

ONTOLOGICAL ASPECTS IN THE IMPLEMENTATION OF NORMS IN AGENT-BASED ELECTRONIC INSTITUTIONS

D. Grossi, H. Aldewereld, J. Vázquez-Salceda, and F. Dignum

Institute of Information and Computing Sciences

Utrecht University

PO Box 80.089

3508TB Utrecht

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 operationalisation 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 analyse this phenomenon introducing a notion of contextual ontology. We will focus on the formal machinery necessary to characterise it as well.*

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, see for instance [8] and [16]; this means that norms are expressed in terms of concepts that are, on purpose, kept vague and ambiguous, cf. [18]. 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 [8]. 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 have focused on declarative aspects of norms, see [10] and [11], formally defining norms by means of some variations of deontic logic that include conditional and temporal aspects, in [5] and [9]. 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 [28]. 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 (we will also call them contextual terminologies) based on [17,?]. Using this framework we will formalise an example in section 4. After this, we will discuss the implementational aspects of our framework in section 5 and we end the paper with some discussion, conclusions and future work.

Throughout the paper, 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 (see [2], [20], [19], [26] and [4]), 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* ([29]), *conceptual rules* ([6]), *qualification norms* ([23]), *definitional norms* ([19]). 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. [25] and [26]. 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 [26], where constitutive rules are viewed as representable via the following type of statements: “X counts as Y *in context C*”.

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. [27], 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 eInstitutions 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 in [21,22], is to include *contexts* 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 (from [14]):

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

With this idea, in [27], 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 regulations¹ *personal data* are defined as “[...] those [data] which

¹ European Parliament created the 95/46/CE Directive ([12]) 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 ([24]), 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.

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:

- *individual's consent*: as a general rule, personal data collection and processing requires the approval of the affected person.
- *rights over the collected information*: each person has the right to access, amend, cancel or be opposed to the collection of her personal data,
- *data maintenance*: Personal data will 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 uses 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 *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 of each EU member state. In such a document there are very specific definitions of which are the *allowed data* and the *allowed data processing* for that particular organisation. Once the allowed data is defined, all regulations on data protection reduce to a single rule: organisations can only store a subset of the relevant data, the *allowed data*, and can only use such data by means of the *allowed data processing*.

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 *allowed 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 *allowed 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 [12].

- 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 [12].

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 *allowed 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 [12].

Therefore, inside the context of a hospital information such as *name*, *address* or *blood type* are usually relevant, and belongs to the set of *allowed 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 Organisation² (ONT) is an interesting, specific subcontext of *EU Health Service*. By Spanish Law, ONT must ensure an equitable 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 *allowed 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

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 languages³. 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 general as possible) for contextualising 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”.

² The Organización Nacional de Trasplantes 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.

³ This is a much acknowledged characteristic of contextual reasoning in general, see [21].

- 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.

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 them⁴.

These intuitions about the semantics of context operators will be clarified and made more rigorous in section 3.2 dedicated to the formal semantics of the framework, and in section 4.1 in which we will formalise an example.

3.1 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 the union of concepts, full existential quantification (we want to deal with concepts such as “either car or bicycle” and “persons which drive cars”)⁵.

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

⁴ Note 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.

⁵ This type of language is usually referred to as \mathcal{ALUE} , or \mathcal{ALC} . Within this type of languages the negation of arbitrary concepts is also enabled, see [3].

(R), i.e., atomic attributes. 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 attribute-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 a contextual subsumption relation symbols “ $. : . \sqsubseteq .$ ” which is used for both concept contextual subsumption (within context c , concept A_1 is a subconcept of concept A_2) and attribute contextual subsumption (within context c , attribute R_1 is a subattribute of attribute R_2). Lastly, the alphabet of \mathcal{L}^{CT} contains also the sentential connectives \sim (negation) and \wedge (conjunction)⁶.

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

$$\xi ::= c \mid \perp_i \mid \top_i \mid \neg_i(\xi) \mid abs_i(\xi) \mid \xi_1 \wedge_i \xi_2 \mid \xi_1 \vee_i \xi_2.$$

Concept constructs and attribute 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 attributes descriptions (ρ) coincides with the set of all atomic attributes.

