

Contextual Taxonomies

D. Grossi, F. Dignum, and J.-J. Ch. Meyer

Utrecht University,
The Netherlands
{davide,dignum,jj}@cs.uu.nl

Abstract. We provide a formal characterization of a notion of contextual taxonomy, that is to say, a taxonomy holding only with respect to a specific context. To this aim, a new proposal for dealing with “contexts as abstract mathematical entities” is set forth, which is geared toward solving some problems arising in the area of normative system specifications for modeling multi-agent systems. Contexts are interpreted as sets of description logic models for different languages, and a number of operations on contexts are defined. Using this framework, a simple scenario taken from the legal domain is modeled, and a formal account of the so called open-texture of legal terms is provided characterizing the notions of “core” and “penumbra” of the meaning of a concept.

1 Introduction

The motivation of this work lies in problems stemming from the domain of normative system specifications for modeling multi-agent systems ([1, 2]). In [3–5] contexts have been advocated to play a central role in the specification of complex normative systems. The notion of context has obtained attention in AI researches since the seminal work [6], and much work has been carried out with regard to the logical analysis of this notion (see [7, 8] for an overview). With this work, we intend to pursue this research line providing a logical framework for dealing with a conception of context specifically derived from the aforementioned application domain. We nevertheless deem that the formal analysis we are going to present may give valuable insights for understanding contexts in general, also outside our specific domain of interest.

In general, the purpose of the present work is to propose a framework for grounding a new formal semantics of expressions such as: “ A counts as B ([9]) in institution c ”, or “ B supervenes A in institution c ” ([10]), or “ A conventionally generates B in institution c ” ([11]), or “ A translates (means) B in institution c ” ([5]). These expressions, known in legal theory as *constitutive rules*, will be interpreted essentially as contextualized subsumption relations establishing taxonomies which hold only with respect to a specific (institutional) context. We came to a notion of *contextual taxonomy* through the analysis of some well known problems of underspecification, or more technically open-texture ([12]), typical of legal terminologies. These vagueness-related issues constitute, more concretely, the direct target of the work. We quote down here an excerpt from [13] neatly exposing this type of problems.

[Suppose a] legal rule forbids you to take a vehicle into the public park. Plainly this forbids an automobile, but what about bicycles, roller skates, toy automobiles? What about airplanes? Are these, as we say, to be called “vehicles” for the purpose of the rule or not? If we are to communicate with each other at all, and if, as in the most elementary form of law, we are to express our intentions that a certain type of behavior be regulated by rules, then the general words we use like “vehicle” in the case I consider must have some standard instance in which no doubts are felt about its application. There must be a *core* of settled meaning, but there will be, as well, a *penumbra* of debatable cases in which words are neither obviously applicable nor obviously ruled out. [...] We may call the problems which arise outside the hard core of standard instances or settled meaning “problems of the penumbra”; they are always with us whether in relation to such trivial things as the regulation of the use of the public park or in relation to the multidimensional generalities of a constitution.

Given a general (regional) rule not allowing vehicles within public parks, there might be a municipality allowing bicycles in its parks, and instead another one not allowing them in. What counts as a vehicle according to the first municipality, and what counts as a vehicle according to the second one then? This type of problems has been extensively approached especially from the perspective of the formalization of defeasible reasoning: the regional rule “all vehicles are banned from public parks” is defeated by the regulation of the first municipality stating that “all vehicles that are bicycles are allowed in the park” and establishing thus an *exception* to the general directive. The formalization of norms via non-monotonic techniques (see [14] for an overview) emphasizes the existence of exceptions to norms while understanding abstract terms in the standard way (all instances of bicycles are always vehicles). It has also been proposed to view the inclusion rules themselves as defaults: “normally, if something is a bicycle, then it is a vehicle” (for example [15, 5]). We deem these approaches, despite being effective in capturing the reasoning patterns involved in these scenarios, to be not adequate for analyzing problems related with the *meaning* of the terms that trigger those reasoning patterns. Those reasoning patterns are defeasible because the meaning of the terms involved is not definite, it is vague, it is -and this is the thesis we hold here- context dependent¹. We propose therefore to analyze these “*problems of the penumbra*” in terms of the notion of context: according to (in the context of) the public parks regulation of the first municipality bicycles are not vehicles, according to (in the context of) the public parks regulation of the second one bicycles are vehicles. This reading will be interpreted as follows: “the subsumption of the concept `bicycle` under the concept `vehicle` holds in the context of the first municipality, but not in the context of the second one”. A defeasible reasoning analysis leads to a quite different reading, which flattens the meaning of concepts and handles its variations by means of the notion of ex-

¹ The issue of the relationship between contextuality and defeasibility has been raised also in [7].

ception: “every exceptional instance of `bicycle` is not an instance of `vehicle`”. Bringing contexts into play will instead allow for a neat characterization of the notions of “core” and “penumbra” of the meaning of a concept, a characterization which is not obtainable via the use of a notion of exception.

The remainder of this paper is structured in accordance with the following outline. In Section 2 we will introduce the notion of *contextual taxonomy* making use of a concrete scenario; in Section 3 we will provide a formal framework based on a very simple type of description logic which accounts for this concept; in Section 4 we will provide a formalization of the scenario introduced, and we will formally characterize the notions of conceptual “core” and “penumbra”; in Section 5 we will discuss relations with other work; finally, in Section 6, some conclusive remarks are made.

2 Contextualizing Taxonomies

Let us now depict a simple scenario in order to state in clear terms the example used in the introduction.

Example 1. (The public park scenario) In the regulation governing access to public parks in region R it is stated that: “vehicles are not allowed within public parks”. In this regulation no mention is made of (possible) subconcepts of the concept `vehicle`, e.g., cars, bicycles, which may help in identifying an instance of `vehicle`. In municipal regulations subordinated to this regional one, specific subconcepts are instead handled. In municipality M1, the following rule holds: “bicycles are allowed to access public parks”. In M2 instead, it holds that: “bicycles are not allowed to access public parks”. In both M1 and M2 it holds that: “cars are not allowed in public parks”.

