Comput Math Organiz Theor (2006) 12:251–275
DOI 10.1007/s10588-006-9546-6

Ontological aspects of the implementation of norms in agent-based electronic institutions

**Davide Grossi · Huib Aldewereld ·
Javier Vázquez-Salceda · Frank Dignum**

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 become instead relevant for the actual functioning of the 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 analyse this phenomenon introducing a notion of contextual ontology. We will focus on the formal machinery necessary to characterise it as well.

Keywords Agents · Norms · Context · Ontology · Institution

D. Grossi (✉) · H. Aldewereld · F. Dignum
Institute of Information and Computing Sciences, Utrecht University,
P.O.Box 80.089, 3508 TB Utrecht, The Netherlands
e-mail: davide@cs.uu.nl

H. Aldewereld
e-mail: huib@cs.uu.nl

F. Dignum
e-mail: dignum@cs.uu.nl

J. Vázquez-Salceda
Knowledge Engineering and Machine Learning Group, Universitat Politècnica de Catalunya,
Jordi Girona 1–3, Barcelona, Spain
e-mail: jvazquez@lsi.upc.edu

1. 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.[1] This means that norms are expressed in terms of concepts that are, on purpose, kept vague and ambiguous.[2] 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 operationalisation of the institution, should be answered in a clear and definite way.

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 as sketched in 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 previous work we focused on declarative aspects of norms, Dignum, Kinny and Sonenberg (2002) and Dignum et al. (2002), formally defining norms by means of some variations of deontic logic that include conditional and temporal aspects, in Broersen et al. (2004) and Dignum et al. (2004). We have also explored some of the operational aspects of norms, by focusing on how norms should be operationally implemented in MAS from an institutional perspective in Vázquez-Salceda, Aldewereld and Dignum (2004). In this paper we extend this line of research, taking into account the ontological aspects of norm implementation.

This work is organised 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 Section 3, we will present a formal framework in which it is possible to represent and reason about divergent ontologies. Using this framework we will formalise and discuss an example in Section 4. In Section 5 we will address the issue of the representation of norms, together with ontologies, in the framework proposed. After this, in Section 6, we will discuss the issue of the practical implementation of the framework in eInstitutions. We end the paper recapitulating our theses and drawing some conclusions in Section 7.

---

[1] See for instance Dignum (2002) and Grossi and Dignum (2004).

[2] Cf. (Hart, 1961).

Throughout the all paper, we will constantly refer to related work especially in the area of formal representation of context (Section 3) and of deontic logic (Section 5). As main example we will use the regulations on personal data protection in several scenarios: the European Union, the Dutch Police, European Hospitals and the Spanish National Transplants Organisation (an organisation for the allocation of human organs and tissues for transplantation purposes).

2. Institutions, ontologies and contexts

In order to properly implement norms in eInstitutions, we should first analyse how norms are handled 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 vice versa. 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 privacy policies. 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 emphasises 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, [3] 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 legal theory, the non-regulative component of the issuing of norms has been labelled in ways that emphasise 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 and 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 therefore been called also *constitutive rules* or *constitutive norms*, cf. Ross (1968) and Searle (1995). In this view, statements to the effect that racial data count as personal data establish that being racial data constitutes, 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*. This feature has been particularly emphasised in Searle (1995), where constitutive rules are viewed as representable via the following type of statements: "X counts as Y *in context* C".

---

[3] See Alchourrón and Bulygin (1986); Jones and Sergot (1992, 1993); Searle (1995); Boella and Van der Torre (2004).

2.1. 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 influences 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, cf. (Vázquez-Salceda, 2004), as the environment may affect its specification (especially 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 an 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 eInstituions have to inter-operate. The source of these problems is that, in most 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 transnational, 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 (1986, 1987), is to include *contexts* as formal objects in the model. Therefore, most theoretical approaches have moved towards having an explicit representation of context.

In Vázquez-Salceda (2004), 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 eInstitution operates are also (super)contexts, being possibly nested (e.g. to model the nesting in regional/national/transnational environments).

2.2. Contextual ontologies

Each normative context should therefore define a vocabulary to be shared by agents in that context. It means that each context is associated with a domain ontology that defines the meaning of the terms that are present in the norms, the actions the agent may perform 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 made between ontologies of inter-related contexts. This problem usually appears in multiagent 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[4] *personal data* are defined as *"[...] those [data] which allow 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"*. This abstract definition of the term *personal data* has been introduced, in more or less extent, in the regulations of the EU member states. EU regulations on personal data protection apply to every data archives structured in a way which allows the easy extraction of personal information, including electronic archives on any computer-readable storage device and format. One important aspect is the rights that EU citizens have over their personal data.

As a general rule, personal data collection and processing requires first of all the approval of the affected person. Each person has then the right to access, amend, cancel or be opposed to the collection of her personal data. Personal data will finally only be kept during the period needed to achieve the aims they were collected for, or the authorised extensions of those aims. If it is desirable to maintain this information long after this period (for historical, statistical or scientific purposes), it must be done in a way that avoids personal identifications.

In practice, this means that any institution within the European Union context should only store a subset of the personal data, the *relevant data*, that is needed for the purposes of the data collection. The definition of relevance is highly contextual, depending on the activity of the institution and/or the purpose of the archive. Therefore, different institutions will have different definitions of *relevant data*: for instance, relevant data about patients in a hospital clearly should include name, address and any medical information details that are important for the patient's treatment, while relevant data that some companies (e.g. shopping centres) keep about their clients may include name, address and a history of items the client used to buy (e.g. to adapt stocks and avoid item shortage), but not medical information, as it is not relevant for that company. There are some special cases where data, although being possibly relevant for a given institution, is not allowed to be stored. For instance, companies would find useful to have full access to the medical records of their employees, in order to ensure the productivity of its staff by reducing the risk of long-period illnesses. Although in this scenario medical information is relevant, in some European countries that information is not allowed at all or it is only allowed in some specific situations (e.g. with the explicit agreement of the person). In order to ensure personal data protection, all organisations that store and/or process personal data should get a certificate given by a National data protection agency. In such a document there are very specific definitions of which are the data officially being relevant for that particular institution. This provides a definition of what *operable data* are for that institution and once this is defined, all

---

[4] European Parliament created the 95/46/CE Directive (Directive ed9546, 1995) with the purpose of homogenising 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 (Regulation er45, 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 control authority.

regulations on data protection reduce to a single rule: organisations can only store and process the *operable data*.

Although any EU citizen has the right to access and check the information that any institution has about himself, this is highly impractical. Let us suppose that in the near future any organisation has an agent-mediated eInstitution to provide information and services to individuals and that any person can have an automated personal agent that keeps track of all personal information that organisations have about the person. This agent would enter in each eInstitution checking, for each case, that only *operable data* is stored, and eventually requesting for amendments or deletions of the information. Such an agent should adapt to the normative and ontological differences between contexts: although the agent may have an ontology defining what *personal data* is, it should be able to adapt its reasoning processes to the regulations and ontologies applying in a given, specific context. For instance, let us focus on two bits of personal information that are protected in the context of European regulations: a person's *blood type* and the person's *race* (Caucasian, Native American, Mongolian, Ethiopian and so on).

