Institutionalized Robots: Using Institutions to Make Robots Socially Aware

Stevan Tomic[1] and Julien Bidot[1] and Pedro Lima[2] and Federico Pecora[1] and Alessandro Saffiotti[1]

***Abstract*— Human interactions within a society are often regulated by social norms which involve assuming certain roles and acting differently depending on the assumed role and social context. In this paper we concern ourselves with how to employ social norms in societies with humans and robots. In particular, we focus on the issue of making social norms amenable to robots for reasoning, and address issues connected to how norms are represented and which cognitive processes should be used by robots to account for social norms in their behavior. We exploit a notion of institution to encapsulate social norms and to associate them with roles, artifacts and actions. The paper also shows an illustrative example and experimental trial involving the use of an institution with a real robotic system.**

I. INTRODUCTION

Social norms are widely used in daily interaction among humans. Some are simple, like being silent when someone is sleeping, but human interactions in society are often regulated by more implicit and complex social norms. These may involve assuming certain roles and acting differently depending on the role and social context.

We believe that there are compelling reasons to make robots understand social norms. Using norms that stem from human societies can lead to better human-robot interaction. Social norms are important for the acceptance of humans into societies, and there is no reason to believe that this would not be the case for robots. Social norms can also facilitate decision-making in large multi-robot systems, as the robots involved can directly use social norms to guide their actions, without having to autonomously reason about which actions should be performed.

As noted by Boella [1], research on social norms spans multiple fields, ranging from sociology [2], [3] to philosophy [4] and computer science [5]. In computer science, norms and social interactions are especially studied in normative multi-agent systems (MAS). Many formal models and frameworks have been proposed in this area, including S-MOISE+ [6], ISLANDER [7], ORA4MAS [8], and OPERA [9]. Dignum & Dignum argue that norms are a way for agents to cope with the challenge of social order [10].

Encoding social norms in an artificial system is complicated by the fact that norms may change depending on the agent's role, on its activity, and on the social situation [3]. One way to bind norms to social context is through concepts like organizations and institutions, mutated from the economics and social sciences [11], [12]. Organizations and institution may guide the agents toward common goals that could not be achieved if they were acting separately [13], [14]. In the area of MAS, Horling and Lesser [15] define an organization as a collection of roles, relationships and authority structures. So and Durfee [16] divide the components of an organization into three main classes: environmental factors, like tasks and resources; structural factors, like roles and norms; and agent factors. Other researchers have formalized the concept of institution for trading and auction scenarios [17], [18].

Institutions describe social situations, define roles in social interactions, and provide a normative dimension by binding roles to obligations, prohibitions and permissions. Institutions thus provide a way to encode possible social dynamics between agents and possibly between robots. Lima and Pereira [19], [20] have developed for the first time an institutional framework that has been used in real robots. In their framework, institutions are defined in terms of Petri nets. While this gives a sound mathematical foundation, it falls short of task planning. The latter is, in our opinion, a crucial cognitive task in ensuring that robots adhere to social norms while acting autonomously in a human environment.

The goal of this paper is to propose a method for making robots *socially aware*, that is, to have some sense of the surrounding social situation and to be able to act socially in a human environment. To do so, we combine ideas from three areas: normative MAS, social robotics, and planning. We build upon our previous work in human-aware planning, where we have extended automated planning algorithms to account for human needs, behaviors, and social context [21], [22], [23]. In that work, social norms were not precisely formalized, but they were encoded as part of the planning domain, which makes modeling of social context cumbersome and counter-intuitive. This paper makes two main contributions compared to our previous work: (1) we use the notion of *institution* to model and encapsulate social norms that belong to a given social context; (2) we provide a framework to represent and reason about institutions in a robotic system. To make the presentation concrete, we illustrate our framework on a simple but real robotic example.

In the next section, we describe our general concept of institution. We then describe in Section III how an institution can be used to achieve adherence of robots to social norms, through the notion of *cognitive tasks*, that is, processes for institutional reasoning. Finally, in Section IV we show an illustrative example and experimental trial involving the use of an institution in a robotic system.

*This work was funded by the EC Seventh Framework Programme (FP7/2007-2013) grant agreement no. 601033 MOnarCH.