In this scenario the concept of `vehicle` is clearly open-textured. Instances of `car` (w.r.t. the taxonomies presupposed by M1 and M2) are “core” instances of `vehicle`, while instances of `bicycle` lay in the “penumbra” of `vehicle`. We will constantly refer back to this example in the remaining of the work. In fact, our first aim will be to provide a formal framework able to account for scenarios formally analogous to the aforementioned one².

Since the statement about the need for addressing “contexts as abstract mathematical entities” in [6], many formalizations of the notion have been proposed (see [7] or [8] for an overview). Our proposal pursues the line of developing a semantic approach to the notion of context according to what was originally presented in [16]. In that work contexts are formalized as sets of first order logic models and then connected via a relation which requires the sets of models constituting the different contexts to satisfy sets of domain specific inter-contextual inference rules (*bridge rules*). This theory has been variously used in work on specification of agent architectures ([17, 18]) where the stress lies in how contexts

² Note that this scenario hides a typical form of contextual reasoning called “categorization” ([8]), or “perspective” ([7]).

influence each other at a proof-theoretical level rather than at a semantical one (what can be inferred in this context, given that something holds in some other context?). We follow the basic intuition of understanding contexts as sets of models. Nevertheless, since we are mainly interested in taxonomies, much simpler models will be used here³). Moreover, we will partly depart from the proposal in [16] trying to characterize also a set of operations meaningfully definable on contexts. In fact, what we are interested in is also an articulate characterization of the interplay between contexts: how can contexts be joined, abstracted, etc. Instead of focusing on bridge rules, which have to be introduced outside and separately from the contexts, we will define some operations on contexts such that all possible compatibility relations will be generated by the semantics of the contexts alone. This will provide intrinsic boundaries within which other bridge rules may later be defined.

To summarize, we will expose an approach to contexts which is driven by intuitions stemming from the analysis of normative terminologies, and which is based on description logic semantics.

3 A Formal Framework

The main requirements of the formal framework that we will develop are the following ones.

1. It should enable the possibility of expressing lexical differences. A much acknowledged characteristic of contextual reasoning is, indeed, that contexts should be specified on different languages ([19–22]). The context of the national regulation about access to public parks should obviously be specified on a vocabulary that radically differs from the vocabulary used to specify the context of regulations about, for instance, immigration law: public park regulations do not talk about immigrants. Moreover, in Example 1, we observed that more concrete contexts make actually use of richer vocabularies: talking about vehicles comes down to talk about cars, bicycles, etc. In a nutshell, different contexts mean different ontologies and therefore different languages.
2. It should provide a formal semantics (as general as possible) for contextualized subsumption expressions, that is to say, for contextual taxonomies.
3. It should enable the possibility of describing operations between contexts.

Following these essential guidelines, a language and a semantics are introduced in what follows. 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 “dinosaurs” and “contemporary reptiles” can be intersected on a language talking about crocodiles generating a common less general context like “crocodiles”.

³ Basically models for description logic languages without roles. See Section 3.

- A *contextual disjunction* operator. Intuitively, it will yield a union of contexts: the contexts “viruses” and “bacterias” can be unified on a language talking about microorganisms generating a more general context like “viral or bacterial microorganisms”.
- A *contextual negation* operator. Intuitively, it will yield the context obtained via subtraction of the context negated: the negation of the context “viruses” on the language talking about microorganisms generates a context like “non viral microorganisms”.
- A *contextual abstraction* operator. Intuitively, it will yield the context consisting of some information extracted from the context to which the abstraction is applied: the context “crocodiles”, for instance, can be obtained via abstraction of the context “reptiles” on the language talking only about crocodiles. In other words, the operator prunes the information contained in the context “reptiles” keeping only what is expressible in the language which talks about crocodiles and abstracting from the rest.

Finally, also *maximum* and *minimum* contexts will be introduced: these will represent the most general, and respectively the less 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 where the semantics of the framework will be presented, and in Section 4.1 where an example will be formalized.

3.1 Language

The language we are interested in defining is nothing but a formal metalanguage for talking about sets of subsumption relations, i.e., what in description logic are called terminological boxes (TBoxes). In fact, we consider only TBoxes specified on very simple languages containing just atomic concepts and boolean operators⁵. We decided to keep the syntax of these languages poor mainly for two reasons: firstly, because the use of concept descriptions alone is enough to model the scenario depicted in Example 1, our core expressive concern being just to talk about concept subsumptions which vary from context to context; secondly, because this is still a preliminary proposal with which we aim to show how contextual reasoning and reasoning about vague notions are amenable to

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 We are going, indeed, to extend the language of propositional logic. Nevertheless, the semantics we are going to use in Section 3.2 is not the semantics of propositional logic, and it is instead of a description logic kind. For this reason we deem instructive to refer to these simple languages also as description logic languages of the type \mathcal{ALC} ([23]) but with an empty set of roles.

being handled on the basis of computationally appealing logics. On this basis it will be natural, in future, to consider also richer languages.

The alphabet of the language \mathcal{L}^{CT} (*language for contextual taxonomies*) 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 a non-empty finite set \mathbf{A}_i of atomic concepts (A), the zeroary operators \perp (bottom concept) and \top (top concept), the unary operator \neg , and the binary operators \sqcap and \sqcup ⁶.