In the generic, *European Union context*, both *blood type* and *race* are *personal data* of a special nature that, in principle, are not *operable data*, unless some specific regulation or a certificate by a National data protection agency allows the storage and treatment of such information for some specific, well-defined purposes:

> "Member States shall prohibit the processing of personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, and the processing of data concerning health or sex life." Article 8.1 in Directive ed9546 (1995).

In the context of any *EU Police Force*, there are special allowances on the use of personal data:

> "Processing of data relating to offences, criminal convictions or security measures may be carried out only under the control of official authority [...]." Article 8.5 in Directive ed9546 (1995).

That means that, outside the context of an official institution, personal information about criminal antecedents of an individual (a criminal record) is completely forbidden, while in the context of, e.g., the Dutch Police, any relevant information about a criminal or a suspect of a crime (name, address, a physical description -including race- or even medical information such as blood type) is *operable data*.

In the context of any *EU Health Service*, there are also special allowances on the use of medical data:

> "Article 8.1 shall not apply where processing of the data is required for the purposes of preventive medicine, medical diagnosis, the provision of care or treatment or the management of health-care services [...]" Article 8.3 in Directive ed9546 (1995).

Therefore, inside the context of a hospital information such as *name*, *address* or *blood type* are usually relevant, and belongs to the set of *operable data* in medical records. On the other hand, *race* is rarely relevant and can only be included in the medical records in those illnesses that are highly related to race.

The context of the Spanish National Transplants Organisation[5] (ONT) is an interesting, specific subcontext of *EU Health Service*. By Spanish Law, ONT must ensure an equitative and fair distribution of organs and tissues by only taking into account clinical and geographical criteria. Therefore, clinical data such as *blood type* are allowed. Physical descriptions of the donor recipient (basically size, age and weight) are also allowed when they are relevant for the allocation. But such anthropometric data can never include *race*, as it is explicitly forbidden for ONT to use racial information during the allocation process. In this complex, multi-contextual scenario, a personal agent checking the use of a person's data on each of those contexts does not need to have a full model of all the regulations that apply in a given context. Reasoning on data allowance can be done in an ontological level, that is, the agent should adapt its reasoning to the ontological definitions of *relevant data* and *operable data* that holds in each context. Some kind of formal model for multi-contextual ontologies is needed, though, in order to properly model the relations between terms in different contexts.

3. Modelling contextual ontologies

This section is devoted to the exposition of a language and a semantics for talking about contextual ontologies. It widely draws on results we exposed in Grossi, Dignum and Meyer (2005a–2005c). More in detail, we aim at devising a formal morphology and a formal semantics meeting the following requirements.

Firstly, it should support reasoning about the validity of TBoxes with respect to contexts giving a semantics to expressions of the type: "the concept `race` is a subconcept of the concept `relevant_data` in the context of ONT". Secondly, it should provide a representation of context interplay. In particular, we will introduce: a *contextual disjunction* operator and a *contextual focus* operator.[6]

The first one yields a union of contexts: the contexts "viruses" and "bacteria" can be unified on a language talking about microorganisms generating a more general context like "viral or bacterial microorganisms". The second one, which plays a central role in our framework, yields the context consisting of some information extracted from the context on which it is focused: the context categorizing "crocodiles", for instance, can be obtained via focusing the context which categorizes all reptiles on the language talking only about crocodiles and disregarding other reptiles. In other words, the operator prunes the information contained in the context "reptiles" focusing only on what is expressible in the language which talks about crocodiles and abstracting from the rest. Also *maximum* and *minimum* contexts will be introduced: these will represent the most general, and respectively the most specific, contexts on a language.[7] It is

---

[5] The Organización Nacional de Transplantes is a technical organisation within the Spanish Department of Health and Consumer Affairs, whose fundamental mission is the promotion, facilitation and coordination of all types of organs, tissues and bone marrow.

[6] In Grossi, Dignum and Meyer (2005a, 2005b) the *focus* operation is called *abstraction*. We decided to modify our terminology in order to avoid confusions with other approaches to notions of abstraction like for instance (Ghidini and Giunchiglia, 2004).

[7] In this paper, we limit the number of context operations to disjunction and focus. More operations are formalized in Grossi, Dignum, and Meyer (2005b). It is worth noticing, in passing, that similar operations and special contexts are discussed in Shoham (1991).

important to notice that all operations explicitly refer to a precise language on which the operation should take place. As we will see in the following section our formal language will be tuned to incorporate this feature.

Finally, our language should represent a generality relation between contexts[8] expressing that a context is at most as general as another one: the context of the European regulation on personal data protection is somehow more general than the concrete regulations governing personal data processing within ONT and the Dutch police force (PF).[9]

Context relations and operations provide the means for a formal description of how different contexts are related and interact. These intuitions about the semantics of context operators will be clarified and made more rigorous in Section 3.2 where the semantics of the framework will be presented, and in Section 4.1 where the example will be formalized deploying all these types of expressions.

In a nutshell, our approach consists in mixing the semantics of description logic with the idea of modeling contexts as sets of models. Description logic is a well-known logical formalism intended for representing knowledge about hierarchies of concepts endowed with a Tarskian semantics and effective reasoning procedures.[10] The idea of modeling contexts as sets of models is borrowed from work developed in the area of formal context representation, in particular from the framework developed in Ghidini and Giunchiglia (2001), and tailored to the description logic paradigm.

The framework delivered is able to represent reasoning about sets of concept subsumptions, i.e., what in description logics are called taxonomical boxes (TBoxes), in a contextual setting, and to characterise forms of contextual interplay like the aforementioned contextual focus and the generality relation between contexts.

3.1. Language

The language we are defining can be seen as a meta-language for TBoxes defined on $\mathcal{ALC}$ description logic languages, i.e., concept languages enabling arbitrary concept negation, union and intersection, and full universal and existential quantification.

The alphabet of the language $\mathcal{L}^{CT}$ (*language for contextual terminologies*) contains therefore the alphabets of a family of languages $\{\mathcal{L}_i\}_{0\leq i\leq n}$. This family is built on the alphabet of a given "global" language $\mathcal{L}$ which contains all the terms occurring in the elements of the family. Moreover, we take $\{\mathcal{L}_i\}_{0\leq i\leq n}$ to be such that, for each non-empty subset of terms of the language $\mathcal{L}$, there exist a $\mathcal{L}_i$ which is built on that set and belongs to the family. Each $\mathcal{L}_i$ contains two non-empty finite sets $\mathbf{A_i}$ of atomic concepts (A), i.e., monadic predicates, and $\mathbf{R_i}$ of atomic roles (R), i.e.,

---

[8] Literature on context theory often addresses this type of relation between contexts. See for instance (McCarthy, 1986; Benerecetti, Bouquet and Ghidini, 2000).

[9] As the discussion of the formalization of the examples will show (Section 4.1), there are some more subtleties to be considered since the context of the European regulation is not only more general but is also specified on a simpler language.

[10] For space reasons we cannot introduce here description logic, although we present its semantics in Section 3.2. For a detailed exposition of description logic we then refer the reader to Baader et al. (2002).

dyadic predicates.[11] These languages contain also concepts constructors, that is, each $\mathcal{L}_i$ contains the zeroary operators $\bot$ (bottom concept) and $\top$ (top concept), the unary operator $\neg$ (complement), and the binary operator $\sqcap$ (conjunction). Finally, the value restriction operator $\forall R.A$ ("the set of elements such that all elements that are in a relation R with them are instances of A") applies to role-concept pairs.