[1]Center for Applied Autonomous Systems (AASS), Orebro University, Sweden, <name.surname>@aass.oru.se

[2]Instituto de Sistemas e Robótica, Instituto Superior Técnico, Lisboa, Portugal pal@isr.ist.utl.pt

II. INSTITUTIONS

While a universally agreed definition of organization or of institution is not available, we adopt the following view on what an institution is: "a stable pattern/structure of joint activity that may constrain or affect the actions and interactions of agents towards some purpose [24]. A social and normative perspective is often seen as a way to simplify the development and maintenance of multi-agent systems [25]: we contend that the same is true for multi-robot systems or for mixed robot-human societies.

Below, we first introduce an institutional representation suitable for describing a *shared social space* [26] among robots or among robots and humans in a physical environment; we then discuss how institutions can be used to enforce adherence to social norms by robots. We keep the discussion general, in the sense that it does not assume how institutions are instantiated — e.g., in an organization centered or agent centered way [24]. We discuss possible concrete representations of the elements that make up an institution after having discussed the cognitive tasks for which institutions are useful.

Definition 1: An *institution* $\mathcal{I}$ is a five element tuple $\mathcal{I} = \langle Arts, Roles, Acts, Norms, Meta \rangle$, where:

- *Arts* is a set of physical or mental objects (*artifacts*) that have some meaning for the institution. E.g., a schoolroom in the case of teaching, virtual currency in the case of an online transaction, a road or a lane separation line in a traffic institution (e.g., [27]).
- *Roles* is a set of responsibilities that are to be assumed by the agents involved in the institution. Roles are associated with specific actions (in *Acts*) to achieve or execute. The association between roles and actions are specified in the *Norms*.
- *Norms* is a set of relations between elements in Arts ∪ *Roles* ∪ *Acts*. Norms tell which (social) actions should be assigned to which role, and define relations between roles, tasks and artifacts.
- *Acts* is a set of (social) actions (or tasks) that the fillers of *Roles* can be engaged in. Each element in this set is associated to a particular role, and can only be carried out by an agents filling that role.
- *Meta* is a set of meta-relations on *Norms*. These relations represent the conditions under which different norms apply.

Norms play a central role in our institution formalism. They specify what kinds of actions (or tasks) must be associated with each role. For example, the role 'Buyer' in a trading institution must be associated with two actions: (1) *pay* and (2) *receive goods*. From the normative point of view, action can be seen as obligations that the agent has to achieve once it assumes particular role: these are sometimes called "positive constraints" [28]. Norms can also describe the relations between different roles, actions, artifacts, etc. For instance, they could specify when it is expected for a buyer to pay, by defining a temporal relation *before* between the buyer's action *pay* and the buyer's action *receive*, so that norm states *the buyer should pay before he/she receives the goods*. By checking whether all relations hold, one can say if the norms are satisfied. We say that an agent *adheres to a norm* (or the norm holds) if all relations involving the agent's role are satisfied. For example, if the norm is specified as the relation

```
before(pay(Buyer,Seller),
        receiveGoods(Buyer,Seller))
```

then the norm is adhered to if the action of paying is carried out by the agent filling the role of buyer before performing the action of receiving goods from the agent filling the role of seller.

The above norms are *regulative*, which specify appropriate and expected behaviors [29]. Norms can also be *constitutive*, e.g., defining *count-as* relations, to associate real world social situation with the specific institution [30], [29].

Roles can also be seen as normative concepts "focusing on what is proper for a person in a particular category to do" [31]. The above roles represent agents that act within the institution. Another type of role is the "supervision" role, whose main purpose is to check that all of the norms in the institution are satisfied, and if not try to enforce them. Agents taking this role are called "Governors" in ISLANDER [7] and "Managers" in S-MOISE+[6]. Norms associated with this kind of roles are called "repair" norms, since they specify how to repair regulative norms (social order) in an institution [32]. Supervision roles and constitutive norms are further discussed in the next section.