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 $\{abs_i\}_{0 \leq i \leq n}$ (context abstraction operators) and $\{\neg_i\}_{0 \leq i \leq n}$ (context negation operators), two families of binary operators $\{\wedge_i\}_{0 \leq i \leq n}$ (context conjunction operators) and $\{\vee_i\}_{0 \leq i \leq n}$ (context disjunction operators), one context relation symbol \preceq (context c_1 “is at most as general as” context c_2) and a contextual subsumption relation symbol “. : . \sqsubseteq .” (within context c , concept A_1 is a subconcept of concept A_2), finally, the sentential connectives \sim (negation) and \wedge (conjunction)⁷. Thus, the set Ξ of context constructs (ξ) is defined through the following BNF:

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

Concepts and concept constructors are then defined in the usual way. The set Γ of concept descriptions (γ) is defined through the following BNF:

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

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

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

Technically, a *contextual taxonomy* in \mathcal{L}^{CT} is a set of subsumption relation expressions which are contextualized with respect to the same context, e.g.: $\{\xi : \gamma_1 \sqsubseteq \gamma_2, \xi : \gamma_2 \sqsubseteq \gamma_3\}$. This kind of sets of expressions are, in a nutshell, what we are interested in. Assertions of the form $\xi_1 \preceq \xi_2$ provide a formalization of the notion of *generality* often touched upon in context theory (see for example [6, 24]). In Section 4.1 the following symbol will be also used “. : . \sqsubseteq .” (within

⁶ It is worth stressing again that, in fact, a language \mathcal{L}_i , as defined here, is just a sub-language of languages of the type \mathcal{ALC} . Besides lacking roles symbols (and therefore role restriction operators), note that it also lacks the subsumption symbol. As we will see later in this section, the subsumption symbol is replaced by a set of contextualized subsumption symbols.

⁷ It might be worth remarking that language \mathcal{L}^{CT} is, then, an expansion of each \mathcal{L}_i language.

context c , concept A_1 is a proper subconcept of concept A_2). It can be defined as follows:

$$\xi : \gamma_1 \sqsubset \gamma_2 \stackrel{def}{=} \xi : \gamma_1 \sqsubseteq \gamma_2 \wedge \sim \xi : \gamma_2 \sqsubseteq \gamma_1.$$

A last category of expressions is also of interest, namely expressions representing what a concept means in a given context: for instance, recalling Example 1, “the concept `vehicle` in context $M1$ ”. These expressions, as it will be shown in Section 3.2, are particularly interesting from a semantic point of view. Let us call them contextual concept descriptions and let us define their set \mathcal{D} through the following BNF:

$$\delta ::= \xi : \gamma.$$

As we will see in Section 3.2, contextual concept descriptions \mathcal{D} play an important role in the semantics of contextual subsumption relations.

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 characterize the class of operations and relations on contexts that will constitute the semantic counterpart of the operators and relation symbols introduced in Section 3.1. Definitions of the formal meaning of our expressions will then follow.

Before pursuing this line, it is necessary to recollect the basic definition of a description logic model for a language \mathcal{L}_i ([23]).

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 $\mathcal{I}_m : \mathbf{A}_i \rightarrow \mathcal{P}(\Delta_m)$, that is, an interpretation of (atomic concepts expressions of) \mathcal{L}_i on Δ_m . This interpretation is extended to complex concept constructs via the following inductive definition:

$$\begin{aligned} \mathcal{I}_m(\top) &= \Delta_m \\ \mathcal{I}_m(\perp) &= \emptyset \\ \mathcal{I}_m(\neg A) &= \Delta_m \setminus \mathcal{I}_m(A) \\ \mathcal{I}_m(A \sqcap B) &= \mathcal{I}_m(A) \cap \mathcal{I}_m(B) \\ \mathcal{I}_m(A \sqcup B) &= \mathcal{I}_m(A) \cup \mathcal{I}_m(B). \end{aligned}$$

Out of technicalities, what a model m for a language \mathcal{L}_i does, is to assign a denotation to each atomic concept (for instance the set of elements of Δ_m that instantiate the concept `bicycle`) and, accordingly, to each complex concept (for instance the set of elements of Δ_m that instantiate the concept `vehicle` \sqcap \neg `bicycle`).

3.3 Models for \mathcal{L}^{CT}

We can now define a notion of *contextual taxonomy 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 . That is, $\forall m \in \mathbf{M}_i$, m is a model for \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 identifier 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. $0 \leq i \leq n$. Function \mathbb{I} can be seen as labeling sets of models on some language i via atomic context identifiers. Notice that \mathbb{I} fixes, for each atomic 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 models m is unique. We establish this constraint simply because we are interested in modeling different (taxonomical) conceptualizations of a same set of individuals.

This can be clarified by means of a simple example. Suppose the alphabet of \mathcal{L}^{CT} to be the set of atomic concepts $\{\text{allowed}, \text{vehicle}, \text{car}, \text{bicycle}\}$ and the set of atomic context identifiers $\{c_{M_1}, c_{M_2}, c_R\}$. The number of possible languages \mathcal{L}_i given the four aforementioned concepts is obviously $2^4 - 1$. A ct-model for this \mathcal{L}^{CT} language would have as domain the set of all sets of models for each of the $2^4 - 1$ \mathcal{L}_i languages, and as interpretation a function \mathbb{I} which assigns to each c_{M_1} , c_{M_2} and c_R an element of that set. We will come back to this in Section 4.1, where we discuss the formalization of the public park scenario.

The key feature of this semantics is that contexts are characterized as sets of models for the same 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 s.t. for each element of C , f and $C \lceil f$ return the same image. The exact definition is the following one: $C \lceil f(x) = \{y \mid y = f(x) \ \& \ x \in C\}$ ([25]).

Definition 3. (Operations on contexts)

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

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

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

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

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

Intuitively, the operations have the following meaning: operation 1 allows for abstracting the relevant content of a context with respect to a specific language; operations 2 and 3 express basic set-theoretical composition of contexts; finally, operation 4 returns, given a context, the most general of all the remaining contexts. Let us now provide some technical observations. First of all notice that operation \lceil_i yields the empty context when it is applied to a context M' the language of which is not an elementary expansion of \mathcal{L}_i . This is indeed very intuitive: the context obtained via abstraction of the context “dinosaurs” on the language of, say, “botanics” should be empty. Empty contexts can be also obtained through the \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 jeopardizing the fairness of the match. The following propositions can be proved with respect to the operations on contexts.