Besides, the alphabet of $\mathcal{L}^{CT}$ contains a finite set of context identifiers **c**, two families of zeroary operators $\{\bot_i\}_{0\leq i\leq n}$ (minimum contexts) and $\{\top_i\}_{0\leq i\leq n}$ (maximum contexts), one family of unary operators $\{\mathsf{fcs}\ _i\}_{0\leq i\leq n}$ (contextual focus operator) , one family of binary operators $\{\curlyvee_i\}_{0\leq i\leq n}$ (contexts disjunction operator), one context relation symbol $\preccurlyeq$ expressing the generality relation between contexts (context c_1 "is at most as general as" context c_2), and finally a contextual subsumption relation symbol " $. : . \sqsubseteq .$" (within context c, concept A_1 is a subconcept of concept A_2) for concept subsumption. Lastly, the alphabet of $\mathcal{L}^{CT}$ contains also the sentential connectives $\sim$ (negation) and $\wedge$ (conjunction).[12]

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

$$\xi ::= c \mid \bot_i \mid \top_i \mid \mathsf{fcs}\ _i(\xi) \mid \xi_1 \curlyvee_i \xi_2.$$

The language contains only atomic roles. The set of atomic roles is denoted by P. As to concept constructs, they are defined in the standard way. The set Γ of concept descriptions (γ) is defined through the following BNF:

$$\gamma ::= A \mid \bot \mid \top \mid \neg\gamma \mid \gamma_1 \sqcap \gamma_2 \mid \forall R.\gamma.$$

Concept union and existential quantification are defined respectively as:

$$\gamma_1 \sqcup \gamma_2 =_{\text{def}} \neg(\neg\gamma_1 \sqcap \neg\gamma_2) \quad \text{and} \quad \exists R.\gamma =_{\text{def}} \neg(\forall R.\neg\gamma).$$

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

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

Technically, a *contextual terminology* in $\mathcal{L}^{CT}$ is a set of subsumption relation expressions on concepts which are contextualised with respect to the same context. This kind of expressions are, in a nutshell, what we are mainly interested in formalising. It is worth noticing that contextual subsumption relations can be viewed as a formal characterisation of statements of the type "X counts as Y in context C". This type of statements is considered in (Searle, 1995) to be essential for the analysis of social and institutional realities.[13]

Throughout the paper the following symbols will be also used " $. : . \sqsubset .$" (within context c, concept A_1 is a proper subconcept of concept A_2), and " $. : . \equiv .$"

---

[11] We use here the term "role" in the technical sense in which it is understood in description logic, i.e., as "attribute" of a concept.

[12] It might be worth remarking that language $\mathcal{L}^{CT}$ is, then, an expansion of each $\mathcal{L}_i$ language. Notice also that all operators on contexts are indexed with the language on which the operation they denote takes place.

[13] We addressed this issue in depth in Grossi, Meyer and Dignum (2005d).

(within context c, concept A_1 is equivalent to concept A_2). They can be defined as follows:

$$\xi : \gamma_1 \sqsubset \gamma_2 =_{\text{def}} \xi : \gamma_1 \sqsubseteq \gamma_2 \wedge \sim \xi : \gamma_2 \sqsubseteq \gamma_1$$
$$\xi : \gamma_1 \equiv \gamma_2 =_{\text{def}} \xi : \gamma_1 \sqsubseteq \gamma_2 \wedge \xi : \gamma_2 \sqsubseteq \gamma_1.$$

3.2. Semantics

As exposed in the previous section, a $\mathcal{L}^{CT}$ consists of four classes of expressions: Ξ (context constructs), P (atomic roles), Γ (concept descriptions) and $\mathcal{A}$ (assertions). Semantics of P and Γ will be the standard description logic semantics of roles and concepts, on which our framework is based. Semantics for Ξ will be given in terms of model theoretic operations on sets of description logic models which constitute the formal characterization of the forms of context interplay we are dealing with. At that stage the semantics of assertions $\mathcal{A}$ will be defined via an appropriate satisfaction relation. The structures obtained, which we call *contextual terminology models* or ct-models, provide a formal semantics for $\mathcal{L}^{CT}$ languages.

The first step is then to provide the definition of a description logic 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:

$$\begin{aligned}
\mathcal{I}_m(\top) &= \Delta_m \\
\mathcal{I}_m(\bot) &= \emptyset \\
\mathcal{I}_m(\neg\gamma) &= \Delta_m \backslash \mathcal{I}_m(\gamma) \\
\mathcal{I}_m(\gamma_1 \sqcap \gamma_2) &= \mathcal{I}_m(\gamma_1) \cap \mathcal{I}_m(\gamma_2) \\
\mathcal{I}_m(\forall\rho.\gamma) &= \{a \in \Delta_m \mid \forall b, < a, b > \in I_m(\rho) \Rightarrow b \in I_m(\gamma)\}.
\end{aligned}$$

A model m for a language $\mathcal{L}_i$ assigns a denotation to each atomic concept and to each atomic role. For example $\mathcal{I}_m$(`personal_data`), i.e., the interpretation of the concept `personal_data` would be a set of individuals from Δ_m; the interpretation $\mathcal{I}_m$(refer) of the role refer would be a set of pairs of individuals from Δ_m, i.e., the pairs of individuals such as the first individual "refers" to the second one. Meaning is then accordingly given to each complex concept. For instance, $\mathcal{I}_m(\texttt{personal\_data} \sqcap \exists\text{refer}.\texttt{race})$ yields the set of individuals which are instances of the intersection of the denotation of the concept `personal_data` with the set of individuals referring to at least one other individual which instantiates the concept `race`, i.e., the denotation of $\mathcal{I}_m(\exists\text{refer}.\texttt{race})$.

3.3. 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:

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

where:

- $\{\mathbf{M_i}\}_{\mathbf{0 \leq i \leq n}}$ 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 description logic model of $\mathcal{L}_i$.
- $\mathbb{I}$ is a function $\mathbb{I} : \mathbf{c} \longrightarrow \mathcal{P}(\mathbf{M}_0) \cup \ldots \cup \mathcal{P}(\mathbf{M}_n)$. 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 modelling different (taxonomical) conceptualisations of a same set of individuals.

Contexts are therefore formalised 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.

3.4. Context focus

We model focus as a specific operation on sets of models which provides the semantic counterpart for the *contextual focus* operator introduced in $\mathcal{L}^{CT}$. Intuitively, abstracting a context ξ to a language $\mathcal{L}_i$ yields a context consisting in that part of ξ which can be expressed in $\mathcal{L}_i$.

Let us first recall a notion of *domain restriction* ($\rceil$) 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 $C\rceil f$ having C as domain and s.t. for each element of C, f and $C\rceil f$ return the same image: $C\rceil f = \{\langle x, f(x)\rangle \mid x \in C\}$.

Definition 3 (Context focus operation: $\rceil_i$). Let M' be a set of models, then: $\rceil_i M' = \{m \mid m = \langle \Delta_{m'}, \mathbf{A}_i \cup \mathbf{R}_i \rceil \mathcal{I}_{m'} \rangle \;\&\; m' \in M'\}$.

The following can be proved.