The set *Meta* specify the relations that hold between norms. It should not be confused with relations defining the norms themselves. For example, it should be possible to specify that either the first or second of two norms applies, or they both apply. This may be the case in an institution regulating road traffic: one norm specifies that cars should drive on the right side of the road, another that they should drive on the left; a third norm states that cars should stop at red lights. The first two norms are mutually exclusive, and apply depending on the country; the third norm applies always. This logic is expressed as relations in the *Meta* set of the traffic institution.

The following institution captures an interaction (or social situation) where two agents are involved in a simple trading situation. This institution can be represented graphically as shown in Figure 1.

- *Arts*: Goods, Money
- *Roles*: Buyer, Seller
- *Acts*: Give(*from*, *to*, *artifact*), Receive(*from*, *to*, *artifact*)
- *Meta*: ∅
- *Norms*:

  1) $from, to \in$ {Buyer, Seller}
  2) $artifact \in$ {Goods, Money}
  3) $from \neq to$
  4) before(Give(Buyer, Seller, Money), Receive(Seller, Buyer, Goods))

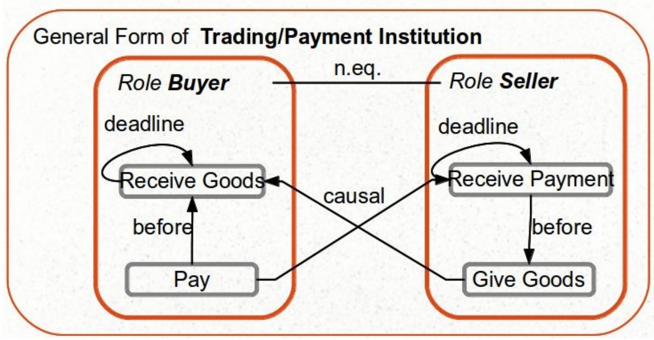

Fig. 1. Example of trading/payment institution. There are two roles with associated tasks to pay and receive goods, and to receive payment and provide the goods. Tasks are bound by relations, e.g., the buyer has to pay *before* receiving the goods, which implies that the seller only gives the goods after receiving the payment. Two *deadline* relations have also been asserted for the two tasks, e.g., the goods should be given no later than 7 days after the payment is received.

5) before(Receive(Buyer, Seller, Money), Give(Seller, Buyer, Goods))
6) deadline(Give(Seller, Buyer, Goods), Receive(Buyer, Seller, Money), 7 days)

There are two roles in this institution, a 'Buyer' and a 'Seller'. The former can pay and receive goods from the latter. Similarly, the seller receives payment from and gives goods to the buyer. Meaningful artifacts are hence 'Goods' and 'Money'. Norms that regulate trading state that the seller and buyer have different fillers, and that payment has to occur before goods are delivered. Also, there is a maximum time within which the seller has to provide the goods to the buyer after payment. These relations constitute the social norms that regulate the trading institution. Since all norms apply concurrently, the set *Meta* of the institution is empty.

Note that the relations stating that payment occurs before receiving goods can be different in different payment situations. It is easy to imagine a more general trading institution in which the temporal order of payment and goods delivery depends on the context. If, for instance, the context is a restaurant, then payment occurs after receiving goods, whereas the opposite holds if the context is a store. In this more general trading institution, one of two alternative sets of norms thus should hold. The theory describing under which circumstances each set of norms holds is codified in the *Meta* element of the institution tuple.

III. Using Institutions for Robots

We now turn our attention to the issue of how and when an institution can be used for robots. Key issues here include: (1) assessing whether an institution is active; (2) assigning institutional roles to agents; (3) assessing whether or not the norms stated in an institution are adhered to; and (4) enforcing that agents adhere to the (regulative) norms of an institution. Each of the above questions maps to a particular *cognitive task* that is to be carried out. Below, we show how the above four cognitive tasks can be seen as processes that reason with institutions, we hint at possible options to realize these processes, and discuss some of the difficulties especially in the case of robotic agents.

To address question 1, we assume a cognitive task that monitors the environment by checking for specific situations of interest. If such a situation occurs, a specific institution is said to be *active*. Imagine a robotic car entering a one-way street. The institution of interest here is the "one-way street" institution, which regulates traffic flow in a one-way street. It is active if a one-way street sign is observed upon entering the street. Hence, context recognition is a fundamental cognitive task for an institutional framework.

