

Ontological Aspects of the Implementation of Norms in Agent-Based Electronic Institutions

Davide Grossi, Huib Aldewereld, Javier Vázquez-Salceda and Frank Dignum

Institute of Information and Computing Sciences
Utrecht University, The Netherlands
{davide,huib,javier,dignum}@cs.uu.nl

Abstract

In order to regulate different circumstances over an extensive period of time, norms in institutions are stated in a vague and often ambiguous manner, thereby abstracting from concrete aspects, which are relevant for the operationalization of institutions. If agent-based electronic institutions, which adhere to a set of abstract requirements, are to be built, how can those requirements be translated into more concrete constraints, the impact of which can be described directly in the institution? We address this issue considering institutions as normative systems based on articulate ontologies of the agent domain they regulate. Ontologies, we hold, are used by institutions to relate the abstract concepts in which their norms are formulated, to their concrete application domain. In this view, different institutions can implement the same set of norms in different ways as far as they presuppose divergent ontologies of the concepts in which that set of norms is formulated. In this paper we analyze this phenomenon introducing a notion of contextual ontology. We will focus on the formal machinery necessary to characterize it as well.

Introduction

Electronic institutions (eInstitutions) are agent environments that can regulate and direct the interactions between agents, creating a safe and stable environment for agents to act. This is accomplished by incorporating a number of norms in the institution which indicate the type of behaviour to which each agent should adhere within that institution. Similar to their human counterparts (legal systems are the eminent example), norms in eInstitutions should be stated in such a form that allows them to regulate a wide range of situations over time without need for modification. To guarantee this stability, the formulation of norms needs to abstract from a variety of concrete aspects, which are instead relevant for the actual implementation of an eInstitution (Dignum 2002; Grossi & Dignum 2004); this means that norms are expressed in terms of concepts that are, on purpose, kept vague and ambiguous (Hart 1961). On the other hand, whether a concrete situation actually falls under the scope of application of a norm is a question that, from the point of view of an effective operationalization of the institution, should be answered in a clear and definite way.

Copyright © 2004, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

The problem is that concrete situations are generally described in terms of ontologies which differ from the abstract ontology in which, instead, norms are specified. This means that, to actually give a concrete operational meaning to the norms, i.e., to implement them, a connection should be made which can integrate the two ontological levels (Dignum 2002). We need to determine what the concepts in the situation mean and somehow check them against the terms used in the norms. In other words, we have to see whether the concepts used to specify the situation are classified by (or counts as) the concepts used in the norm formulations; we have to formulate them in an ontology which makes the relation between the concrete and the abstract specifications explicit.

In our previous work we have done ... and now we are going to work on this. cites

This work is organized as follows. In the next section we will elaborate on how ontologies are used in institutions to determine the meaning of the concepts used in the norms under different contexts. Then, in **Modeling Contextual Ontologies**, we will present a formal framework in which it is possible to represent and reason about divergent ontologies (we will call them contextual terminologies) based on (Grossi, Dignum, & Meyer 2004a; 2004b). Using this framework we will formalize an example in section **Contextual Terminologies at Work**. After this, we will discuss the implementational aspects of our framework in section **Implementation** and we end the paper with some discussion, conclusions and future work.

Through the paper, we will use the regulations on personal data protection in several scenarios: the European Union, the Dutch Police, European Hospitals and Eurotransplant (a trans-national human organ allocation organization for transplantation purposes).

Institutions, Ontologies and Context

In order to properly implement norms in eInstitutions, we should first analyze the role of norms in human institutions. It is our thesis that institutions provide structured interpretations of the concepts in which norms are stated. In a nutshell, institutions do not only consist of norms, but are also based on ontologies of the to-be-regulated domain. For instance, whether something within a given institution counts as *personal data* and should be treated as such depends on how

that institution interprets the term *personal data*. What counts as personal data in a hospital, might not count as personal data in a police register and viceversa. Nevertheless, in both hospitals and police registers, if some piece of information is personal data, it should then be treated in accordance to the regional, national and/or international policies regulating the use of such data. That is to say, hospitals and police registers, although providing potentially inconsistent understanding of what personal data is, do share the normative consequences (rights, duties, prohibitions, etc.) attached to the classification of information as personal data.

This perspective on institutions, which emphasizes the semantic dependence of norm implementation, goes hand in hand with widely acknowledged positions on the normative nature of social reality. Institutions can be indeed seen as normative systems of high complexity, which consist of regulative as well as non-regulative components (Alchourrón & Bulygin 1986; Jones & Sergot 1993; 1992; Searle 1995; Boella & Van der Torre 2004), that is to say, which do not only regulate existing forms of behaviour, but they actually specify and create -via classification- new forms of behaviour. In the literature of legal theory, the non-regulative component of the issuing of norms has been labeled in ways that emphasize a classificatory, as opposed to a normative/regulative, character: *determinative rules* (Von Wright 1963), *conceptual rules* (Bulygin 1992), *qualification norms* (Peczenik 1989), *definitional norms* (Jones & Sergot 1992). This characteristic of the non-regulative, or classificatory, components of normative systems is intermingled with a second feature, namely the *constitutive, conventional* character of these components that have been therefore called also *constitutive rules* or *constitutive norms* (Ross 1968; Searle 1995). In this view, statements to the effect that racial data count as personal data establish that being a racial data constitute, in the sense of being a sufficient condition, for being personal data. However, this “constitution” is not absolute. It being conventional, it only holds within the specific institution in which that relation of constitution is effective, it is *contextual*.