Proposition 1. *(Properties of context focus)*
Operation $\rceil_i$ *is: surjective, idempotent* ($\rceil_i(\rceil_i M) = \rceil_i M$)*, normal* ($\rceil_i \emptyset = \emptyset$)*, additive* ($\rceil_i(M_1 \cup M_2) = \rceil_i M_1 \cup \rceil_i M_2$)*, monotonic* ($M_1 \subseteq M_2 \Rightarrow \rceil_i M_1 \subseteq \rceil_i M_2$) .

Proof: A proof is worked out in (Grossi, Dignum and Meyer, 2005a). □

The operation of focus allows for shifting from richer to simpler languages and it is, as we would intuitively expect: surjective (every context, even the empty one, can be seen as the result of focusing a different richer context, in the most trivial case, a focus of itself), idempotent (focusing on a focus yields the same first focus), normal (focusing the empty context yields the empty context), additive (the focus of a context obtained via joining of two contexts can be obtained also joining the focuses of the two contexts), monotonic (if a context is less general then another one, the focus of the first is also less general than the focus of the second one). Notice also that operation $\rceil_i$ yields the empty set of models when it is applied to a context M' the language of which is not an expansion of $\mathcal{L}_i$. This is indeed very intuitive: the context obtained via focus of the context "dinosaurs" on the language of, say, "gourmet cuisine" should be empty.

It is instructive to notice that our notion of focus can be perfectly framed within the formal characterization of the notion of abstraction exposed in Ghidini and Giunchiglia (2004). It can therefore be regarded as a special case of the notion of abstraction as intended in that work. A detailed comparison of our account of focus with the aforementioned work is discussed in Grossi, Dignum and Meyer (2005a).

3.5. Operations on contexts

We are now in a position to give a semantics to context constructs as introduced in Section 3.1. In Definition 2 atomic contexts are interpreted as sets of models on some language $\mathcal{L}_i$ for $0 \leq i \leq n$: $\mathbb{I}(c) = M \in \mathcal{P}(\mathbf{M}_0) \cup \cdots \cup \mathcal{P}(\mathbf{M}_n)$. The semantics of context constructs Ξ can be defined via inductive extension of that definition.

Definition 4 (Semantics of context constructs). Let ξ, ξ_1, ξ_2 be context constructs, then:

$$\mathbb{I}(\mathsf{fcs}_i\ \xi) = \rceil_i \mathbb{I}(\xi)$$

$$\mathbb{I}(\bot_i) = \emptyset$$

$$\mathbb{I}(\top_i) = \mathbf{M}_i$$

$$\mathbb{I}(\xi_1 \curlyvee_i \xi_2) = \rceil_i(\mathbb{I}(\xi_1) \cup \mathbb{I}(\xi_2)).$$

The focus operator fcs_i is interpreted on the contextual focus operation introduced in Definition 3, i.e., as the restriction of the interpretation of its argument to language $\mathcal{L}_i$. The $\bot_i$ 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

$\curlyvee_i$-composition of contexts is interpreted as the lowest upper bound of the restriction of the interpretations of the two contexts on $\mathcal{L}_i$.[14]

3.6. Assertions

Semantics for the assertions $\mathcal{A}$ is based on the function $\mathbb{I}$. In what follows we denote with $\delta(\mathcal{I})$ the domain of an interpretation function $\mathcal{I}$.

Definition 5 (Semantics of assertions: $\models$). The semantics of assertions is defined as follows:

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

A contextual concept subsumption relation between γ_1 and γ_2 holds iff concepts γ_1 and γ_2 are defined in the models constituting context ξ, i.e., they receive a denotation in those models, and all the 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, but conditioned to the fact that the concepts involved are actually meaningful in that context. This further condition in the clause is necessary because our contexts have different languages. Intuitively, we interpret contextual subsumption relations as inherently presupposing the meaningfulness of their terms.[15] The $\preccurlyeq$ relation between context constructs is interpreted as a standard subset relation: $\xi_1 \preccurlyeq \xi_2$ means that the 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.

Notions of validity and logical consequence are classically defined. An assertion α is valid if all ct-models $\mathbb{M}$ satisfy it. An assertion α is a logical consequence of the set of assertions $\alpha_1, \ldots, \alpha_n$ if all ct-models satisfying $\alpha_1, \ldots, \alpha_n$ satisfy α.

4. Contextual ontologies at work

4.1. Formalising an example

We are now able to provide a formalisation of a fragment of the scenario presented in the first part of the paper, making use of the formal semantic machinery just exposed.

---

[14] It is worth noticing, in passing, that it can be proved that this semantics of context constructs forces contexts to be structured according to a special kind of Boolean Algebras. See (Grossi, Dignum and Meyer, 2005a) for details.

[15] For a more detailed discussion of these clauses we refer the reader to Grossi, Dignum and Meyer (2005a, 2005b).

Example 1 (Personal data in transplant organisations and police forces). We will formalise how the use of personal data is regulated in the two different contexts of Dutch police force (PF) and of the Spanish national transplant organisation (ONT) in accordance with the directives applying to the superordinate European context. We will see how the two concrete contexts PF and ONT implement the same European norm differently: personal data that are allowed to be operated by an institution are only those which are strictly relevant for the execution of the purpose of that institution. The two concrete contexts PF and ONT presuppose a different understanding of what counts as operable data, because their understanding of the norm lies in divergent ontologies of the concepts involved.

A language $\mathcal{L}$ is needed, which contains the following atomic concepts: `personal_data`, `relevant_data`, `operable_data`, `blood_type`, `race`, `anthropometric_properties`; and the following atomic role: refer. Three atomic contexts are at issue here: the context of the superordinate European regulation, let us call it c_{SUP}; the contexts of the subordinate regulations ONT and PF, let us call them c_{ONT} and c_{PF} respectively. These contexts should be interpreted on two relevant languages $\mathcal{L}_0$, i.e., the language of the context of European regulation, and $\mathcal{L}_1$, i.e., the language of the two concrete contexts PF and ONT. Languages $\mathcal{L}_0$ and $\mathcal{L}_1$ are such that:

$$\begin{aligned}
\mathbf{A}_0 &= \{\mathtt{personal\_data}, \mathtt{relevant\_data},\\
&\quad \mathtt{operable\_data}\},\\
\mathbf{R}_0 &= \emptyset\\
\mathbf{A}_1 &= \{\mathtt{personal\_data}, \mathtt{relevant\_data},\\
&\quad \mathtt{operable\_data}, \mathtt{blood\_type},\\
&\quad \mathtt{race}, \mathtt{anthropometric\_properties}\},\\
\mathbf{R}_1 &= \{\mathrm{refer}\}.
\end{aligned}$$

That is to say, an abstract language concerning only personal, relevant and operable data, and a more detailed language concerning, besides personal, relevant and operable data, also blood type, race, anthropometric properties and the *refer* role.