Context recognition may require non-trivial inference. Suppose that the sign is missing, or not visible: in order to recognize that the current context is one for a "one-way street" institution, the robotic car may have to observe that all cars move in the same direction on both sides of the road. Robotic systems have inherently limited perceptual capabilities, which suggests that reasoning may be necessary in order to infer the current context from observations, and hence to recognize which institution is active. This is especially true when the social context should be inferred from human behavior of from communicative acts. Context inference can be performed in a model-driven way using explicitly represented criteria, e.g., rules, that make a particular institution active. These criteria describe when a social situation *counts-as* a given institution, and can be seen as a type of constitutive norms (see previous section).

Once an institution is active, the role assignment problem should be solved — issue (2) above. In particular, how does a robot become a *participant* in an institution with a specific role in it? Several strategies can be used. One can model the criteria for role assignment, e.g., through a fitness function that is used to perform distributed role assignment via a bidding scheme. Another possibility is to define rules for each role that define the privileges, the physical abilities and the cognitive capabilities necessary to fill that role and to carry out its corresponding tasks. For example, not every person can assume the role of seller if the artifacts for sale is medication. Note that role assignment is potentially very related to the social norms specified in an institution. For instance, it may not be socially acceptable for one agent to be both buyer and seller, which in turn affects how roles are assigned. Role assignment may also include a scheduling problem: e.g., it may be possible for a given agent to cover two roles but not concurrently.

The third cognitive task consists in checking if the relations in all the norms in the institution hold. This can be seen as a continuous context inference processes, which is often referred to as *monitoring*. A robot being part of an institution with a given role can only satisfy norms concerning its own role. To satisfy these norms, robots may manipulate the environment, alter their own course of action, or interact with other agents that are assigned to other roles.

In order to ensure adherence to norms, monitoring should be complemented by the additional cognitive tasks of *goal generation* and *planning*, to use "repair norms" for supervisor roles — task 4. The generated goals should be inferred

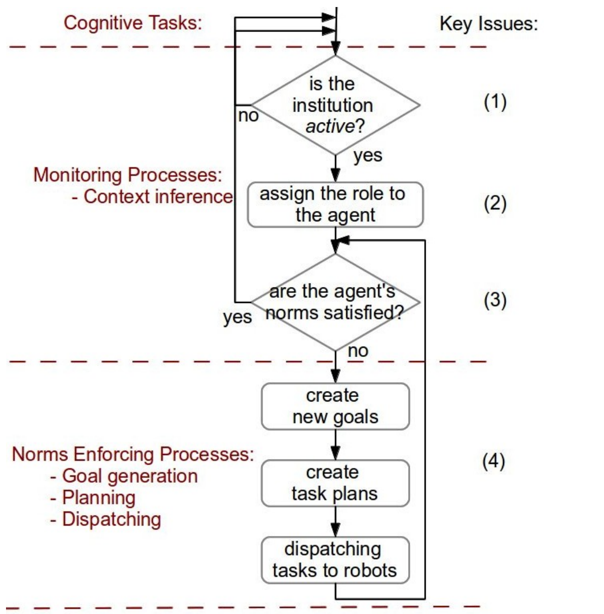

Fig. 2. How institutions are reasoned about through the four cognitive tasks. Context inference is involved in inferring if an institution is active (task 1), assigning robots to specific roles (task 2) and checking if norms are satisfied (task 3). Bringing about the adherence to norms is the purpose of goal generation and planning (task 4). In a robotic system, plans are then generated, dispatched to robots and monitored.

starting from the situations in which the robot does not adhere to role's norms. Examples in the literature suggest how model-driven approaches may be used to generate goals [22], [33]. Whenever rules describing norms are not satisfied, the corresponding goal should be generated to enforce the unsatisfied norm. The generated goal alters how agents behave in the institution, and is a result of institutional reasoning and norm maintenance.

Generated goals correspond to norms in the institution. Goals may be more than just states (predicates) to be achieved. For example, in the trading institution, a buyer robot must realize the tasks of paying and receiving goods, but it must also make sure that the right temporal and spatial relations between these tasks hold. Moreover, performing tasks may require plans in which actions are connected by complex causal, temporal, spatial and resource relations: e.g., paying may require to make money from the wallet, using a credit card, or going to a bank; and receiving the goods may require going to post office or arranging a transportation service. The literature on constraint-based planning offers approaches to address many of these issues [22], [34]. In the case of robots, planning also require a subsequent phase of plan dispatching and execution, where actions are assigned to robots and their execution is monitored.