Proposition 1. (Structure of contexts on a given language)

The structure of contexts $\langle \mathcal{P}(\mathbf{M}_i), \cup_i, \pitchfork_i, -_i, \mathbf{M}_i, \emptyset \rangle$ on a language \mathcal{L}_i is a Boolean Algebra.

Proof. The proof follows straightforwardly from Definition 3. ■

Proposition 2. (Abstraction operation on contexts)

Operation \lceil_i is surjective and idempotent.

Proof. That \lceil_i is surjective can be proved per absurdum. First notice that this operation is a function of the following type: $\lceil_i : \mathcal{P}(\mathbf{M}_0) \cup \dots \cup \mathcal{P}(\mathbf{M}_n) \longrightarrow \mathbf{M}_i$. If it is not surjective then $\exists M'' \subseteq \mathbf{M}_i$ s.t. $\nexists M'$ in the domain of \lceil_i s.t. $\lceil_i M' = M''$. This means $\nexists \{m \mid m = \langle \Delta_{m'}, \mathbf{A}_i \lceil \mathcal{I}'_{m'} \rangle \ \& \ m' \in M'\}$, which is impossible. The proof of the equation for idempotency $\lceil_i(\lceil_i M) = \lceil_i M$ is straightforward. ■

These propositions clarify the type of conception of context we hold here: contexts are sets of models on different taxonomical languages; on each language the set of possible contexts is structured in a boolean algebra; the operation of abstraction allows for shifting from richer to simpler languages and it is, as we would intuitively expect, idempotent (abstracting from an abstraction yields the same first abstraction) and surjective (every context, even the empty one, can be seen as an abstraction of a different richer context, in the most trivial case, an abstraction of itself).

3.5 Formal meaning of Ξ , \mathcal{D} , 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:

$$\begin{aligned} \mathbb{I}(c) &= M \in \mathcal{P}(\mathbf{M}_0) \cup \dots \cup \mathcal{P}(\mathbf{M}_n) \\ \mathbb{I}(\perp_i) &= \emptyset \\ \mathbb{I}(\top_i) &= \mathbf{M}_i \\ \mathbb{I}(\xi_1 \wedge_i \xi_2) &= \mathbb{I}(\xi_1) \cap_i \mathbb{I}(\xi_2) \\ \mathbb{I}(\xi_1 \vee_i \xi_2) &= \mathbb{I}(\xi_1) \cup_i \mathbb{I}(\xi_2) \\ \mathbb{I}(\neg_i \xi) &= \neg_i \mathbb{I}(\xi) \\ \mathbb{I}(\text{abs}_i \xi) &= \downarrow_i \mathbb{I}(\xi). \end{aligned}$$

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

Definition 5. (Semantics of contextual concept descriptions: $\|\cdot\|_{\mathbb{M}}$)

The semantics of contextual concept description is defined as follows:

$$\|\xi : \gamma\|_{\mathbb{M}} = \{\mathcal{I}_m(\gamma) \mid m \in \mathbb{I}(\xi)\}.$$

The meaning of a concept γ in a context ξ is the set of denotations attributed to that concept by the models constituting that context.

It is worth noticing that if concept γ is not expressible in the language of context ξ , then $\|\xi : \gamma\|_{\mathbb{M}} = \emptyset$, that is, concept γ gets no denotation at all in context ξ . This happens simply because no interpretation for that concept is defined in the models constituting ξ . This shows also how Definition 5 allows to capture the intuitive distinction between concepts which lack denotation ($\|\xi : \gamma\|_{\mathbb{M}} = \emptyset$), and concepts which have a denotation which is empty ($\|\xi : \gamma\|_{\mathbb{M}} = \{\emptyset\}$): a concept that lacks denotation is for example the concept **immigrant** in the context of public park access regulation; in the same context, a concept with empty denotation is for example the concept **car** \sqcap **¬car**.

In what follows we will often use the notation $\mathbb{I}(\xi : \gamma)$ instead of the heavier $\|\xi : \gamma\|_{\mathbb{M}}$.

Definition 6. (Semantics of assertions: \models)

The semantics of assertions is defined as follows:

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

A contextual subsumption relation between γ_1 and γ_2 holds iff $\mathbb{I}(\xi)$ makes the meaning of γ_1 and γ_2 not empty and all models m of $\mathbb{I}(\xi)$ 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. The \preceq relation between context constructs is interpreted as a standard subset relation: $\xi_1 \preceq \xi_2$ means that context denoted by ξ_1 contains at most all the models that ξ_2 contains, that is to say, ξ_1 is *at most as general as* ξ_2 . Note that this relation, being interpreted on the \subseteq relation, is reflexive, symmetric and transitive. In [5] 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 [11]. 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.

The satisfaction clause of contextual subsumption relations deserves some more remarks. We observed that the satisfaction is conditioned to the meaningfulness of the terms involved with respect to the context. This condition is necessary because our contexts have different languages. Another way to deal with this would be to impose syntactic constraints on the formation of $\xi : \gamma_1 \sqsubseteq \gamma_2$ expressions, in order to distinguish the well-formed ones from the ill-formed ones. However, this would determine a dependence of the definition of well-formed expressions of \mathcal{L}^{CT} on the models \mathbb{M} of the language itself. Alternatively, the satisfaction relation itself might be restricted to consider only those subsumptions between concepts that, given the interpretation of the context, are interpreted as meaningful. Nevertheless, this option too determines a weird dependence, namely between the definition of the satisfaction relation and the models: the scope of the satisfaction would vary according to the models⁸. We chose for yet another solution, exploiting the possibility that our semantics enables of distinguishing meaningless concepts from concepts with empty extension (see Definition 5). By means of this feature it is possible to constrain the satisfaction of $\xi : \gamma_1 \sqsubseteq \gamma_2$ formulas, in such a way that, for them to be true, concepts γ_1 and γ_2 have to be meaningful in context ξ . Intuitively, we interpret contextual subsumption relations as inherently presupposing the meaningfulness of their terms.