Context

Human institutions hardly operate in isolation and therefore frequent references are made to other regulations and institutions. Institutions and their environment are interdependent, and each influence the other. In human societies the context of an institution includes regulations that are applied to the institution's internal and/or external behaviour. Therefore, when building eInstitutions, special attention should be given to the environment where the eInstitution will operate (?), as the environment may affect its specification (specially in the normative aspects of the eInstitution) and design; the regulations that apply to the environment should be considered and included by the designer inside the designing process of the eInstitution.

In agent-based eInstitutions, the agents should be provided with a model of the norms that may apply inside the institution and an ontology giving a semantic interpretation of the terms used. From the point of view of a single eInstitution, a single norm model and ontology are enough in order

to define the boundaries between acceptable and unacceptable behaviour. But problems may arise when agents have to operate in more than one eInstitution, each one having its own norms and norm interpretation, or when two eInstitutions have to inter-operate. The source of these problems is that, in most of real domains, norms are not universally valid but bounded to a given *context*. This is the case of norms, for instance, in Health Care, as they are bounded to trans-national, national and regional regulations, each of them defining a different normative context.

In those scenarios where more than one normative context should be modelled, trying to force a single vocabulary, theory and representation to model and reason about any situation on any context is not a good option. The alternative, first proposed by McCarthy in (McCarthy 1986)(?), is to include *context* as formal objects in the model. Therefore, most theoretical approaches have moved towards having an explicit representation of context. One of the most used approaches is the *box metaphor*, that is, considering context as a box:

[...] *Each box has its own laws and draws a sort of boundary between what is in and what is out.(?)*

With this idea, in (?), context in eInstitutions is defined formally as a subset of possible worlds where there is a shared vocabulary and a normative framework to be followed by a certain group of agents. In this view, an eInstitution is a context defining a) its vocabulary (by means of an ontology) and b) the norms that apply in that context. In parallel, the environments where the e-institution operates are also (super)contexts, being possibly nested (e.g. to model the nesting in regional/national/transnational environments).

Contextual Ontologies

Each normative context should therefore define a vocabulary to be shared by agents in a given context. It means that each context is associated with a domain ontology that defines the meaning of the terms that are both present in the norms, the actions the agent may carry and the terms in the communication with others. However, standard ontologies are not enough. As we have mentioned, contexts may be nested. Each context (defining their norms and an ontology) may contain other (sub)contexts inside (extending and/or modifying the norms and the ontology) or belong to one or several (super)contexts. Some kind of connection should be build between ontologies of inter-related contexts. This problem usually appears in multi-agent systems that should operate in a transnational, multi-lingual environment such as Europe. To illustrate this problem, let us return to the regulations on personal data protection. In European Union regulations ¹ *personal data* are defined as “those which al-

¹European Parliament created the 95/46/CE Directive (?) with the purpose of homogenizing legal cover on data protection, in order to warrant an appropriate protection level on each transfer inside the European Union. At the end of year 2000, the European Parliament extended the personal data regulations initiated by this norm by means of Regulation (CE) 45/2001 (?), which covers all that was already established by the Directive 95/46/CE, determines the penalty mechanism at the European level, and creates the figure of the Data Protection European Supervisor as an independent

low the identification of a person, and which reveal racial or ethnic origins, political opinions, religious or philosophic beliefs, trade union's affiliation, as well as data related to health or sexuality"... [example explanation here]

Modeling Contextual Terminologies

We will develop our formal framework keeping the following requirements in mind.

1. The formal framework should enable the possibility of expressing lexical differences, because institutions yield terminologies defined on different languages². In particular, in the institutional normative domain, we observe that more concrete contexts mean richer terminologies: talking about personal data comes down to talk about racial data, health data, etc.
2. It should provide a formal semantics (as more general as possible) for contextualizing terminological expressions.

Following these essential guidelines, a language and a semantics are introduced in this section. The language will make use of part of description logic syntax, as regards the concept constructs, and will make use of a set of operators aimed at capturing the interplay of contexts. In particular, we will introduce:

- A *contextual conjunction* operator. Intuitively, it will yield a composition of contexts: the contexts "personal data in hospitals" and "personal data in police registers" can be intersected on a language talking about data concerning the date of birth and alike generating a common less general context like "anagraphic data in hospitals and police registers".
- A *contextual disjunction* operator. Intuitively, it will yield a union of contexts: the contexts "personal data in hospitals" and "personal data in police registers" can be unified on a language talking about personal data generating a more general context like "personal data in hospitals or police registers".
- A *contextual negation* operator. Intuitively, it will yield the context obtained via subtraction of the context negated: the negation of the context "personal data in hospitals" on the language talking about data in general generates a context like "data which are not personal data in hospitals".
- A *contextual abstraction* operator. Intuitively, it will yield the context consisting in some information extracted from the context to which the abstraction is applied: the abstraction of the context "personal data in hospitals" on the language talking only about anagraphic data generates a context like "anagraphic data in hospitals". In other words, the operator prunes the information contained in the context "personal data in hospitals" keeping only what is expressible in the language which talks about anagraphic data and abstracting from the rest.

control authority.