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

$$\begin{aligned} \alpha ::= & \xi : \gamma_1 \sqsubseteq \gamma_2 \mid \xi : \rho_1 \sqsubseteq \rho_2 \mid \xi_1 \preceq \xi_2 \mid \sim \alpha \\ & \mid \alpha_1 \wedge \alpha_2. \end{aligned}$$

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. Contextual subsumption relations are the expression by mean of which we give a rigorous characterisation of searlean statements: “X counts as Y in context C”, [26]. This kind of expressions are, in a nutshell, what we are interested in formalising.

In the formalisation of the example, the following symbols will be also used “ $. : . \sqsubset .$ ” (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:

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

⁶ It might be worth remarking that language \mathcal{L}^{CT} is, then, an expansion of each \mathcal{L}_i language.

3.2 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 characterise the class of operations on contexts that will constitute the semantic counterpart of the context operators symbols introduced in the language. 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 , cf. [3].

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

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

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:

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

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 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 Operations on contexts

Before getting to this, let us first recall a notion of *domain restriction* (\lceil) 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: $C \lceil f(x) = \{y \mid y = f(x) \ \& \ x \in C\}$, cf. [7].

Definition 3. (Operations on contexts)

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

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

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

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

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

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 \pitchfork_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 jeopardising the fairness of the match.

$$\begin{aligned}
abs_0(c_{CONT}) \curlyvee_0 abs_0(c_{PPF}) &\preceq c_{SUP} & (5) \\
c_{SUP} : \mathbf{personal_data} \sqcap \mathbf{relevant_data} &\sqsubseteq \mathbf{allowed_data} & (6) \\
c_{CONT} \curlyvee_1 c_{PPF} : \mathbf{personal_data} \sqcap \exists \text{refer.blood.type} &\sqsubseteq \mathbf{relevant_data} & (7) \\
c_{CONT} \curlyvee_1 c_{PPF} : \mathbf{personal_data} \sqcap \exists \text{refer.anthropometric.properties} &\sqsubseteq \mathbf{relevant_data} & (8) \\
c_{CONT} \curlyvee_1 c_{PPF} : \mathbf{personal_data} \sqcap \exists \text{refer.race} &\sqsubseteq \mathbf{personal_data} & \\
&\sqcap \exists \text{refer.anthropometric.properties} \sqcup \neg \mathbf{relevant_data} & (9) \\
c_{CONT} : \mathbf{race} &\sqsubseteq \neg \mathbf{anthropometric.properties} & (10) \\
c_{PPF} : \mathbf{race} &\sqsubseteq \mathbf{anthropometric.properties}. & (11)
\end{aligned}$$

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

3.5 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:

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

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} . 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:

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

A contextual concept subsumption relation between γ_1 and γ_2 holds iff 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 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, 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. Perfectly analogous observations hold also for the clause regarding contextual attribute subsumption relations. The \preceq relation between context constructs is interpreted as a standard subset relation: $\xi_1 \preceq \xi_2$ 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 . Note that this relation, being interpreted on the \subseteq relation, is reflexive, antisymmetric and transitive. In [16] 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 [15]. 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.

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.

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 allowed data, because their understanding of the norm lies in divergent ontologies of the concepts involved⁷.

To formalise the scenario 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 attribute: `refer`. From this language we obtain $2^6 - 1 \cdot 2$ languages \mathcal{L}_i ⁸. Three atomic contexts are at issue here: the context of the superordinate European regulation, let us call it c_{SUP} ; the contexts of the municipal regulations ONT and PF, let us call them c_{ONT} , c_{PF} and c_{M3} 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:

$$\begin{aligned} \mathbf{A}_{\sigma} &= \{\text{personal_data, relevant_data,} \\ &\quad \text{allowed_data}\}, \\ \mathbf{R}_{\sigma} &= \emptyset \end{aligned}$$

and

$$\begin{aligned} \mathbf{A}_{\tau} &= \{\text{personal_data, relevant_data,} \\ &\quad \text{allowed_data, blood_type,} \\ &\quad \text{race, anthropometric_properties}\}, \\ \mathbf{R}_{\tau} &= \{\text{refer}\}. \end{aligned}$$

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 attribute.

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

Formula (5) 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 (6)-(11) all express contextual subsumption relations. It is worth stressing that they can all be seen as formalising counts-as

⁷ It is instructive to notice, in passing, that no deontics is actually enabled in our formalism. Indeed, the norm according to which only relevant personal data can be operated will be treated as a subsumption statement. This might be regarded as simplistic, but notice that our attention here does not focus on normative reasoning problems such as reasoning about violations at the level on individuals (ABox), and therefore no deontics is strictly required here.