The cognitive tasks discussed above, and their relation to institutions, are summarized in Figure 2.

IV. Illustrative Experiment

We illustrate the above concepts with a simple experiment conducted using a real robot. This experiment shows a preliminary implementation of an institutional framework for use with robots that operate in a social environment, and it shows some concrete choices made for representing institutions in a specific "school teaching" scenario. We also select specific reasoning tools for achieving the four cognitive tasks described above. We focus here on the "supervisor" role, and how it achieves the adherence to the social norms defined in our teaching institution.

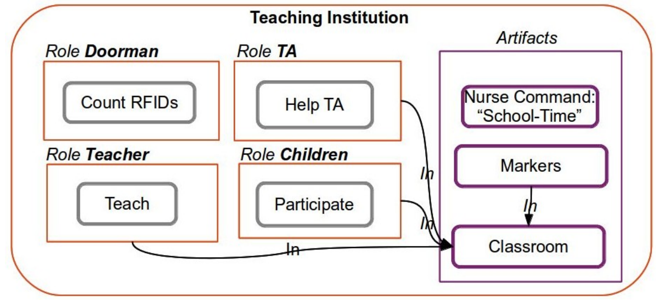

Fig. 3. Teaching institution. Norms in the institution are given through temporal and spatial relations. When the institution is active, markers, teacher, TA and children, should all be in the classroom. The supervisor role is not shown.

A. Scenario

The following scenario is borrowed from the ongoing EU project Monarch [35]. The project's goal is to help staff members in a pediatric ward to entertain children through games and education activities. Since the hospital is an environment which is rich in social norms, one of the project's sub-goals is to use ideas from institutional robotics [36], adapting them to planning for robots.

Occasionally, children are gathered in the hospital's classroom in order to engage in a teaching activity. In order for the class to start, all of the involved children have to be in the classroom. Also there should be a teacher and a teaching assistant present in the classroom, and the classroom should contain fresh markers for writing. The class is usually started when a nurse announces that it is class time and then gathers the children. These norms constitute what we call the "teaching" institution — see Figure 3. Norms and roles in this institution include that the teacher should have a teaching assistant (role TA); that the children, teacher and TA should be in the classroom; and that there should be markers in the classroom. Every child in the ward wears a bracelet with a Radio Frequency Identification (RFID) tag. Assessing that all children are in the classroom is done by a "doorman" using an RFID reader at the door of the classroom: a suitable norm states that this check should occur. The teaching activity should start when a nurse informs a robot that it is a "school time". Note that norms in this institution include both temporal and spatial relations.

B. Experimental Setup and Methods

a) Environment and hardware setup.: The experiment was executed on a laptop (8x Intel(R) Core(TM) i7-3740QM CPU and 8GB RAM) running Ubuntu 12.04.5 LTS (64-bit) and ROS Hydro. It was carried out in the PEIS Home environment [37] (see Figure 4), of which 2 rooms were used to represent a classroom and a playroom. We used two robotic platforms developed by the Monarch project consortium [38], called `mbot1` and `mbot2`.

b) Planning approach.: The reasoning capabilities needed for the cognitive tasks underlying our institutional framework are context inference, goal generation and planning. We chose a constraint based planning approach (CBP) extended for robotic systems [22], [39] for realizing all three reasoning capabilities. This approach allows us to represent temporal constraints (both qualitatively and quantitatively), to define custom constraints, but most importantly, it allows us to perform integrated context inference and planning. Specifically, the planner can generate new goals depending on inferred context. The capabilities of the planner are sufficient for dealing with the discussed mechanisms underlying the use of institutions, in the particular case of the institution at hand. The particular techniques developed to achieve integrated context inference, goal generation and planning, are outside the scope of this paper.

c) Implementation.: We have partially implemented a version of the concepts discussed in the previous section, in a software module called Institution Manager (IM). In our implementation, the IM is realized as a centralized coordinator, which uses institutions to govern the robots. In practice, it is realized as a ROS node which encapsulates the planner. Limitation in the current implementation include the fact that institutional norms (temporal and spatial relations) are specified directly into the planners domain, as well as the rules for context inference for assessing whether an institution is active or norms are adhered to. Richer implementations are the subject of our future work.