²This is a much acknowledged characteristic of contextual reasoning in general (McCarthy 1986).

Finally, also *maximum* and *minimum* contexts will be introduced: these will represent the most general, and respectively the least general, contexts on a language. As it appears from this list of examples, operators will need to be indexed with the language where the operation they denote takes place. The point is that contexts always belong to a language, and so do operations on them³.

These intuitions about the semantics of context operators will be clarified and made more rigorous in Section where the semantics of the framework will be presented, and in Section where an example will be formalized.

Language

In a nutshell, the language we are interested in defining can be seen as a meta-language for TBoxes defined on \mathcal{AL} description logic languages, which also handle concept union, full existential quantification (we want to deal with concepts such as "either car or bicycle" and "persons which drive cars")⁴.

The alphabet of the language \mathcal{L}^{CT} (language for contextual terminologies) therefore contains the alphabets of a family of languages $\{\mathcal{L}_i\}_{0 \leq i \leq n}$. We take this family to be such that $\{\mathcal{L}_i\}_{0 \leq i \leq n} = \mathcal{P}^+(\mathcal{L})$, that is to say, each language \mathcal{L}_i is expanded by the "global" language \mathcal{L} .

Each \mathcal{L}_i contains a non-empty finite set \mathbf{A}_i of monadic predicates (A), i.e., atomic concepts, and a (possibly empty) set \mathbf{R}_i of dyadic predicates (R), i.e., atomic roles. These languages contain also concept constructors: each \mathcal{L}_i contains the zeroary operators \perp (bottom concept) and \top (top concept), the unary operator \neg (complement), and the binary operators \sqcap and \sqcup . Finally, operators \forall . (universal quantification) and \exists . (existential quantification) apply to role-concept pairs.

Besides, the alphabet of \mathcal{L}^{CT} contains a finite set of context identifiers \mathbf{c} , two families of zeroary operators $\{\perp_i\}_{0 \leq i \leq n}$ (minimum contexts) and $\{\top_i\}_{0 \leq i \leq n}$ (maximum contexts), two families of unary operators $\{abs_i\}_{0 \leq i \leq n}$ (contextual abstraction operator) and $\{\neg_i\}_{0 \leq i \leq n}$ (contextual negation operator), two families of binary operators $\{\wedge_i\}_{0 \leq i \leq n}$ (contexts conjunction operator) and $\{\vee_i\}_{0 \leq i \leq n}$ (contextual disjunction operator), one context relation symbol \preceq (context c_1 "is less general than" context c_2), and finally two contextual subsumption relation symbols " $\cdot : \cdot \sqsubseteq \cdot$ " (within context c , concept A_1 is a sub-concept of concept A_2) and " $\cdot : \cdot \sqsubseteq \cdot$ " (within context c , role R_1 is a subrole of role R_2) for, respectively, concept and role subsumption. Lastly, the alphabet of \mathcal{L}^{CT} contains also the sentential connectives \sim (negation) and \wedge (conjunction)⁵.

³Note that indexes might be avoided considering operators interpreted on operations taking place on one selected language, like the largest common language of the languages of the two contexts. However, this would result in a lack of expressivity that we prefer to avoid for the moment.

⁴This type of language is usually referred to as *ALUE*, or *ALC*. Within this type of languages the negation of arbitrary concepts is also enabled (Baader *et al.* 2002).

⁵It might be worth remarking that language \mathcal{L}^{CT} is, then, an expansion of each \mathcal{L}_i language.

Thus, the set Ξ of context constructs (ξ) is defined through the following BNF:

$$\xi ::= c \mid \perp_i \mid \top_i \mid \neg_i \mid \text{abs}_i \mid \xi_1 \wedge_i \xi_2 \mid \xi_1 \vee_i \xi_2.$$

Concept constructs and role constructs are defined in the standard way. The set Γ of concept descriptions (γ) is defined through the following BNF:

$$\gamma ::= A \mid \perp \mid \top \mid \neg\gamma \mid \gamma_1 \sqcap \gamma_2 \mid \gamma_1 \sqcup \gamma_2 \mid \forall\rho.\gamma \mid \exists\rho.\gamma.$$

The set P of roles descriptions (ρ) coincides with the set of all atomic roles.

The set \mathcal{A} of assertions (α) is then defined through the following BNF:

$$\alpha ::= \xi : \gamma_1 \sqsubseteq \gamma_2 \mid \xi : \rho_1 \sqsubseteq \rho_2 \mid \xi_1 \preceq \xi_2 \mid \sim \alpha \mid \alpha_1 \wedge \alpha_2..$$

Technically, a *contextual terminology* in \mathcal{L}^{CT} is a set of subsumption relation expressions on concepts or roles, which are contextualized with respect to the same context. This kind of sets of expressions are, in a nutshell, what we are interested in formalizing.