⁸ Though in a completely different formal setting, this way is pursued in [21, 22].

4 Contextual Taxonomies, “Core” and “Penumbra”

4.1 Formalizing an example

We are now able to provide a formalization of the simple scenario introduced in Example 1 based on the formal semantic machinery just exposed.

Example 2. (The public park scenario formalized) To formalize the public park scenario within our setting a language \mathcal{L}^{CT} is needed, which contains the following atomic concepts: **allowed**, **vehicle**, **car**, **bicycle**. Three atomic contexts are at issue here: the context of the main regulation **R**, let us call it c_R ; the contexts of the municipal regulations **M1** and **M2**, let us call them c_{M1} and c_{M2} respectively. These contexts should be interpreted on two relevant languages. A language \mathcal{L}_0 for c_R s.t. $\mathbf{A}_0 = \{\mathbf{allowed}, \mathbf{vehicle}\}$; and a language \mathcal{L}_1 for c_{M1} and c_{M2} s.t. $\mathbf{A}_1 = \mathbf{A}_0 \cup \{\mathbf{car}, \mathbf{bicycle}\}$ (an abstract language concerning only vehicles and objects allowed to get into the park, and a more concrete one concerning, besides this, also cars and bicycles). A formalization of the scenario by means of \mathcal{L}^{CT} formulas is the following one:

$$abs_0(c_{M1}) \curlyvee_0 abs_0(c_{M2}) \preceq c_R \quad (5)$$

$$c_R : \mathbf{vehicle} \sqsubseteq \neg\mathbf{allowed} \quad (6)$$

$$c_{M1} \curlyvee_1 c_{M2} : \mathbf{car} \sqsubseteq \mathbf{vehicle} \quad (7)$$

$$c_{M1} : \mathbf{bicycle} \sqsubseteq \mathbf{vehicle} \quad (8)$$

$$c_{M2} : \mathbf{bicycle} \sqsubseteq \neg\mathbf{vehicle} \quad (9)$$

$$c_{M1} \curlyvee_1 c_{M2} : \mathbf{bicycle} \sqsubseteq \mathbf{vehicle} \sqcup \mathbf{allowed}. \quad (10)$$

Formula (5) plays a key role, stating that the two contexts c_{M1} , c_{M2} are concrete variants of context c_R . It tells this by saying that the context obtained by joining the two concrete contexts on language \mathcal{L}_0 (the language of c_R) is at most as general as context c_R . As we will see in discussing the logical consequences of this set of formulas, formula (5) makes c_{M1} , c_{M2} inherit what holds in c_R . Formula (6) formalizes the abstract rule to the effect that vehicles belongs to the category of objects not allowed to access public parks. Formula (7) states that in both contexts cars count as vehicles. Formulas (8) and (9) state the two different conceptualization of the concept of bicycle holding in the two concrete contexts at issue. These formulas show where the two contextual taxonomies diverge. Formula (10), finally, tells that bicycles either are vehicles or should be allowed in the park. Indeed, it might be seen as a clause avoiding “cheating” classifications such as: “bicycles counts as cars”.

It is worth listing and discussing some logical consequences of the formalization.

$$\begin{aligned} (5), (6) &\models_{c_{M1}} \mathbf{vehicle} \sqsubseteq \neg\mathbf{allowed} \\ (5), (6), (7) &\models_{c_{M1}} \mathbf{car} \sqsubseteq \neg\mathbf{allowed} \\ (5), (6), (8) &\models_{c_{M1}} \mathbf{bicycle} \sqsubseteq \neg\mathbf{allowed} \end{aligned}$$

$$\begin{aligned} (5), (6) &\models_{c_{M2}} \mathbf{vehicle} \sqsubseteq \neg\mathbf{allowed} \\ (5), (6), (7) &\models_{c_{M2}} \mathbf{car} \sqsubseteq \neg\mathbf{allowed} \\ (5), (6), (9), (10) &\models_{c_{M2}} \mathbf{bicycle} \sqsubseteq \mathbf{allowed} \end{aligned}$$

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. They formalize the two contextual taxonomies at hands in our scenario. Let us have a closer look. 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_{M1} and c_{M2} inherit the general rule stating that vehicles are not allowed to access public parks. Via this inherited rule, and via (7), it is shown that, in all concrete contexts, cars are also not allowed to access the park. As to cars then, all contexts agree. Where differences arise is in relation with how the concept of bicycle is handled. In context c_{M1} , since bicycles count as vehicles (8), bicycles are also not allowed. In context c_{M2} , instead, bicycles constitute an allowed class because they are not considered to be vehicles (9) and there is no bicycle which does not count as a vehicle and which does not belong to that class of allowed objects. In the following section we show in some more detail how a model for the formalization just exposed looks like.