C. The Trials

In addition to planning and context recognition capabilities, the system includes a situation awareness module which continuously gathers information about the environment and publishes it as topics to which the planner subscribes to. In our simple experiment, we manually input some of these topics. All of the inputs are represented as sensor values that are further processed by the planner. This is a way to simulate the *monitoring processes*. The most important inputs for this discussion are: nurseCommand; markerLocationSensor; childLocationSensor; mbot1Location; mbot2Location; teacherLocationSensor. In the presented scenario we set these inputs so that all norms of the "teaching institution"' are not adhered to for the purpose of demonstrating the generation of a plan that enforces institutional norms. Thus, the initial settings are as follows: the children are in the playroom, and the teacher and the markers are not in the classroom. Also, no robot is in position at the door.

When the nurse announces school time, the planner's context inference processes will infer that the "teaching institution" is active. From the rest of the inputs and inference mechanisms the planner knows that this institution's norms are not valid, so it generates goals for both robots and starts to plan to achieve them, thus, at the same

Fig. 4. The PEIS-home environment using in the experimental trial, with a child and the two "mbot" robots.

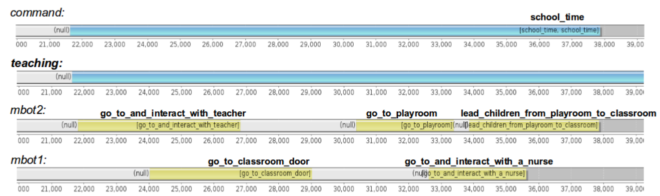

Fig. 5. Timelines. The figure shows four timelines. The first represents the command given by the nurse. When the command is given, the second timeline (representing activation of the teaching institution) also becomes active as a result of *context inference*. Once this happens, the planner plans how to enforce the norms defined by that institution. As a result of goal generation and planning, behaviors are dispatched for execution on the two robots mbot1 and mbot2. This can be observed in the last two timelines, representing intervals of the executed behaviors, from the time they are dispatched to the time when they are finished.

time, enforcing the institutional norms (see 5). The plan consists of five actions that are assigned (by the planner itself) to the two available robots. In this run, the robot `mbot2` performs behavior 'go_to_and_interact_with_teacher', as the teacher is not in the classroom and the robot must enforce norm in(teacher, classroom). Similarly, the robot `mbot1` is assigned behavior 'go_to_and_interact_with_nurse', which consists in reaching the nurse to inform him/her about the missing markers in the classroom. Upon finishing these behaviors, other behaviors are dispatched to enforce the remaining social norms. Thus, for `mbot1` the planner dispatches 'go_to_classroom_door', and for `mbot2` it dispatches 'go_to_playroom'. Upon finishing the latter, `mbot2` carries out 'lead_children_from_playroom_to_classroom'. The last two behaviors are implemented on real robot. As a result, the robot moved to reach to the playroom (a predefined area in the PEIS home) and after finishing this action, it moved back to the schoolroom executing a simplified version of 'lead_children_from_playroom_to_classroom'.

V. Conclusions

We have presented an initial investigation on the representation and use of institutions in robotic systems. Interestingly, our notion of institution allows us to define the cognitive tasks that are relevant for robotic agents operating in social contexts. We have argued that planning and context recognition are fundamental reasoning capabilities needed for the realization of these cognitive tasks. Those abilities seem to lack in the frameworks reported in the literature on normative multi-agent systems so far. By contrast, planning, context recognition and normative reasoning are tightly coupled together in our proposed approach. We have shown this coupling in a concrete institutional robotic system, in which a unique planning process accounts for norms (expressed as constraints), for obligations (expressed as goals), and for context recognition (used for norm violation detection and for activating institutions).