In Section the following symbols will be also used “ \sqsubseteq ” (within context c , concept A_1 is a proper subconcept of concept A_2), and “ \equiv ” (within context c , concept A_1 is equivalent to concept A_2). They can be obviously defined as follows:

$$\begin{aligned} \xi : \gamma_1 \sqsubseteq \gamma_2 &=_{def} \xi : \gamma_1 \sqsubseteq \gamma_2 \wedge \sim \xi : \gamma_2 \sqsubseteq \gamma_1 \\ \xi : \gamma_1 \equiv \gamma_2 &=_{def} \xi : \gamma_1 \sqsubseteq \gamma_2 \wedge \xi : \gamma_2 \sqsubseteq \gamma_1. \end{aligned}$$

Semantics

In order to provide a semantics for \mathcal{L}^{CT} languages, we will proceed as follows. First we will define a class of structures which can be used to provide a formal meaning to those languages. We will then characterize the class of operations on contexts that will constitute the semantic counterpart of the context operators symbols introduced in Section . Definitions of the formal meaning of our expressions and of the semantics of assertions will then follow.

Before pursuing this line, it is necessary to recollect the basic definition of a model for a language \mathcal{L}_i (Baader *et al.* 2002).

Definition 1. (Models for \mathcal{L}_i 's)

A model m for a language \mathcal{L}_i is defined as follows:

$$m = \langle \Delta_m, \mathcal{I}_m \rangle$$

where:

- Δ_m is the (non empty) domain of the model;
- \mathcal{I}_m is a function $\mathcal{I}_m : \mathbf{A}_i \cup \mathbf{R}_i \longrightarrow \mathcal{P}(\Delta_m) \cup \mathcal{P}(\Delta_m \times \Delta_m)$, such that to every element of \mathbf{A}_i and \mathbf{R}_i an element of $\mathcal{P}(\Delta_m)$ and, respectively, of $\mathcal{P}(\Delta_m \times \Delta_m)$ is associated. This interpretation of atomic concepts and roles of

\mathcal{L}_i on Δ_m is then inductively extended:

$$\begin{aligned} \mathcal{I}_m(\top) &= \Delta_m \\ \mathcal{I}_m(\perp) &= \emptyset \\ \mathcal{I}_m(\neg\gamma) &= \Delta_m \setminus \mathcal{I}_m(\gamma) \\ \mathcal{I}_m(\gamma_1 \sqcap \gamma_2) &= \mathcal{I}_m(\gamma_1) \cap \mathcal{I}_m(\gamma_2) \\ \mathcal{I}_m(\gamma_1 \sqcup \gamma_2) &= \mathcal{I}_m(\gamma_1) \cup \mathcal{I}_m(\gamma_2) \\ \mathcal{I}_m(\forall\rho.\gamma) &= \{a \in \Delta_m \mid \forall b, \langle a, b \rangle \in I_m(\rho) \\ &\quad \Rightarrow b \in \mathcal{I}_m(\gamma)\} \\ \mathcal{I}_m(\exists\rho.\gamma) &= \{a \in \Delta_m \mid \exists b, \langle a, b \rangle \in I_m(\rho) \\ &\quad \& b \in \mathcal{I}_m(\gamma)\}. \end{aligned}$$

Models for \mathcal{L}^{CT}

We can now define a notion of *contextual terminology model* (ct-model) for languages \mathcal{L}^{CT} .

Definition 2. (ct-models)

A ct-model \mathbb{M} is a structure:

$$\mathbb{M} = \langle \{\mathbf{M}_i\}_{0 \leq i \leq n}, \mathbb{I} \rangle$$

where:

- $\{\mathbf{M}_i\}_{0 \leq i \leq n}$ is the family of the sets of models \mathbf{M}_i of each language \mathcal{L}_i . In other words, $\forall m \in \mathbf{M}_i$, m is a basic description logic model of \mathcal{L}_i .
- \mathbb{I} is a function $\mathbb{I} : \mathbf{c} \longrightarrow \mathcal{P}(\mathbf{M}_0) \cup \dots \cup \mathcal{P}(\mathbf{M}_n)$. In other words, this function associates to each atomic context in \mathbf{c} a subset of the set of all models in some language \mathcal{L}_i : $\mathbb{I}(c) = M$ with $M \subseteq \mathbf{M}_i$ for some i s.t. $0 \leq i \leq n$. Notice that \mathbb{I} fixes, for each context identifier, the language on which the context denoted by the identifier is specified. We could say that it is \mathbb{I} itself which fixes a specific index i for each c .
- $\forall m', m'' \in \bigcup_{0 \leq i \leq n} \mathbf{M}_i$, $\Delta_{m'} = \Delta_{m''}$. That is, the domain of all basic description logic models m is unique. We establish this constraint simply because we are interested in modeling different (taxonomical) conceptualizations of a same set of individuals.

Contexts are therefore formalized as sets of models for the same language i.e., a set of instantiations of a terminology on that language. This perspective allows for straightforward model theoretical definitions of operations on contexts.

Operations on contexts

Before getting to this, let us first recall a notion of *domain restriction* ($\lceil \cdot \rceil$) of a function f w.r.t. a subset C of the domain of f . Intuitively, a domain restriction of a function f is nothing but the function $C \lceil f$ having C as domain and such that for each element of C , f and $C \lceil f$ return the same image. The exact definition is the following one: $C \lceil f(x) = \{y \mid y = f(x) \& x \in C\}$ (Casari 2002).

