the 10th European Workshop on Multi-Agent Systems, Dublin, Ireland* (2012)
26. Nardin, L.G., Balke-Visser, T., Ajmeri, N., Kalia, A.K., Sichman, J.S., Singh, M.P.: Classifying sanctions and designing a conceptual sanctioning process model for socio-technical systems. *The Knowledge Engineering Review* **31**(2), 142–166 (2016)
27. Sabater-Mir, J., Paolucci, M., Conte, R.: Repage: Reputation and image among limited autonomous partners. *Journal of Artificial Societies and Social Simulation* **9**, 1–18 (2006)
28. dos Santos Neto, B.F., da Silva, V.T., de Lucena, C.J.: Developing goal-oriented normative agents: The nbdi architecture. In: *Agents and Artificial Intelligence: Third International Conference, ICAART 2011, Rome, Italy, January, 28-30, 2011. Revised Selected Papers 3*. pp. 176–191. Springer (2013)
29. Ullmann-Margalit, E.: *The emergence of norms*. Oxford University Press (1977)
30. Wooldridge, M.: *An Introduction to MultiAgent Systems*. Wiley Publishing, 2nd edn. (2009)
31. Yan, E., Burattini, S., Hübner, J.F., Ricci, A.: Towards a multi-level explainability framework for engineering and understanding bdi agent systems. In: Falcone, R., Castelfranchi, C., Sapienza, A., Cantucci, F. (eds.) *Proceedings of the 24th Workshop “From Objects to Agents” (WOA 2023)*. CEUR Workshop Proceedings, vol. 3579. CEUR-WS.org, Rome (2023)