4.2 A model of the formalization

Formulas (5)-(10) constrain ct-models in the following way:

$$\begin{aligned} \bigcup_0 \mathbb{I}(c_{M1}) \cup \bigcup_0 \mathbb{I}(c_{M2}) &\subseteq \mathbb{I}(c_R) \\ \forall m \in \mathbb{I}(c_R), \mathcal{I}_m(\mathbf{vehicle}) &\subseteq \Delta_1 \setminus \mathcal{I}_m(\mathbf{allowed}) \\ \forall m \in \mathbb{I}(c_{M1}) \cup \mathbb{I}(c_{M2}), \mathcal{I}_m(\mathbf{car}) &\subset \mathcal{I}_m(\mathbf{vehicle}) \\ \forall m \in \mathbb{I}(c_{M1}), \mathcal{I}_m(\mathbf{bicycle}) &\subset \mathcal{I}_m(\mathbf{vehicle}) \\ \forall m \in \mathbb{I}(c_{M2}), \mathcal{I}_m(\mathbf{bicycle}) &\subseteq \Delta_1 \setminus \mathcal{I}_m(\mathbf{vehicle}) \\ \forall m \in \mathbb{I}(c_{M1}) \cup \mathbb{I}(c_{M2}), \mathcal{I}_m(\mathbf{bicycle}) &\subseteq \mathcal{I}_m(\mathbf{vehicle}) \cup \mathcal{I}_m(\mathbf{allowed}) \\ \mathbb{I}(\mathbf{allowed}), \mathbb{I}(\mathbf{vehicle}), \mathbb{I}(\mathbf{car}), \mathbb{I}(\mathbf{bicycle}) &\neq \emptyset. \end{aligned}$$

Besides the ones above, a model of the scenario can be thought of requiring two more constraints. Although the formal language as it is defined in 3.1 cannot express them, we show that they can be perfectly captured at a semantic level and therefore that new appropriate symbols might be accordingly added to the syntax.

- $\mathbb{I}(c_{M1} : \mathbf{bicycle}) = \mathbb{I}(c_{M2} : \mathbf{bicycle}) = \{\{a, b\}\}$ ⁹ (c_{M1} and c_{M2} agree on the interpretation of **bicycle**, say, the set of objects $\{a, b\}$);
- $\mathbb{I}(c_{M1} : \mathbf{car}) = \mathbb{I}(c_{M2} : \mathbf{car}) = \{\{c\}\}$ ¹⁰ (c_{M1} and c_{M2} agree on the interpretation of **car**, say, the singleton $\{c\}$).

Let us stipulate that the models m that will constitute our interpretation of contexts identifiers consist of a domain $\Delta_m = \{a, b, c, d\}$ and let us call the sets of all models for \mathcal{L}_0 and \mathcal{L}_1 on this domain respectively \mathbf{M}_0 and \mathbf{M}_1 . Given the restrictions, a ct-model of the scenario can consist then of the domain $\mathcal{P}(\mathbf{M}_0) \cup \mathcal{P}(\mathbf{M}_1)$ and of the function \mathbb{I} s.t.:

- $\mathbb{I}(c_{M1}) = \{m_1, m_2\} \subseteq \mathbf{M}_1$ s.t. $\mathcal{I}_{m_1}(V) = \{a, b, c\}$, $\mathcal{I}_{m_1}(B) = \{a, b\}$, $\mathcal{I}_{m_1}(C) = \{c\}$ and $\mathcal{I}_{m_2}(V) = \{a, b, c, d\}$, $\mathcal{I}_{m_2}(B) = \{a, b\}$, $\mathcal{I}_{m_2}(C) = \{c\}$.
In c_{M1} the concept **vehicle** is interpreted in two possible ways; notice that in this case no exact categorization of d can be enabled;
- $\mathbb{I}(c_{M2}) = \{m_3\} \subseteq \mathbf{M}_1$ s.t. $\mathcal{I}_{m_3}(V) = \{c, d\}$, $\mathcal{I}_{m_3}(C) = \{c\}$, $\mathcal{I}_{m_3}(B) = \{a, b\}$.
In c_{M2} , which is constituted by a single model, the concept **vehicle** strictly contains **car**, and excludes **bicycle**. Notice also that here individual d is categorized as a vehicle.
- $\mathbb{I}(c_R) = \{m \mid \mathcal{I}_m(\mathbf{vehicle}) \subseteq \mathcal{I}_m(\mathbf{car})\}$.
In c_R , concepts **vehicle** and **allowed** get all possible interpretations that keep them disjoint.

We can now get to the main formal characterizations at which we have been aiming in this work.

4.3 Representing conceptual “core” and “penumbra”

What is the part of a denotation of a concept which remains context independent? What is the part which varies instead? “Core” and “penumbral” meaning are formalized in the two following definitions.

Definition 7. ($\mathbf{Core}(\gamma, \xi_1, \xi_2)$)

The “core meaning” of concept γ w.r.t. contexts ξ_1, ξ_2 on language \mathcal{L}_i is defined as:

$$\mathbf{Core}(\gamma, \xi_1, \xi_2) =_{def} \bigcap (\mathbb{I}(\xi_1 : \gamma) \cup \mathbb{I}(\xi_2 : \gamma)).$$

Intuitively, the definition takes just the conjunction of the union of the interpretations of γ in the two contexts. Referring back to Example 2: $\mathbf{Core}(V, c_{M1}, c_{M2}) = \{c\}$, that is, the core of the concept **vehicle** coincides, in those contexts, with the denotation of the concept **car**. The notion of “penumbra” is now easily definable.

⁹ It might be worth recalling that the meaning of a concept in a context is a set of denotations, which we assume to be here, for the sake of simplicity (and in accordance with our intuitions about the scenario), a singleton.

¹⁰ See previous footnote.

Definition 8. ($\mathfrak{P}enumbra(\gamma, \xi_1, \xi_2)$)

The “penumbra” of concept γ w.r.t. contexts ξ_1, ξ_2 on language \mathcal{L}_i is defined as:

$$\mathfrak{P}enumbra(\gamma, \xi_1, \xi_2) =_{def} (\mathbb{I}(\xi_1 : \gamma) \cup \mathbb{I}(\xi_2 : \gamma)) \setminus \mathbf{Core}(\gamma, \xi_1, \xi_2).$$

A “penumbral meaning” is then nothing else but the set of individuals on which the contextual interpretation of the concept varies. Referring back again to Example 2: $\mathfrak{P}enumbra(V, c_{M1}, c_{M2}) = \{a, b, d\}$, that is to say, the penumbra of the concept `vehicle` ranges over those individuals that are not instances of the core of `vehicle`, i.e., the concept `car`. Notice that the definitions are straightforwardly generalizable by induction to formulations with more than two contexts.