Definition 3. (Operations on contexts)

Let M' and M'' be sets of models:

$$\lceil_i M' = \{m \mid m \in M' \& m = \langle \Delta_m, \mathbf{A}_i \lceil \mathcal{I}_m \rangle\} \quad (1)$$

$$M' \pitchfork_i M'' = \lceil_i M' \cap \lceil_i M'' \quad (2)$$

$$M' \cup_i M'' = \lceil_i M' \cup \lceil_i M'' \quad (3)$$

$$-_i M' = \mathbf{M}_i \setminus \lceil_i M'' \quad (4)$$

Intuitively, the operations have the following meaning: operation 1 allows for abstracting the relevant content of a context with respect to a specific language; operations 2 and 3 express basic set-theoretical composition of contexts; finally, operation 4 returns, given a context, the most general of all the remaining contexts. Let us now provide some technical observations. First of all notice that operation \lceil_i yields the empty context when it is applied to a context M' the language of which is not an elementary expansion of \mathcal{L}_i . This is indeed very intuitive: the context obtained via abstraction of the context “dinosaurs” on the language of, say, “botanics” should be empty. Empty contexts can be also obtained through the \mathbb{m}_i operation. In that case the language is shared, but the two contexts simply do not have any interpretation in common. This happens, for example, when the members of two different football teams talk about their opponents: as a matter of fact, no interpretation of the concept opponent can be shared without jeopardizing the fairness of the match.

Formal meaning of Ξ and \mathcal{A}

The semantics of contexts constructs Ξ can be now defined.

Definition 4. (Semantics of contexts constructs)

The semantics of context constructors is defined as follows:

$$\begin{aligned} \mathbb{I}(c) &= M \in \mathcal{P}(\mathbf{M}_0) \cup \dots \cup \mathcal{P}(\mathbf{M}_n) \\ \mathbb{I}(\perp_i) &= \emptyset \\ \mathbb{I}(\top_i) &= \mathbf{M}_i \\ \mathbb{I}(\xi_1 \wedge_i \xi_2) &= \mathbb{I}(\xi_1) \mathbb{m}_i \mathbb{I}(\xi_2) \\ \mathbb{I}(\xi_1 \vee_i \xi_2) &= \mathbb{I}(\xi_1) \mathbb{u}_i \mathbb{I}(\xi_2) \\ \mathbb{I}(\neg_i(\xi)) &= \neg_i \mathbb{I}(\xi) \\ \mathbb{I}(abs_i(\xi)) &= \lceil_i \mathbb{I}(\xi). \end{aligned}$$

As anticipated, atomic contexts are interpreted as sets of models on some language \mathcal{L}_i ; the \perp_i context is interpreted as the empty context (the same on each language); the \top_i context is interpreted as the greatest, or most general, context on \mathcal{L}_i ; the binary \wedge_i -composition of contexts is interpreted as the greatest lower bound of the restriction of the interpretations of the two contexts on \mathcal{L}_i ; the binary \vee_i -composition of contexts is interpreted as the lowest upper bound of the restriction of the interpretations of the two contexts on \mathcal{L}_i ; context negation is interpreted as the complement with respect to the most general context on that language; finally, the unary abs_i operator is interpreted just as the restriction of the interpretation of its argument to language \mathcal{L}_i .

Semantics for the assertions \mathcal{A} and for the contextual concept description \mathcal{D} in \mathcal{L}^{CT} is based on the function \mathbb{I} .

Definition 5. (Semantics of assertions: \models)

The semantics of assertions is defined as follows:

$$\begin{aligned} \mathbb{M} \models \xi : \gamma_1 \sqsubseteq \gamma_2 &\text{ iff } \forall m \in \mathbb{I}(\xi) \mathcal{I}_m(\gamma_1) \subseteq \mathcal{I}_m(\gamma_2) \\ \mathbb{M} \models \xi : \rho_1 \sqsubseteq \rho_2 &\text{ iff } \forall m \in \mathbb{I}(\xi) \mathcal{I}_m(\rho_1) \subseteq \mathcal{I}_m(\rho_2) \\ \mathbb{M} \models \xi_1 \preceq \xi_2 &\text{ iff } \mathbb{I}(\xi_1) \subseteq \mathbb{I}(\xi_2) \\ \mathbb{M} \models \sim \alpha &\text{ iff } \text{not } \mathbb{M} \models \alpha \\ \mathbb{M} \models \alpha_1 \wedge \alpha_2 &\text{ iff } \mathbb{M} \models \alpha_1 \text{ and } \mathbb{M} \models \alpha_2. \end{aligned}$$

A contextual concept subsumption relation between γ_1 and γ_2 holds iff all the basic description logic models constituting that context interpret γ_1 as a subconcept of γ_2 . Note that this is precisely the clause for the validity of a subsumption relation in standard description logics. Perfectly analogous observation holds also for the clause regarding contextual role subsumption relations. The \preceq relation between context constructs is interpreted as a standard subset relation: $\xi_1 \preceq \xi_2$ means that context denoted by ξ_1 contains at most all the models that ξ_2 contains, that is to say, ξ_1 is *at most as general as* ξ_2 . Note that this relation, being interpreted on the \subseteq relation, is reflexive, symmetric and transitive. In (Grossi & Dignum 2004) a generality ordering with similar properties was imposed on the set of context identifiers, and analogous properties for a similar relation have been singled out also in (Goldman 1976). The interesting thing is that such an ordering is here emergent from the semantics. Note also that this relation holds only between contexts specified on the same language. Clauses for boolean connectives are the obvious ones.