⁸ See section 3.3 in which the language \mathcal{L}^{ct} is presented.

statements which specify the ontologies holding in the contexts at issue. Formula (6) formalises the abstract rule to the effect that personal data which are relevant for the accomplishment of the aim of the organisation 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, 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 (10) and (11) 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 (5)-(11) in figure 2. We will focus on subsumptions contextualised 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_{PF} . These are indeed the formulas

$$\begin{aligned}
& (5), (6) \models_{C_{ONT}} \text{personal_data} \sqcap \text{relevant_data} \sqsubseteq \text{allowed_data} \\
(5), (6), (7) \models_{C_{ONT}} & \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{relevant_data} \\
(5), (6), (7) \models_{C_{ONT}} & \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{allowed_data} \\
(5), (6), (8) \models_{C_{ONT}} & \text{personal_data} \sqcap \exists \text{refer.anthropometric_properties} \sqsubseteq \text{relevant_data} \\
(5), (6), (8) \models_{C_{ONT}} & \text{personal_data} \sqcap \exists \text{refer.anthropometric_properties} \sqsubseteq \text{allowed_data} \\
(8), (10) \models_{C_{ONT}} & \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \text{personal_data} \\
& \sqcap \exists \text{refer.}\neg\text{anthropometric_properties} \\
(5), (6), (9), (10) \models_{C_{ONT}} & \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \neg\text{relevant_data} \\
\\
& (5), (6) \models_{C_{PF}} \text{personal_data} \sqcap \text{relevant_data} \sqsubseteq \text{allowed_data} \\
(5), (6), (7) \models_{C_{PF}} & \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{relevant_data} \\
(5), (6), (7) \models_{C_{PF}} & \text{personal_data} \sqcap \exists \text{refer.blood_type} \sqsubseteq \text{allowed_data} \\
(5), (6), (8) \models_{C_{PF}} & \text{personal_data} \sqcap \exists \text{refer.anthropometric_properties} \sqsubseteq \text{relevant_data} \\
(5), (6), (8) \models_{C_{PF}} & \text{personal_data} \sqcap \exists \text{refer.anthropometric_properties} \sqsubseteq \text{allowed_data} \\
(8), (11) \models_{C_{PF}} & \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \text{personal_data} \\
& \sqcap \exists \text{refer.anthropometric_properties} \\
(5), (6), (9), (11) \models_{C_{PF}} & \text{personal_data} \sqcap \exists \text{refer.race} \sqsubseteq \text{relevant_data}
\end{aligned}$$

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

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 (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_{PF} 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. Those differences determine a variation in the interpretation of the abstract norm expressed in (6).

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 (10) and by means of the “no-cheating” clause (9). In fact, in this context, data about race are not anthropometric data. Context c_{PF} , instead, expresses a different view. Since race counts as anthropometric information (11), data about race are actually relevant data and, as such, can be used.

Before ending the section, we 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 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 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 dependent⁹. 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 also has 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.

⁹ The issue of the relationship between contextuality and defeasibility has been raised also in [1].

5 Specifying Contextual Ontologies for eInstitutions

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 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.¹⁰ 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 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 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

¹⁰ More details on the implementation of norm enforcement mechanisms can be found in [28].

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 Eurotransplant¹¹, 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.

6 Conclusions and Future Work

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 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

¹¹ The Eurotransplant International Foundation is responsible for the mediation and allocation of organ donation procedures in Austria, Belgium, Germany, Luxembourg, the Netherlands and Slovenia.

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.

With respect to future work, we firstly intend to develop an extension of the framework which can enable a full-fledged interaction of context and ontological features with the more standard normative reasoning issues (eminently, reasoning with violations). This requires to focus also on aspects concerning reasoning with instances of concepts (what in description logics is called ABox), and of course on the inclusion of some deontic logic. Secondly, this extension should be brought into practice and applied in the development of eInstitutions and Normative Agents.