5 Related Work

We already showed, in Section 2, how the present proposal relates to work developed in the area of logical modeling of the notion of context. Contexts have been used here in order to propose a different approach to vagueness (especially as it appears in the normative domain). In this section some words will be spent in order to put the present proposal in perspective with respect to some more standard approaches to vagueness, namely approaches making use of fuzzy sets ([26]) or rough sets ([27]).

The most characteristic feature of our approach, with respect to fuzzy or rough set theories, consists in considering vagueness as an inherently semantic phenomenon. Vagueness arises from the referring of a language to structures modeling reality, and not from those structures themselves. That is to say, the truth denotation of a predicate is, in our approach, always definite and crisp, even if multiple. Consequently, no degree of membership is considered, as in fuzzy logic, and no representation of sets in terms of approximations is used, as in rough set theory. Let us use a simple example in order to make this distinction evident. Consider the vague monadic predicate or, to use a description logic terminology, the concept `tall_person`. Fuzzy approaches would determine the denotation of this predicate as a fuzzy set, i.e., as the set of elements with membership degree contained in the interval $]0, 1]$. Standard rough set theory approaches would characterize this denotation not directly, but on the basis of a given partition of the universe (the set of all individuals) and a lower and upper approximation provided in terms of that partition. For instance, a trivial partition might be the one consisting of the following three concepts: `tall > 2m`, `1.60m ≤ tall ≤ 2m`, `tall < 1.60m`. Concept `tall_person` would then be approximated by means of the lower approximation `tall > 2m` (the elements of a set that are definitely also members of the to be approximated set), and the upper approximation `1.60m ≤ tall ≤ 2m` \sqcup `tall > 2m` (the elements of a set that may be also members of the to be approximated set). In this rough set representation, set `1.60m ≤ tall ≤ 2m` constitutes the so called *boundary* of `tall_person`. Within our approach instead, the set `tall_person` can be represented crisply and without approximations. The key feature is that `tall_person` obtains multiple crisp interpretations, at least one for each context: in the context of dutch standards,

concept `tall_person` does not subsume concept `1.60m ≤ tall ≤ 2m`, whereas it does in the context of pygmy standards. According to our approach, vagueness resides then in the contextual nature of interpretation rather than in the concepts themselves¹¹.

It is nevertheless easy to spot some similarities, in particular with respect to rough set theory. The notions of “core” and “penumbra” have much in common with the notions of, respectively, *lower approximation* and *boundary* developed in rough set theory: each of these pairs of notions denotes what is always, and respectively, in some cases, an instance of a given concept. But the characterization of the last pair is based on a partition of the universe denoting the equivalence classes imposed by a set of given known properties. The notions of “core” and “penumbra”, instead, are yielded by the consideration of many contextual interpretations of the concept itself. With respect to fuzzy approaches, notice that sets \mathbf{Core} can be viewed exactly as the sets of instances having a membership degree equal to one, while sets $\mathbf{Penumbra}$ can be viewed as the sets of instances with degree of membership between zero and one. Besides, sets $\mathbf{Penumbra}$ could be partitioned in sets X_n each containing instances that occur in a fixed number n of models constituting the “penumbra”, thus determining a total and, notice, discrete ordering on membership: instances occurring in only one model in the “penumbra” will belong to the denotation of the concept at the minimum degree of membership, while instances occurring in the “core” at the maximum one.

Another relevant feature of our proposal, which we deem worth stressing, consists in the use of a fragment of predicate logic. This allows, first of all, the intra-contextual reasoning to be classical. Furthermore, the use of description logic, even if not yet fully elaborate in this work, allows for its well known inter-

¹¹ A clear position for our thesis can also be found within those analysis of vagueness, developed in the area of philosophical logic, which distinguish between *de re* and *de dicto* views of vagueness ([28]), the first holding that referents themselves are vague and therefore that vagueness constitutes something objective, whereas the second holding that it is the way referents are established that determines vagueness. Fuzzy set approaches lie within a *de re* conception of vagueness, while our approach is grounded on the alternative *de dicto* view (rough sets approaches have instead more to do with insufficient information issues). In philosophical logic, a formal theory has been developed which formalizes this *de dicto* approach to vagueness, the so called *superevaluationism* ([29]). On this view, when interpreting vague terms, we consider the many possible ways in which those terms can be interpreted:

“Whatever it is that we do to determine the ‘intended’ interpretation of our language determines not one interpretation but a range of interpretations. (The range depends on context [...])” ([30]).

As it is evident from Section 3.2, this intuition backs also our semantics. What our approach adds to formal accounts of *superevaluationism* such as [29, 31] consists in the explicit use of contexts as specific formal objects clustering the possible ways terms can be interpreted: contexts are precisely the set of admissible interpretations of the concepts at issue, and as such are endowed with precise formal properties.

esting computability properties to be enabled at the intra-contextual reasoning level, thus making the framework appealing also in this respect.

6 Conclusions

Our aim was to account for a notion of contextual taxonomy, and by means of that, to rigorously characterize the notions of “core” and “penumbra” of a concept, that is to say, to define what is invariant and what is instead context dependent in the meaning of a concept. We did this contextualizing of a standard description logic notion of taxonomy by means of a formal semantics approach to contexts which provides also an account of a variety of forms of contexts interactions.