Contextual Terminologies at Work

Formalizing an example

We are now able to provide a formalization of a fragment of the scenario presented in Section ?? making use of the formal semantic machinery just exposed.

Example. (Personal data in transplant organizations and police registers) *To formalize a fragment of the scenario within our setting a language \mathcal{L} is needed, which contains the following atomic concepts: `personal_data`, `relevant_data`, `allowed_data`, `blood_type`, `race`, `anthropometric_properties`; and the following atomic role: `refer`. From this language we obtain $2^6 - 1 \cdot 2$ languages \mathcal{L}_i (see Section). Three atomic contexts are at issue here: the context of the superordinate regulation, let us call it c_{SUP} ; the contexts of the municipal regulations ONT and PR , let us call them c_{ONT} , c_{PR} and c_{M3} respectively. These contexts should be interpreted on two relevant languages (let us call them \mathcal{L}_0 and \mathcal{L}_1) s.t.:*

$$\mathbf{A}_0 = \{\text{personal_data, relevant_data, allowed_data}\},$$

$$\mathbf{R}_0 = \emptyset$$

and

$$\mathbf{A}_1 = \{\text{personal_data, relevant_data, allowed_data, blood_type, race, anthropometric_properties}\},$$

$$\mathbf{R}_1 = \{\text{refer}\}.$$

That is to say, an abstract language concerning only personal, relevant and allowed data, and a more detailed language concerning, besides personal, relevant and allowed data, also blood type, race, anthropometric properties and the refer role. The sets of all models for \mathcal{L}_0 and \mathcal{L}_1 are then respectively \mathbf{M}_0 and \mathbf{M}_1 .

To model the desired situation, our ct-model should then at least satisfy the \mathcal{L}^{CT} formulas listed in figure 1.

$$\begin{aligned}
abs_0(c_{ONT}) \gamma_0 abs_0(c_{PR}) &\preceq c_{SUP} & (5) \\
c_{SUP} : \text{personal_data} \sqcap \text{relevant_data} &\sqsubseteq \text{allowed_data} & (6) \\
c_{ONT} \gamma_1 c_{PR} : \text{personal_data} \sqcap \exists \text{refer.blood_type} &\sqsubseteq \text{relevant_data} & (7) \\
c_{ONT} \gamma_1 c_{PR} : \text{personal_data} \sqcap \exists \text{refer.ant hr opom et r ic_pr opert i es} &\sqsubseteq \text{relevant_data} & (8) \\
c_{ONT} \gamma_1 c_{PR} : \text{personal_data} \sqcap \exists \text{refer.race} &\sqsubseteq \text{personal_data} \sqcap & \\
&\sqsubseteq \exists \text{refer.ant hr opom et r ic_pr opert i es} \sqcup \neg \text{relevant_data} & (9) \\
c_{ONT} : \text{race} &\sqsubseteq \neg \text{ant hr opom et r ic_pr opert i es} & (10) \\
c_{PR} : \text{race} &\sqsubseteq \text{ant hr opom et r ic_pr opert i es} . & (11)
\end{aligned}$$

Figure 1: \mathcal{L}^{CT} formalization of the scenario

Formula (5) plays a key role, stating that the two contexts c_{ONT} , c_{PR} are concrete variants of context c_{SUP} . It tells this by saying that the context obtained by joining the two concrete contexts on language \mathcal{L}_0 (the language of c_{SUP}) is at most as general as context c_{SUP} . As we will see in the following section, this makes c_{ONT} , c_{PR} inherit what holds in c_{SUP} . Formula (6) formalizes the abstract rule to the effect that personal data which are relevant for the accomplishment of the aim of the organization are allowed to be recorded and used. Formulas (7) and (8) express subsumptions holding in both contexts. Formula (9) tells something interesting, namely that data about race, in order to be used, they have to be considered as anthropomorphic information. Indeed, it might be seen as a clause avoiding “cheating” classifications such as: “data about race count as data about blood type”. Finally, formulas (10) and (11) describe about what precisely the taxonomies holding in the two contexts diverge.

Discussing the formalization

To discuss in some more depth the proposed formalization, let us first list some interesting logical consequences of formulas (5)-(11) in figure 2. We will focus on subsumptions contextualized to monadic contexts, that is to say, we will show what the consequences of formulas (5)-(11) are at the level of the two contexts c_{ONT} , c_{PR} . These are indeed the formulas that we would intuitively expect to hold in our scenario. The list displays two sets of formulas grouped on the basis of the context to which they pertain. Let us have a closer look to them. The first consequence of each group results from the generality relation expressed in (5), by means of which, the content of (6) is shown to hold also in the two concrete contexts: in simple words, contexts c_{ONT} , c_{PR} inherit the general rule stating that only relevant personal data can be included and used. Via this inherited rule, and via (7) and (8), it is shown that, in all contexts, data about blood type and anthropometric properties are always allowed. As to data about blood type and anthropometric properties, all contexts agree. Differences arise in relation with how the concept of race is handled.