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

$$c_{ONT} \curlyvee_0 c_{PF} \preccurlyeq c_{SUP} \tag{1}$$

$$c_{SUP} : \mathtt{personal\_data} \sqcap \mathtt{relevant\_data} \equiv \mathtt{operable\_data} \tag{2}$$

$$c_{ONT} \curlyvee_1 c_{PF} : \mathtt{personal\_data} \sqcap \exists \mathrm{refer}.\mathtt{blood\_type} \sqsubset \mathtt{relevant\_data} \tag{3}$$

$$c_{ONT} \curlyvee_1 c_{PF} : \mathtt{personal\_data} \sqcap \exists \mathrm{refer}.\mathtt{anthropometric\_properties} \sqsubset \mathtt{relevant\_data} \tag{4}$$

$$\begin{aligned} c_{ONT} \curlyvee_1 c_{PF} : \mathtt{personal\_data} \sqcap \exists \mathrm{refer}.\mathtt{race} &\sqsubset \mathtt{personal\_data}\\ &\sqcap \exists \mathrm{refer}.\mathtt{anthropometric\_properties} \sqcup \neg \mathtt{relevant\_data} \end{aligned} \tag{5}$$

$$c_{ONT} : \mathtt{race} \sqsubseteq \neg \mathtt{anthropometric\_properties} \tag{6}$$

$$c_{PF} : \mathtt{race} \sqsubset \mathtt{anthropometric\_properties}. \tag{7}$$

Fig. 1 $\mathcal{L}^{CT}$ formalisation of the scenario

$$\begin{aligned}
(1),(2) &\models c_{ONT} : \texttt{personal\_data} \sqcap \texttt{relevant\_data} \sqsubseteq \texttt{operable\_data}\\
(1),(2),(3) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{blood\_type} \sqsubset \texttt{relevant\_data}\\
(1),(2),(3) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{blood\_type} \sqsubset \texttt{operable\_data}\\
(1),(2),(4) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{anthropometric\_properties} \sqsubset \texttt{relevant\_data}\\
(1),(2),(4) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{anthropometric\_properties} \sqsubset \texttt{operable\_data}\\
(4),(6) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{race} \sqsubset \texttt{personal\_data}\\
&\quad \sqcap \exists \mathrm{refer}.\neg\texttt{anthropometric\_properties}\\
(1),(2),(5),(6) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{race} \sqsubset \neg\texttt{relevant\_data}\\
(1),(2),(5),(6) &\models c_{ONT} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{race} \sqsubset \neg\texttt{operable\_data}
\end{aligned}$$

$$\begin{aligned}
(1),(2) &\models c_{PF} : \texttt{personal\_data} \sqcap \texttt{relevant\_data} \sqsubseteq \texttt{operable\_data}\\
(1),(2),(3) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{blood\_type} \sqsubset \texttt{relevant\_data}\\
(1),(2),(3) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{blood\_type} \sqsubset \texttt{operable\_data}\\
(1),(2),(4) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{anthropometric\_properties} \sqsubset \texttt{relevant\_data}\\
(1),(2),(4) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{anthropometric\_properties} \sqsubset \texttt{operable\_data}\\
(4),(7) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{race} \sqsubset \texttt{personal\_data}\\
&\quad \sqcap \exists \mathrm{refer}.\texttt{anthropometric\_properties}\\
(1),(2),(5),(7) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{race} \sqsubset \texttt{relevant\_data}\\
(1),(2),(5),(7) &\models c_{PF} : \texttt{personal\_data} \sqcap \exists \mathrm{refer}.\texttt{race} \sqsubset \texttt{operable\_data}
\end{aligned}$$

Fig. 2 Logical consequences of formulas (1)–(7)

Formula (1) plays a key role, stating that the two contexts c_{ONT}, c_{PF} 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_{PF} inherit what holds in c_{SUP}. Formulas (2)–(7) all express contextual subsumption relations. It is worth stressing that they can all be seen as formalising counts-as statements which specify the ontologies holding in the contexts at issue. Formula (2) formalises an abstract definition of what type of data can be recorded and used, a kind of definitional core for the notion of `operable_data`. Formulas (3) and (4) express subsumptions holding in both contexts. Formula (5) tells something interesting, namely that data about race, in order to be used, has to be considered as anthropometric information. Indeed, it might be seen as a clause avoiding "cheating" classifications such as: "data about race counts as data about blood type". Finally, formulas (6) and (7) describe how precisely the ontologies holding in the two contexts diverge.

4.2. Discussing the formalisation

To discuss in some more depth the proposed formalisation, let us first list some interesting logical consequences of formulas (1)–(7) in Fig. 2. We will focus on subsumptions contextualised to monadic contexts, that is to say, we will show what the consequences of formulas (1)–(7) are at the level of the two contexts c_{ONT}, c_{PF}. 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 at them; the first consequence of each group results from the generality relation expressed in (1), by means of which, the content of (2) is shown to hold also in the two concrete contexts: in simple words, contexts c_{ONT}, c_{PF} inherit the abstract definition stating that only relevant personal data can be included and used. This displays an important aspects of contextual ontologies in the institutional domain. Although contexts may have divergent ontologies concerning the same concepts, they can share, when subordinated to a same abstract context, a same definitional core for those concepts. In the example considered this is evident for the concept `operable_data`, but a definitional core could be stated by c_{SUP} also for other concepts like `personal_data`, etc. Via the inheritance of the abstract definition, and via (3) and (4), it is shown that, in all contexts, data about blood type and anthropometric properties are always operable. As to data about blood type and anthropometric properties, all contexts agree. Differences arise in relation with how the concept of race is handled. Those differences determine the variation in how the abstract definition expressed in (2) is interpreted.

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 anthropometric information (6) and by means of the "no-cheating" clause (5). In fact, in this context, data about race are not anthropometric data and consequently they are not operable. Context c_{PF}, instead, expresses a different view. Since race counts as anthropometric information (7), data about race are actually relevant data and, as such, can be operated.

Before ending the section, we briefly confront this context-based approach 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 operable 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 analysing 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 dependent.[16] Our proposal consists instead in analysing these issues in terms of the notion of context: according to (in the context of) PF 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.

5. Ontologies and norms

Up to now, the focus of the present paper has deliberately been on dealing exclusively with how concepts occurring in institutional regulations are structured to form the

---

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

terminologies which institutions use in order to conceptualise the domains they are supposed to regulate. In this section we show how the framework exposed in the previous sections enables sufficient expressivity to deal also with explicit representations of norms, i.e., of the regulative component of normative systems, this meaning that within such a framework it is possible to represent and give semantics to both the regulative and non-regulative aspects characterizing complex normative systems.[17]

The representation of norms in logic, and in general of deontic concepts such as obligations, rights, permissions, prohibitions, has a long history and gave rise to a plethora of considerably different approaches.[18]

To our knowledge, despite this long tradition, no attempt to represent deontic notions in a description logic setting has ever been proposed. Aim of this section is to discuss possible representations of basic deontic notions in such a logical setting in order to readily incorporate them in our framework for contextual ontologies. Two main straightforward options are introduced and discussed.