There is a number of issues which would be worth investigating in future work. First of all, it would be of definite interest to provide formal rigorous comparisons of our framework with:

- Related work in the area of context logics, like especially the *local model semantics* proposed in [16] to which we referred in Section 2.
- Related work in the area of fuzzy or rough sets treatment of conceptual ambiguities ([27, 26]), which have been informally touched upon in Section 5.
- Related work in the area of logic for normative systems specification, and in particular [32] where a modal logic semantics is used to account for expressions such as “A counts as B in context (institution) s”. To this aim, we plan to apply the notion of contextual subsumption relation to modal logic semantics in order to contextualize accessibility relations. For example, it would be interesting to investigate applications to dynamic logic semantics in order to provide a formal account of the contextual meaning of actions: raising a hand in the context of a bidding means something different than rising a hand in the context of a scientific workshop. Some results on this issue has been presented in [33].

Secondly, we would like to enrich the expressivity of our framework considering richer description logic languages admitting also roles (or attributes) constructs. This would allow for a formal characterization of “*contextual terminologies*” in general, enabling the full expressive power description logics are able to provide. A first step along this line has been proposed in [34].

References

1. Dignum, F.: Agents, markets, institutions, and protocols. In: Agent Mediated Electronic Commerce, The European AgentLink Perspective., Springer-Verlag (2001) 98–114
2. Vázquez-Salceda, J., Dignum, F.: Modelling electronic organizations. In V. Marik, J.M., Pechoucek, M., eds.: Proceedings CEEMAS’03. LNAI 2691, Berlin, Springer-Verlag (2003) 584–593

3. Dignum, F.: Abstract norms and electronic institutions. In: Proceedings of the International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA '02), Bologna. (2002) 93–104
4. Vázquez-Salceda, J.: The role of Norms and Electronic Institutions in Multi-Agent Systems. Birkhuser Verlag AG (2004)
- 5.
6. McCarthy, J.: Notes on formalizing contexts. In Kehler, T., Rosenschein, S., eds.: Proceedings of the Fifth National Conference on Artificial Intelligence, Los Altos, California, Morgan Kaufmann (1986) 555–560
7. Akman, V., Surav., M.: Steps toward formalizing context. *AI Magazine* **17** (1996) 55–72
8. Benerecetti, M., Bouquet, P., Ghidini, C.: Contextual reasoning distilled. *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)* **12** (2000) 279–305
9. Searle, J.: *The Construction of Social Reality*. Free Press (1995)
10. Hage, J., Verheij, B.: The law as a dynamic interconnected system of states of affairs. *IJHCS: International Journal of Human-Computer Studies* **51** (1999) 1043–1077
11. Goldman, A.I.: *A Theory of Human Action*. Princeton University Press, Princeton (1976)
12. Hart, H.L.A.: *The Concept of Law*. Clarendon Press, Oxford (1961)
13. Hart, H.L.A.: Positivism and the separation of law and morality. *Harvard Law Review* **71** (1958) 593–629
14. Prakken, H.: *Logical Tools for Modelling Legal Arguments*. Kluwer (1997)
15. Royakkers, L., Dignum, F.: Defeasible reasoning with legal rules. In Nute, D., ed.: *Defeasible Deontic Logic*, Dordrecht, Kluwer (1997) 263–286
16. Ghidini, C., Giunchiglia, F.: Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence* **127** (2001) 221–259
17. Parsons, S., Jennings, N.J., Sabater, J., Sierra, C.: Agent specification using multi-context systems. *Foundations and Applications of Multi-Agent Systems* (2002)
18. Casali, A., Godo, L., Sierra, C.: Graded bdi models for agent architectures. In Leite, J., Toroni, P., eds.: *Proceedings of CLIMA V Workshop*, Lisbon, September. (2004) 18–33
19. Giunchiglia, F., Serafini, L.: Multilanguage hierarchical logics or: How we can do without modal logics. *Artificial Intelligence* **65** (1994) 29–70
20. Giunchiglia, F.: Contextual reasoning. *Epistemologia, special issue on “I Linguaggi e le Macchine”* **16** (1993) 345–364
21. Buvač, S.V., Mason, I.A.: Propositional logic of context. *Proceedings AAAI'93* (1993) 412–419
22. Buvač, S., Buvač, S.V., Mason, I.A.: The semantics of propositional contexts. *Proceedings of the eight ISMIS. LNAI-869* (1994) 468–477
23. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: *The Description Logic Handbook*. Cambridge University Press, Cambridge (2002)
24. McCarthy, J.: Generality in artificial intelligence. *Communications of the ACM* **30** (1987) 1030–1035
25. Casari, E.: *La Matematica della Verità*. Privately distributed (2002)
26. Wygralak, M.: *Vaguely Defined Objects*. Kluwer Academic Publishers (1996)
27. Lin, T.Y., Cercone, N.: *Rough Sets and Data Mining. Analysis of Imprecise Data*. Kluwer Academic Publishers (1997)
28. Varzi, A.: Vague names for sharp objects. In Obrst, L., Mani, I., eds.: *Proceedings of the Workshop on Semantic Approximation, Granularity, and Vagueness, AAAI Press* (2000) 73–78

29. van Fraassen, B.C.: Singular terms, truth-value gaps, and free logic. *Journal of Philosophy* **63** (1966) 481–495
30. Lewis, D.: Many, but almost one. In: *Papers in Metaphysics and Epistemology*, Cambridge University Press (1999) 164–182
31. Fine, K.: Vagueness, truth and logic. *Synthese* **30** (1975) 265–300
32. Jones, A.J.I., Sergot, M.: A formal characterization of institutionalised power. *Journal of the IGPL* **3** (1996) 429–445
33. Grossi, D., Aldewereld, H., Vázquez-Salceda, J., Dignum: Ontological aspects of the implementation of norms in agent-based electronic institutions. Submitted to AISB symposium on Normative Multi-agent Systems, Hatfield, England, April (2005)
34. Grossi, D., Meyer, J.-J. Ch., Dignum, F.: Modal logic investigations in the semantics of counts-as. Submitted to AISB symposium on Mind-Minding Agents, Hatfield, England, April (2005)