In context c_{ONT} , we have that data about race should not be taken as relevant, and this conclusion is reached restricting the interpretation of what counts as an anthropometric information (10) and by means of the “no-cheating” clause

(9). In fact, in this context, data about race are not anthropometric data. Context c_{PR} , instead, expresses a different view. Since race counts as anthropometric information (11), data about race are actually relevant data and, as such, they can be used.

Before ending the section, we deem worth putting this context-based approach in perspective with the more standard ones based instead on the defeasible reasoning paradigm. In a non-monotonic reasoning setting, the key point of the example (the fact that the two contexts diverge in the classification of the concept *race*) would be handled by means of a notion of exception: “normally, race is an anthropometric property and is then an allowed type of personal data” and “every exceptional anthropometric property is a forbidden type of personal data”. We deem these approaches, despite being effective in capturing the reasoning patterns involved in this type of scenarios, to be inadequate for analyzing problems related with the *meaning* of the terms that trigger those reasoning patterns. Those reasoning patterns are defeasible because the meaning of the terms involved is not definite, it is vague, it is -and this is the thesis we hold here- context dependent⁶. Our proposal consists instead in analysing these issues in terms of the notion of context: according to (in the context of) *PR* race is an anthropometric property; according to (in the context of) *ONT* race does not count as an anthropometric property. Besides enabling the possibility of representing semantic discrepancies, such an approach has also the definite advantage of keeping the intra-contextual reasoning classical, framing non-monotonicity as emergent property at the level of inter-contextual reasoning. Furthermore, the use of description logic allows for its well known interesting computability properties to be enabled at the intra-contextual reasoning level, thus making the framework appealing in this respect as well.

Implementing Norms using Contextual Ontologies

In the previous sections we have given an idea of how ontologies and context are used in institutions in order to determine whether or not norms apply to a given situation. We

⁶The issue of the relationship between contextuality and defeasibility has been raised also in (Akman & Surav. 1996).

$$\begin{aligned}
(5), (6) &\models c_{ONT} : \text{personal_data} \sqcap \text{relevant_data} \sqsubseteq \text{allowed_data} \\
(5), (6), (7) &\models c_{ONT} : \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{relevant_data} \\
(5), (6), (7) &\models c_{ONT} : \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{allowed_data} \\
(5), (6), (8) &\models c_{ONT} : \text{personal_data} \sqcap \exists \text{refer.ant hr opom et r i c_pr opert i es} \sqsubseteq \text{relevant_data} \\
(5), (6), (8) &\models c_{ONT} : \text{personal_data} \sqcap \exists \text{refer.ant hr opom et r i c_pr opert i es} \sqsubseteq \text{allowed_data} \\
(8), (10) &\models c_{ONT} : \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \text{personal_data} \\
&\quad \sqcap \exists \text{refer.}\neg \text{ant hr opom et r i c_pr opert i es} \\
(5), (6), (9), (10) &\models c_{ONT} : \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \neg \text{relevant_data} \\
\\
(5), (6) &\models c_{PR} : \text{personal_data} \sqcap \text{relevant_data} \sqsubseteq \text{allowed_data} \\
(5), (6), (7) &\models c_{PR} : \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{relevant_data} \\
(5), (6), (7) &\models c_{PR} : \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{allowed_data} \\
(5), (6), (8) &\models c_{PR} : \text{personal_data} \sqcap \exists \text{refer.ant hr opom et r i c_pr opert i es} \sqsubseteq \text{relevant_data} \\
(5), (6), (8) &\models c_{PR} : \text{personal_data} \sqcap \exists \text{refer.ant hr opom et r i c_pr opert i es} \sqsubseteq \text{allowed_data} \\
(8), (11) &\models c_{PR} : \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \text{personal_data} \\
&\quad \sqcap \exists \text{refer.ant hr opom et r i c_pr opert i es} \\
(5), (6), (9), (11) &\models c_{PR} : \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \text{relevant_data}
\end{aligned}$$

Figure 2: Logical consequences of formulas (5)-(11)

have given a formal framework to formalize the contexts and have shown how this framework can be used to represent and reason about norms in an eInstitution. Although an implementation covering all the aspects of the the formal machinery proposed in the previous sections would be computationally expensive, an optimal implementation of the ontological aspects of norms can be far less complex.

It is important to note here that implementing the contextual ontological aspects does not mean implementing some sort of model-checker to verify the formal models of the norms and situations described in our formal framework (or non-monotonic logical descriptions of those), since one is only going to encounter a limited number of contexts at a given time. From the institutional perspective, as we can consider an eInstitution as a single context, all contextual ontological issues are solved during the design process of the eInstitution when defining its ontology. From the agents' perspective, the contextual ontological problems should be solved on-line; agents that are joining the eInstitution need to know in which context they are supposed to work, and need to be informed of the ontology and norms applicable in the eInstitution.