5.1. First option: a reduction strategy

The most straightforward way to deal with deontics in a $\mathcal{L}^{CT}$ language consists in developing a version of the reduction approach to deontic logic which was first proposed in Anderson (1957, 1958) within a modal logic setting.[19] The reduction strategy is based on the intuition according to which the fact that something is obligatory means that its negation "necessarily" implies a violation (of the relevant set of norms or deontic constraints). The nature of the reduction lies in how this reference to a "necessity" is formally modeled. Various alternative reductions are studied in d'Altan, Meyer and Wieringa (1993), Krabbendam and Meyer (2003), Lomuscio and Sergot (2003) and Grossi, Meyer and Dignum (2005d). Recalling the example discussed in Section 4, a reduction approach can be exemplified in our language as follows:

$$c_{\mathrm{ONT}} : \neg\mathtt{operable\_data} \sqcap \mathtt{operated\_data} \sqsubseteq \mathtt{violating}.$$

Intuitively: "all data which are not operable and have been actually operated count as something violating the relevant regulation in the context of ONT". More in general, the application of a reduction strategy to our framework for representing norms would lead to a new class of contextual subsumption statements defining within each context the ontology of a notion of violation, i.e., of the special concept `violating`. From this standpoint the regulative component of normative systems can be viewed as a specification of how each institution categorizes non-compliant objects. We explored this very same idea within a modal logic setting in Grossi, Meyer and Dignum (2005d).

Along these lines, a yet easier way to represent deontics in a $\mathcal{L}^{CT}$ language from a reduction perspective would be to consider sorts of "ideal" counterparts of contexts. Given a context ξ providing classification of the concept `violating`, its

---

[17] See Section 2.

[18] The first pioneering paper settling this field is considered to be (Wright, 1951). For an overview of the most recent developments in this area of logic see (Lomuscio and Nute, 2004).

[19] Modal logic has always been the most exploited tool for the representation of deontic notions. For a comprehensive introduction to modal logic we refer the reader to (Blackburn, de Rijke and Venema, 2001).

ideal counterpart ξ_{ideal} would consist in the set of models $m \in \mathbb{I}(\xi)$ and such that $\mathcal{I}_m(\texttt{violating}) = \emptyset$. That is to say:

$$\xi_{\text{ideal}} : \texttt{violating} \equiv \bot.$$

This condition states that in such contexts the violation corresponds to an inconsistency and therefore, whatever happens to be classified as yielding a violation is, in such contexts, a non existent entity. In other words, whatever holds in an ideal counterpart ξ_{ideal} can be consequently viewed as what ideally is the case. An approach investigating, in modal logic, this idea of obligation as truth in the ideal context is Lomuscio and Sergot (2003).

5.2. Second option: "deontic roles"

Another option would consist in exploiting well established results about the correspondence between description logic and modal logic (Schild, 1991; Blackburn, de Rijke and Venema, 2001; Baader et al., 2002). Doing this, we would be able to represent standard deontic logic (Wright, 1951), i.e. the standard system formalizing reasoning about "the distinction between what *ideally* is the case on the one hand, and what *actually* is the case on the other" (Jones and Sergot, 1992), in a description logic fashion.

This can be achieved introducing a special role R_{ideal} denoting the relation between individuals and their "ideal" counterparts, in the exact same fashion in which Kripke semantics relates possible worlds in the standard interpretation of deontic logic. This special role needs then to be axiomatized as follows. For all i s.t. $1 \leq i \leq n$:

$$\begin{aligned}
&\top_i : \forall R_{\text{ideal}}.\top \equiv \top \\
&\top_i : \forall R_{\text{ideal}}.(\gamma_1 \sqcap \gamma_2) \equiv \forall R_{\text{ideal}}.(\gamma_1) \sqcap \forall R_{\text{ideal}}.(\gamma_2) \\
&\top_i : \exists R_{\text{ideal}}.\top \equiv \top.
\end{aligned}$$

Intuitively the axioms state that, for all top contexts on each language i, it holds that: R_{ideal} defines a relation between individuals (first and second axioms), and that each individual is always R_{ideal}-related to at least one individual, that is, each individual has always at least one ideal counterpart (third axiom).

Within such a setting the norm of ONT to the effect that non-operable data ought not to be operated can be expressed in $\mathcal{L}^{CT}$ as follows:

$$c_{\text{ONT}} : \neg\texttt{operable\_data} \sqsubseteq \forall R_{\text{ideal}}.\neg\texttt{operated\_data}.$$

Literally, according to c_{ONT} if something is a non-operable data, then it is classified as something the ideal counterpart of which is a non-operated data, that is, it is classified as something which ought not to be operated.

5.3. Reduction or explicit deontics?

We have now two ways of readily expressing norms in a $\mathcal{L}^{CT}$ language. We do not take a stance in favour of any of them, but rather provide a comparison. The main difference resides in the expressivity of the two solutions. Namely, the reduction approach sketched in Section 5.1 does not allow for expressing nested obligations: "it ought to be the case that it ought to be the case that data are not operated". Obligations being classifications of a plain atom representing violation, cannot appear as a term in classifications themselves. In other words, we cannot express the fact that an obligation is obligatory for the simple reason that obligations are themselves classifications.

This can be done instead within the second approach making use of special roles (Section 5.2). Special roles allow to encode a notion of obligation out of the classification itself at the syntactic level of compound concepts: the concept $\forall R_{\text{ideal}}.\neg\mathtt{operated\_data}$ denotes literally all individuals which should not be operated data. Given this, the fact that such an obligation is obligatory can be easily represented nesting the application of the special role: $\forall R_{\text{ideal}}.(\forall R_{\text{ideal}}.\neg\mathtt{operated\_data})$. Intuitively, this latter concept denotes all individuals all the ideal counterparts of which are such that they ought not to be operated data.

The comparison of the two approaches could be further developed especially at a technical level. This falls out of the scope of the present paper though. What is worth stressing, which neatly displays the core difference between the two variants, is their intuitive reading. Representing norms via reduction amounts to classify forbidden properties under a violation atom: what ought not to be the case is what happens to be subsumed under such a concept, i.e., what counts as a violating instance. Representing norms via special roles, instead, amounts to represent obligatory properties via complex concepts built up from those special roles: what ought to be the case is what holds for all "ideal" individuals.

It is finally worth devoting some words to related work in the representation of deontic notions in a contextual setting. We already mentioned that the approach sketched in Section 5.1 is directly connected with the deontic logic reduction proposed, within a modal logic setting, also in Grossi, Meyer and Dignum (2005d). To our knowledge, in the literature on deontic logic, there are only two more attempts to represent deontics together with some notion of context: Krabbendam and Meyer (2003) and van der Torre (2003).

Besides the fact that the aforementioned approaches are developed in modal logic, the main difference that we can point out in relation with our work consists in the way the notion of "context" is interpreted and captured in the formalism. In Krabbendam and Meyer (2003) a contextual deontic logic is proposed which makes use of a reduction approach within a release logic framework (Krabbendam and Meyer, 1997). Release logics are a way to represent context as information about what is relevant or not for asserting the truth of a proposition. In that framework, formulas such as $\Delta_{\mathcal{P}}\,\alpha$ express that α holds while $\mathcal{P}$, i.e., a chosen set of parameters, is not relevant. The notion of context is therefore interpreted in terms of a notion of relevance.

In van der Torre (2003) a logic for tryadic deontic modal logic is introduced. Formulas such as $O_{\gamma}(\alpha \mid \beta)$ denote that α is obligatory given γ in context β. A semantics for these formulas is provided which extends the standard preferred worlds semantics for dyadic deontic logics (Hansson, 1969) labeling the preference relations of the semantics in order to model a notion of context. Contexts are therefore seen as labels on preference orderings on worlds in a preferred world semantics: "according to context β, α is preferred givent that γ is true".