Our approach provides a normative framework for programming multi-robot systems and for defining social structures via natural concepts like roles, norms and artifacts. It also allows us to uniformly represent norms of different types, including social, causal, temporal, or spatial constraints. We believe that this will be pivotal to the use of our framework in real-world environements, where robots and humans may live in complex shared social space that involve awareness of all the above aspects.

The work we have presented is still in its early stages. In the future, we will address the many open questions described in this paper. In particular, issues related to assigning agents to roles and formal languages for representing and reasoning about social norms will be addressed first. In the medium term, we will investigate the issue of relations between institution, including deontic relations and taxonomic relations. The long-term goal of this research is to develop a complete framework for institution-based reasoning for use in social robotic systems.

REFERENCES

[1] G. Boella, L. Torre, and H. Verhagen, "Introduction to normative multiagent systems," *Computational & Mathematical Organization Theory*, vol. 12, no. 2-3, pp. 71–79, 2006.

[2] J. P. Gibbs, "Norms: The problem of definition and classification," *American Journal of Sociology*, vol. 70, no. 5, pp. 586–594, 1965.

[3] G. Therborn, "Back to norms! On the scope and dynamics of norms and normative action," *Current Sociology*, vol. 50, no. 6, pp. 863–880, 2002.

[4] C. E. Alchourrón and E. Bulygin, *Normative systems*. Springer Wien, 1971, vol. 5.

[5] J.-J. C. Meyer and R. J. Wieringa, *Deontic logic in computer science: normative system specification*. John Wiley and Sons Ltd., 1994.

[6] J. Hübner, J. Sichman, and O. Boissier, "S-Moise+: A middleware for developing organised multi-agent systems," in *Coordination, organizations, institutions, and norms in multi-agent systems*, ser. LNAI. Springer, 2006, no. 3913, pp. 64–77.

[7] M. Esteva, D. De la Cruz, and C. Sierra, "ISLANDER: An electronic institutions editor," in *Proc of the Int Conf on Autonomous Agents*, 2002, pp. 1045–1052.

[8] J. Hübner, O. Boissier, R. Kitio, and A. Ricci, "Instrumenting multi-agent organisations with organisational artifacts and agents: Giving the organisational power back to the agents," *Autonomous Agents and Multi-Agent Systems*, vol. 20, no. 3, pp. 369–400, 2010.

[9] M. Dignum, "A model for organizational interaction: based on agents, founded in logic," Ph.D. dissertation, Utrecht University, NL, 2004.

[10] V. Dignum and F. Dignum, "Modelling agent societies: Co-ordination frameworks and institutions," in *Progress in artificial intelligence*, P. Brazdil and A. Jorge, Eds. Springer, 2001, pp. 191–204.

[11] J. R. Searle, "What is an institution," *Journal of institutional economics*, vol. 1, no. 1, pp. 1–22, 2005.

[12] E. Ostrom, *Understanding institutional diversity*. Princeton university press, 2009.

[13] G. Weiss, Ed., *Multiagent Systems - second edition*. MIT press, 2013.

[14] J. Ferber, O. Gutknecht, and F. Michel, "From agents to organizations: An organizational view of multi-agent systems," in *Agent-Oriented Software Engineering IV*. Springer, 2004, pp. 214–230.

[15] B. Horling and V. Lesser, "A survey of multi-agent organizational paradigms," *The Knowledge Engineering Review*, vol. 19, no. 04, pp. 281–316, 2004.

[16] Y.-p. So and E. H. Durfee, "Designing organizations for computational agents," in *Simulating Organizations*, M. Prietula, K. Carley, and L. Gasser, Eds. MIT Press, 1998.

[17] J. A. Rodriguez-Aguilar *et al.*, "On the design and construction of agent-mediated institutions," Ph.D. dissertation, Universidad Autónoma de Barcelona, 2001.

[18] P. Noriega, "Agent mediated auctions: the fishmarket metaphor," Ph.D. dissertation, Universidad Autónoma de Barcelona, 1997.

[19] J. N. Pereira, P. Silva, P. U. Lima, and A. Martinoli, "Formalization, Implementation, and Modeling of Institutional Controllers for Distributed Robotic Systems," *Artificial Life*, vol. 20, no. 1, 2014.