References

1. V. Akman and M. Surav. Steps toward formalizing context. *AI Magazine*, 17(3):55–72, 1996.
2. C. E. Alchourrón and E. Bulygin. *Normative Systems*. Springer Verlag, Wien, 1986.
3. F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, and P.F. Patel-Schneider. *The Description Logic Handbook*. Cambridge University Press, Cambridge, 2002.
4. G. Boella and L. Van der Torre. Regulative and constitutive norms in normative multiagent systems. In *Proceedings of KR2004, Whistler, Canada*, pages 255–266, 2004.
5. J. Broersen, F. Dignum, V. Dignum, and J.-J. Ch. Meyer. Designing a Deontic Logic of Deadlines. In *7th Int. Workshop on Deontic Logic in Computer Science (DEON'04)*, Portugal, May 2004.
6. E. Bulygin. On norms of competence. *Law and Philosophy* 11, pages 201–216, 1992.
7. E. Casari. *La Matematica della Verità*. Privately distributed, 2002.
8. F. Dignum. Abstract norms and electronic institutions. In *Proceedings of the International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA '02)*, Bologna, pages 93–104, 2002.
9. F. Dignum, J. Broersen, V. Dignum, and J.-J. Ch. Meyer. Meeting the Deadline: Why, When and How. In *3rd Goddard Workshop on Formal Approaches to Agent-Based Systems (FAABS)*, Maryland, April 2004.
10. F. Dignum, D. Kinny, and L. Sonenberg. From Desires, Obligations and Norms to Goals. *Cognitive Science Quarterly*, 2(3-4):407–430, 2002.
11. V. Dignum, J.-J.Ch. Meyer, F. Dignum, and H. Weigand. Formal Specification of Interaction in Agent Societies. In *2nd Goddard Workshop on Formal Approaches to Agent-Based Systems (FAABS)*, Maryland, Oct. 2002.

12. 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, October 1995.
13. Eurotransplant International Foundation.
<http://www.eurotransplant.nl>.
14. F. Giunchiglia and P. Bouquet. Introduction to contextual reasoning. *Perspectives on Cognitive Science*, 3, 1997.
15. A. I. Goldman. *A Theory of Human Action*. Princeton University Press, Princeton, 1976.
16. D. Grossi and F. Dignum. From abstract to concrete norms in agent institutions. In M. G. Hinchey, J. L. Rash, W. F. Truszkowski, and et al., editors, *Formal Approaches to Agent-Based Systems: Third International Workshop, FAABS 2004*, LNCS 3228, pages 12–29. Springer-Verlag, April 2004.
17. D. Grossi, F. Dignum, and J.-J. Ch. Meyer. Contextual taxonomies. In J. Leite and P. Toroni, editors, *Proceedings of CLIMA V Workshop, Lisbon, September*, LNAI 3487, pages 33–51. Springer, 2005.
18. H. L. A. Hart. *The Concept of Law*. Clarendon Press, Oxford, 1961.
19. A. J. I. Jones and M. Sergot. Deontic logic in the representation of law: towards a methodology. *Artificial Intelligence and Law 1*, 1992.
20. A. J. I. Jones and M. Sergot. On the characterization of law and computer systems. *Deontic Logic in Computer Science*, pages 275–307, 1993.
21. J. McCarthy. Notes on formalizing contexts. In T. Kehler and S. Rosenschein, editors, *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 555–560, Los Altos, California, 1986. Morgan Kaufmann.
22. J. McCarthy. Generality in artificial intelligence. *Communications of the ACM*, 30(12):1030–1035, 1987.
23. A. Peczenik. *On Law and Reason*. Kluwer, Dordrecht, 1989.
24. 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, January 12 2001.
25. A. Ross. *Directives and Norms*. Routledge & Kegan Paul, London, 1968.
26. J. Searle. *The Construction of Social Reality*. Free Press, 1995.
27. J. Vázquez-Salceda. *The Role of Norms and Electronic Institutions in Multi-Agent Systems. The HARMONIA framework*. Whitestein Series in Software Agent Technology. Birkhäuser Verlag, 2004.
28. J. Vázquez-Salceda, H. Aldewereld, and F. Dignum. Implementing norms in multiagent systems. In G. Lindemann, J. Denzinger, I.J. Timm, and R. Unland, editors, *Multagent System Technologies*, LNAI 3187, pages 313–327. Springer-Verlag, 2004.
29. G. H. Von Wright. *Norm and Action. A Logical Inquiry*. Routledge, London, 1963.