On our side, the semantics of context is not related to any specific notion such as relevance or preference.[20] Instead, both the variants we sketched for representing deontics in $\mathcal{L}^{CT}$ languages capture context only as means for *localizing* reasoning -in our case subsumptions- along the lines of various work carried out in the area of formal context representation (Benerecetti, Bouquet and Ghidini, 2000; Ghidini and Giunchiglia, 2001).[21]

5.4. The example revisited: Ontologies + norms

We are now in the position to provide a formalization of scenarios such as the one we addressed in Section 4.1 in which the normative aspects, besides the ontological, are made explicit. Suppose then that we are interested in formalizing within the same scenario also the norm, pertaining to the context of European regulation on personal data protection (c_{SUP} as we denoted it above), to the effect that non-operable data should not be operated. To do this, the languages $\mathcal{L}_0$ and $\mathcal{L}_1$ defined in Example should be first appropriately expanded adding the concept `operated_data` and `violating` or alternatively the role R_{ideal}. The following expressions should then be added to the formalization depicted in Fig. 1:

$$c_{\text{SUP}} : \neg\texttt{operable\_data} \sqcap \texttt{operated\_data} \sqsubseteq \texttt{violating}$$

$$c_{\text{SUP}} : \neg\texttt{operable\_data} \sqsubseteq \forall R_{\text{ideal}}.\neg\texttt{operated\_data}$$

respectively making use of the reduction approach sketched in Section 5.1 and of the special roles approach exposed in Section 5.2.

Being contextualized with respect to the context c_{SUP}, these subsumptions are inherited by the subordinated contexts c_{ONT} and c_{PF} by virtue of formula (6): $c_{\text{ONT}} \curlyvee_0 c_{\text{PF}} \preccurlyeq c_{\text{SUP}}$.[22] This is a key aspect to be captured in a formal characterization of normative reasoning in institutions: norms are issued at an abstract level (c_{SUP} in the example) and then further refined at more concrete levels c_{ONT} and c_{PF} in the example) via specifying ontologies for the concepts involved in the norms. As it appears in the example discussed, these ontologies can diverge prescribing completely different constraints which nevertheless stem from the very same norm. In fact, with respect to personal data referring to race the two concrete concepts prescribe two different behaviours. Using the reduction representation, it can be proved that:

$$c_{\text{ONT}} : \texttt{personal\_data} \sqcap \exists\text{refer } \texttt{race} \sqcap \texttt{operated\_data} \sqsubseteq \texttt{violating}$$

$$c_{\text{PF}} : \texttt{personal\_data} \sqcap \exists\text{refer } \texttt{race} \sqcap \texttt{operated\_data} \sqsubseteq \texttt{violating}$$

---

[20] See Section 3.3.

[21] See Section 3.

[22] This aspect was discussed in detail in Section 4.2 in relation with the inheritance of the abstract classification: $c_{\text{SUP}} : \texttt{personal\_data} \sqcap \texttt{relevant\_data} \sqsubseteq \texttt{operable\_data}$.

Or alternatively using special roles:

$$c_{\mathrm{ONT}} : \mathtt{personal\_data} \sqcap \exists \mathtt{refer\ race} \sqsubseteq \forall R_{\mathrm{ideal}} \neg \mathtt{operated\_data}$$

$$c_{\mathrm{PF}} : \mathtt{personal\_data} \sqcap \exists \mathtt{refer\ race} \sqsubseteq \forall R_{\mathrm{ideal}} \neg \mathtt{operated\_data}$$

To sum up, ontologies provide interpretations of the concepts involved in normative statements and make them concrete. In this very sense, ontologies guide the "implementation" of norms within institutions allowing at the same time normative statements to remain vague and therefore stable.

6. Contextual ontologies in eInstitutions

In earlier 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 have given a formal framework to formalise 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 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 that can be described in a formal framework such as ours, 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 point of 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.[23] From the ontological perspective, the most complex step is the definition of its ontology, as several 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 that one needs to create some kind of link from the ontologies of different supercontexts to the institutional ontology. In our approach (which is ongoing work), the links between

---

[23] More details on the implementation of norm enforcement mechanisms can be found in Vázquez-Salceda, Aldewereld and Dignum (2004).

ontologies are explicitly defined by the designer by means of different kinds of abstraction and 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 Organisation (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 transnational institutions such as Eurotransplant[24], where different ontological definitions of terms may appear in each of the countries where the institution should operate. In this case, when inheriting different, 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.

From the individual agents' perspective, the ontological aspects of norms and the issue of multi-contextual ontologies influences the on-line reasoning cycle of the agent. That is, when an agent tries to enter an eInstitution it is told which ontologies and norms are used in the eInstitution. However, the ontology used by the eInstitution need not be the same as that of the agent, and concepts in the norms used in the eInstitution might be unclear to the agent. In this case, the eInstitution and agent need to obtain a common understanding of the concepts such that it provides the agent with a clear meaning of the norms used in the institution. This can be done by finding a common supercontext and using the ontology's abstraction and inheritance relations to this supercontext.

7. Conclusions

The motivating question of our research was how 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"). As institutions are supposed to regulate completely different domains, the ontologies they are based on are also different. They can be specified on completely different vocabularies, or, if they share a set of terms, they may interpret it in divergent ways (which is the case of the concept of "relevant data" we discussed in our example). To get a grip on this phenomenon, we made use of contexts as means to localise these ontological discrepancies: institutions are based on ontologies, and these ontologies are contextual. This is also the analytical setting in which we provided a clear understanding of the so called *counts-as* phenomenon; counts-as statements

---

[24] The Eurotransplant International Foundation is responsible for the mediation and allocation of organ donation procedures in Austria, Belgium, Germany, Luxembourg, the Netherlands and Slovenia.

are nothing but contextual subsumption relations: they are the basic brick by means of which institutions establish their ontologies.

This analysis has then been framed in a rigorous setting. The formal framework exposed is based on a specific understanding of the notion of *context* as set of models for particular description logic languages, and provides a formal characterisation of the notion of *contextual ontology*. This framework is also used for formalising an example. At the end of the paper we also provided some general ideas on how these contextual ontologies can be concretely used in order to specify and reason about eInstitutions.

References

Akman V, Surav M (1996) Steps toward formalizing context. AI Magazine 17(3):55–72

Alchourrón CE Bulygin E (1986) normative systems, Wien, Springer Verlag

Anderson A (1957) The formal analysis of normative concepts. Amer Sociol Rev 22:9–17

Anderson A (1958) A reduction of deontic logic to alethic modal logic. Mind 22:100–103

Baader F, Calvanese D, McGuinness D, Nardi D, Patel-Schneider P (2002) The description logic handbook. Cambridge University Press, Cambridge

Benerecetti M, Bouquet P, Ghidini C (2000) Contextual reasoning distilled. J Exper Theor Art Intell (JETAI) 12(3):279–305

Blackburn P, de Rijke M, Venema Y (2001) Modal logic. Cambridge University Press, Cambridge

Boella G, Van der Torre L (2004) Regulative and constitutive norms in normative multiagent systems. In: Proceedings of KR2004, Whistler, Canada, pp 255–266

Broersen J, Dignum F, Dignum V, Meyer J-JC (2004) Designing a deontic logic of deadlines. In: 7th Int. workshop on deontic logic in computer science (DEON'04), Portugal