[20] J. N. Pereira, "Advancing social interactions among robots: an institutional economics-based approach to distributed robotics systems," Ph.D. dissertation, IST-EPFL Joint Doctoral Initiative, 2013.

[21] M. Cirillo, L. Karlsson, and A. Saffiotti, "Human-aware task planning: an application to mobile robots," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 1, no. 2, p. 15, 2010.

[22] F. Pecora, M. Cirillo, F. Dell'Osa, J. Ullberg, and A. Saffiotti, "A constraint-based approach for proactive, context-aware human support," *Journal of Ambient Intelligence and Smart Environments*, vol. 4, no. 4, pp. 347–367, 2012.

[23] S. Tomic, F. Pecora, and A. Saffiotti, "Too cool for school – adding social constraints in human aware planning," in *Proc of the Int Workshop on Cognitive Robotics at ECAI (CogRob)*, 2014.

[24] O. Boissier, J. F. Hübner, and J. S. Sichman, "Organization oriented programming: From closed to open organizations," in *Engineering Societies in the Agents World*. Springer, 2007, pp. 86–105.

[25] N. Alechina, N. Bassiliades, M. Dastani, M. De Vos, B. Logan, S. Mera, A. Morris-Martin, and F. Schapachnik, "Computational models for normative multi-agent systems," Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Tech. Rep., 2013.

[26] P. Noriega, J. Padget, H. Verhagen, and M. d'Inverno, "The challenge of artificial socio-cognitive systems," in *Proc of the Int Conf on Autonomous Agents and Multi-Agent Systems*, 2014.

[27] A. Ricci, M. Viroli, and A. Omicini, "Give agents their artifacts: the a&a approach for engineering working environments in MAS," in *Int Conf on Autonomous Agents and Multi-Agent Systems*, 2007, p. 150.

[28] N. Criado, E. Argente, and V. Botti, "A BDI architecture for normative decision making," in *Proc of the Int Conf on Autonomous Agents and Multi-Agent Systems*, 2010, pp. 1383–1384.

[29] G. Boella and L. van der Torre, "Substantive and procedural norms in normative multiagent systems," *Journal of Applied Logic*, vol. 6, no. 2, pp. 152–171, 2008.

[30] H. Aldewereld, S. Alvarez-Napagao, F. Dignum, and J. Vázquez-Salceda, "Engineering social reality with inheritance relations," in *Engineering Societies in the Agents World*, 2009, pp. 116–131.

[31] R. Harré and P. F. Secord, *The explanation of social behaviour*. Rowman & Littlefield, 1972.

[32] S. Panagiotidi, J. Vázquez-Salceda, and F. Dignum, "Reasoning over norm compliance via planning," in *Coordination, Organizations, Institutions, and Norms in Agent Systems*. Springer, 2013, pp. 35–52.

[33] J. Grosinger, F. Pecora, and A. Saffiotti, "Robots and bananas: Exploring deliberation in cognitive robots," in *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

[34] M. Mansouri and F. Pecora, "More knowledge on the table: planning with space, time and resources for robots," in *IEEE Int Conf on Robotics and Automation (ICRA)*, 2014, pp. 647–654.

[35] MOnarCH, "http://monarch-fp7.eu/," retrived May 2015.

[36] P. Silva and P. U. Lima, "Institutional robotics," in *Advances in Artificial Life*. Springer, 2007, pp. 595–604.

[37] A. Saffiotti, M. Broxvall, M. Gritti, K. LeBlanc, R. Lundh, J. Rashid, B. Seo, and Y.-J. Cho, "The PEIS-ecology project: vision and results," in *IEEE/RSJ Int Conf on Intelligent Robots and Systems*. IEEE, 2008, pp. 2329–2335.

[38] J. Messias, R. Ventura, P. Lima, J. Sequeira, P. Alvito, C. Marques, and P. Carriço, "A robotic platform for edutainment activities in a pediatric hospital," in *IEEE Int Conf on Autonomous Robot Systems*. IEEE, 2014, pp. 193–198.

[39] M. Di Rocco, F. Pecora, P. K. Sivakumar, and A. Saffiotti, "Configuration planning with multiple dynamic goals," in *AAAI Spring Symposium on Designing Intelligent Robots*. AAAI Press, 2013.