From the eInstitution's view, the ontological aspects of norms mainly impact two steps in the eInstitution's implementation: a) the definition of the *eInstitution's ontology*, giving an interpretation of all the terms in the norms, and b) the implementation of the *norm enforcement mechanisms*, following the norm interpretation given by the ontology.⁷ From the ontological perspective, the most complex step is the definition of its ontology, as sev-

⁷More details on the implementation of norm enforcement mechanisms can be found in (?).

eral contextual ontologies should be taken into account. That is, not only does one need to look at the concepts and norms necessary for the eInstitution's context, but one also has to consider the (super)contexts in which the eInstitution is to operate, which are possibly nested (e.g. regional/national/transnational/international contexts). In practice this means, to create some kind of link the ontologies of different supercontexts to the institutional ontology. In our approach (which is an ongoing work), the links between ontologies are explicitly defined by the designer by means of different kinds of ontology abstraction and ontology inheritance relations. The simplest scenario is when an eInstitution has a set of non-conflicting nested supercontexts. For instance, in the case of an eInstitution for the Spanish National Transplant Organization (ONT), in order to define ONT's ontology we can inherit terms from its supercontexts: The Spanish National Health System, the Spanish Law and the European Union Law. It is important to note that an explicit link for all inherited terms should be kept in the ontologies' representation. Then the inherited terms can be extended in ONT's ontology with extra terms and/or re-defined, if needed, for the particular context of the institution. A more complex scenario appears when an eInstitution has disjoint nested supercontexts with conflicting definitions of terms. This is the case of trans-national institutions such as Eurotransplant⁸, where different ontological definitions of terms may appear in each of the countries where the institution should operate. In this case, when inheriting different,

⁸The Eurotransplant International Foundation (?) is responsible for the mediation and allocation of organ donation procedures in Austria, Belgium, Germany, Luxemburg, the Netherlands and Slovenia.

conflicting definitions of the same term into the ontology, the designers should solve the conflict by precisely agreeing on and defining the precise meaning of the term that will apply inside the context of the eInstitution.

This procedure can be very useful to obtain a common ontology for an agent that tries to enter the eInstitution. If an agent tries to enter an eInstitution, and it becomes clear that there are concepts in the ontology of the eInstitution that the agent does not understand, the links to the ontologies of the (super)contexts, which were made when the eInstitution was implemented, can be used in order to get a shared understanding of the concept in question.

Discussion, Conclusions & Future Work

The motivating question of our research was: how do institutions make their norms operative in the domain they are supposed to regulate, i.e., how do institutions implement norms? The thesis we held here is that institutions are based on ontologies. Via these ontologies they translate norms, which are usually formulated in abstract terms (for instance, the concept of “relevant data”), into concrete constraints which are instead understandable in the terms used to describe the situations they regulate (for example, “data about blood type”).

References

- Akman, V., and Surav., M. 1996. Steps toward formalizing context. *AI Magazine* 17(3):55–72.
- Alchourrón, C. E., and Bulygin, E. 1986. *Normative Systems*. Wien: Springer Verlag.
- Baader, F.; Calvanese, D.; McGuinness, D.; Nardi, D.; and Patel-Schneider, P. 2002. *The Description Logic Handbook*. Cambridge: Cambridge University Press.
- Boella, G., and Van der Torre, L. 2004. Regulative and constitutive norms in normative multiagent systems. In *Proceedings of KR2004, Whistler, Canada*, 255–266.
- Bulygin, E. 1992. On norms of competence. *Law and Philosophy* 11 201–216.
- Casari, E. 2002. *La Matematica della Verità*. Privately distributed.
- Dignum, F. 2002. Abstract norms and electronic institutions. In *Proceedings of the International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA '02), Bologna*, 93–104.
- Goldman, A. I. 1976. *A Theory of Human Action*. Princeton: Princeton University Press.
- Grossi, D., and Dignum, F. 2004. From abstract to concrete norms in agent institutions. In *Proceedings of FAABS III Workshop, Washington, april*.
- Grossi, D.; Dignum, F.; and Meyer, J.-J. C. 2004a. Contextual taxonomies. In leite, J., and Toroni, P., eds., *Proceedings of CLIMA V Workshop, Lisbon, September*, 2–17.
- Grossi, D.; Dignum, F.; and Meyer, J.-J. C. 2004b. Contextual terminologies. Draft.
- Hart, H. L. A. 1961. *The Concept of Law*. Oxford: Clarendon Press.

Jones, A. J. I., and Sergot, M. 1992. Deontic logic in the representation of law: towards a methodology. *Artificial Intelligence and Law* 1.

Jones, A. J. I., and Sergot, M. 1993. On the characterization of law and computer systems. *Deontic Logic in Computer Science* 275–307.

McCarthy, J. 1986. Notes on formalizing contexts. In Kehler, T., and Rosenschein, S., eds., *Proceedings of the Fifth National Conference on Artificial Intelligence*, 555–560. Los Altos, California: Morgan Kaufmann.

Peczenik, A. 1989. *On Law and Reason*. Dordrecht: Kluwer.

Ross, A. 1968. *Directives and Norms*. London: Routledge & Kegan Paul.

Searle, J. 1995. *The Construction of Social Reality*. Free Press.

Von Wright, G. H. 1963. *Norm and Action. A Logical Inquiry*. London: Routledge.