Bulygin E (1992) On norms of competence. Law and Philosoph 11:201–216

d'Altan P, Meyer J-JC, Wieringa R (1993) An integrated framework for ought-to-be and ought-to-do constraints. In: Tan YH (ed) Working papers of the workshop om deontic and non-monotonic logics, Rotterdam, pp 14–45

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, pp 93–104

Dignum F, Kinny D, Sonenberg L (2002) From desires, obligations and norms to goals. Cogn Sci Quart 2(3-4):407–430

Dignum F, Broersen J, Dignum V, Meyer J-JCh (2004) Meeting the deadline: Why, when and how. In: 3rd goddard workshop on formal approaches to agent-based systems (FAABS), Maryland

Dignum V, Meyer J-JCh, Dignum F, Weigand H (2002) Formal specification of interaction in agent societies. In: 2nd goddard workshop on formal approaches to agent-based systems (FAABS), Maryland

Directive ed9546 (1995) Directive 95/46/CE of the European parliament and of the council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and of the free movement of such data

Eurotransplant Eurotransplant international foundation http://www.eurotransplant.nl

Giunchiglia F, Bouquet P (1997) Introduction to contextual reasoning. Perspectives on Cognitive Science 3

Ghidini C, Giunchiglia F (2001) Local models semantics, or contextual reasoning = locality + compatibility. Art Intell 127(2):221–259

Ghidini C, Giunchiglia F (2004) A semantics for abstraction. In: Mántaras Rde, Saitta L (eds) Proceedings of ECAI'2004, including PAIS 2004, pp 343–347

Grossi D, Dignum F (2004) From abstract to concrete norms in agent institutions. In: Hinchey MG, Rash JL, Truszkowski WF et al (eds) Formal approaches to agent-based systems: 3rd international workshop, FAABS 2004, LNCS 3228, Springer-Verlag, pp 12–29

Grossi D, Dignum F, Meyer J-JCh (2005a) Context in categorization. In: Proceedings of CRR'05, international workshop on context representation and reasoning, paris

Grossi D, Dignum F, Meyer J-JCh (2005b) Contextual taxonomies. In: Leite J, Toroni P (eds) Proceedings of CLIMA V workshop, Lisbon, September, LNAI 3487, Springer, pp 33–51

Grossi D, Dignum F, Meyer J-JCh (2005c) Contextual terminologies. In: Toni F Toroni P (eds) Proceedings of CLIMA VI, 6th international workshop on computational logic in multi-agent systems, London

Grossi D, Meyer J-JCh, Dignum F (2005d) Modal logic investigations in the semantics of counts-as. In: Proceedings of the 10th international conference on artificial intelligence and law (ICAIL'05), ACM, pp 1–9
Hansson B (1969) An analysis of some deontic logics. Nous (3):373–398
Hart HLA (1961) The concept of law. Clarendon Press, Oxford
Jones AJI, Sergot M (1992) Deontic logic in the representation of law: towards a methodology. Art Intell Law 1
Jones AJI, Sergot M (1993) On the characterization of law and computer systems. Deontic Logic in Comp Sci, pp 275–307
Krabbendam J, Meyer J-JCh (1997) Release logics for temporalizing dynamic logic. In: Barringerand H, Fischer M, Gabbay D, Gough G (eds) Proceedings 2nd International Conference on Temporal Logic (ICTL '97), Manchester, pp 19–42
Krabbendam J, Meyer J-JCh (2003) Contextual deontic logics. In: McNamara P, Prakken H (eds) Norms, Logics and Information Systems, IOS Press, Amsterdam, pp 347–362
Lomuscio A, Nute D (eds) (2004) Deontic logic in computer science, 7th international workshop on deontic logic in computer science, deon 2004, madeira, portugal, may 26–28, 2004. Proceedings. vol. 3065 of Lecture Notes in Computer Science, Springer
Lomuscio A, Sergot M (2003) Deontic intepreted systems. Studia Logica 75:63–92
McCarthy J (1986) Notes on formalizing contexts. In: Kehler T, Rosenschein S (eds) Proceedings of the 5th national conference on artificial intelligence, Los Altos, Morgan Kaufmann, California, pp 555–560
McCarthy J (1987) Generality in artificial intelligence. Communi ACM 30(12):1030–1035
Peczenik A (1989) On law and reason. Kluwer, Dordrecht
Regulation er45 (2001) Regulation (EC) No. 45/2001 of the European parliament and of the council of 18 December 2000 on the protection of individuals with regard to the processing of personal data by the Community institutions and bodies and on the free movement of such data, Official Journal of the European Communities
Ross A (1968) Directives and norms. Routledge, & Kegan Paul London
Schild K (1991) A correspondence theory for terminological logics: Preliminary report. In: Proceedings of IJCAI-91, 12th international joint conference on artificial intelligence, Sidney, AU, pp 466–471
Searle J (1995) The construction of social reality. Free Press
Shoham Y (1991) Varieties of context. Academic Press Professional, Inc., pp 393–407
van der Torre, LWN (2003) Contextual deontic logic: Normative agents, violations and independence. Annals Math Art Intelli 37(1–2):33–63
Vázquez-Salceda J (2004) The role of norms and electronic institutions in multi-agent systems. The HARMONIA framework, Whitestein Series in Software Agent Technology, Birkhäuser Verlag
Vázquez-Salceda J, Aldewereld H, Dignum F (2004) Implementing norms in multiagent systems. In: Lindemann G, Denzinger J, Timm I, Unland R (eds) Multagent System Technologies, LNAI 3187, Springer-Verlag, pp 313–327
Von Wright GH (1963) Norm and action. A logical inquiry. Routledge, London
Wright GHV (1951) Deontic logic. Mind 60:1–15

Davide Grossi studied Philosophy at Università di Pisa, Italy. Besides, he obtained the diploma on philosophical disciplines of Scuola Normale Superiore, an Italian university center for teaching and research. Since 2003 he is Ph.D. student at the Intelligent Systems Group at Universiteit Utrecht. His research focuses on the analysis of the formal aspects of norms, institutions and organizations.

Huib Aldewereld studied Cognitive Artificial Intelligence at Utrecht University in the Netherlands, graduating in 2002 on the subject of logics for BDI-like multiagent systems. At the moment he is finishing his Ph.D. at the Intelligent Systems group in Utrecht, researching norm enforcement in agent-mediated electronic institutions and the verification of norm compliance of protocols.

Javier Vázquez-Salceda studied Computer Science at the Politechnical University of Catalonia (UPC), in Barcelona, Spain. After his master studies he became research assistant in the KEMLg Group at UPC. In 2003 he presented his Ph.D. dissertation (with honors). For two years he was Postdoc researcher in the Intelligent Systems Group at the Universiteit Utrecht. Nowadays he is back at UPC as researcher. His research is focused on theoretical and applied issues of norms and institutions, and the conflict between software and physical agents' autonomy and social control.

Frank Dignum received his PhD in 1989 from the Free University in Amsterdam. After that he has been a lecturer in Swaziland and Portugal before settling down at the Technical University of Eindhoven. Since 2000 he became associate professor at the Universiteit Utrecht. His main research interests are in the social aspects of multi-agents systems and the regulation of mult-agent systems through norms and institutions.