

The LKIF Core Ontology of Basic Legal Concepts

Rinke Hoekstra, Joost Breuker, Marcello Di Bello, Alexander Boer
Leibniz Center for Law, University of Amsterdam
breuker@science.uva.nl, hoekstra@uva.nl, mbello@science.uva.nl, aboer@uva.nl

Abstract. In this paper we describe a legal core ontology that is part of a generic architecture for legal knowledge systems, which will enable the interchange of knowledge between existing legal knowledge systems. This *Legal Knowledge Interchange Format*, is under development in the Estrella project and has two main roles: 1) the translation of legal knowledge bases written in different representation formats and formalisms and 2) a knowledge representation formalism that is part of a larger architecture for developing legal knowledge systems. A legal (core) ontology can play an important role in the translation of existing legal knowledge bases to other representation formats, in particular into LKIF as the basis for articulate knowledge serving. We describe the methodology underlying the LKIF core ontology, introduce the concepts it defines, and discuss its use in the formalisation of an EU directive.

Keywords: ontology, legal ontology, legal concept, LKIF, knowledge representation, framework

1. Introduction

In this paper we describe a legal core ontology that is part of a generic architecture for legal knowledge systems, which will enable the interchange of knowledge between existing legal knowledge systems. This *Legal Knowledge Interchange Format* (LKIF), is currently being developed in the Estrella project.¹ LKIF has two main roles: enable the translation between legal knowledge bases written in different representation formats and formalisms and secondly, as a knowledge representation formalism that is part of a larger architecture for developing legal knowledge systems. These use-cases for LKIF bring us to the classical trade-off between tractability and expressiveness, as in e.g. KIF (Knowledge Interchange Format, (Genesereth and Fikes, 1992)). An additional requirement is that LKIF should comply with current Semantic Web standards to enable legal information serving via the web: the core of LKIF consists of a combination of OWL-DL and SWRL, offering a classical hybrid solution. How these two formalisms have to be combined still is an important issue in the development of LKIF, and for details the reader is referred to (Boer et al., 2007).

¹ Estrella is a 6th European Framework project (IST-2004-027665). See also: <http://www.estrellaproject.org>. The views and work reported here are those of the authors.

Proposing the OWL-DL subset of SWRL as its core does not make LKIF a formalism tuned to *legal* knowledge and reasoning: how do we get the ‘L’ into LKIF? To “legalize” LKIF it needs to be constrained in two ways. The first is a *meta-component* that controls the *reasoning* as to gear it to typical legal tasks. For instance, legal assessment and argumentation provide control structures for legal reasoning that put specific demands on the knowledge to be obtained from a legal knowledge base. The second constraint is not specialised to legal reasoning, but to *legal knowledge*. Typical legal concepts may be strongly interrelated and thereby provide the basis for computing equivalencies (paraphrases) and implications. For instance, by representing an obligation as the opposite of a prohibition, a (legal) knowledge system can make inferences that are specialised to these terms. In our view, specialised legal inference should be based on definitions of concepts involved in an ontology. Concept definitions should make all necessary and sufficient interrelationships explicit; the inference engine can then generate all implied consequences.²

A legal ontology can play an important role in the translation of existing legal knowledge bases to other representation formats, in particular into LKIF as the basis for articulate knowledge serving. Similar to a translation between different natural languages, a formal, ‘syntactic’ translation may clash with the semantics implied by the original knowledge representation. An ontology, as representation of the semantics of terms, allows us to keep track of the use of terms in a knowledge base. Furthermore, and more importantly, an ontology can support the process of knowledge acquisition and modelling in legal domains. Defining concepts like ‘norm’, ‘judge’, ‘liability’, ‘document’, ‘claim’, etc. helps to structure the process of knowledge acquisition. Earlier experience, as in e.g. (Breuker and Hoekstra, 2004b; Breuker and Hoekstra, 2004a), suggests a commonsense basis for distinguishing main categories in an ontology for law.

The following sections describe the theoretical and methodological framework against which the LKIF core ontology has been developed (Section 2 and 3). Section 4 describes the different modules of the ontology, and introduces its most important concepts. Section 5 gives an example of how the ontology can be used in the formalisation of a regulation.

² For an ontology cast in OWL-DL these inference engines are description classifiers, e.g. Pellet, <http://pellet.owldl.com/>

2. Frameworks and Ontologies

We adhere to a rather restrictive view on what an ontology should contain: terminological knowledge, i.e. intensional definitions of concepts, represented as classes with which we interpret the world. The distinction between terminological knowledge (T-Box) and assertional knowledge (A-Box) has already been around for a long time. As a rule, terminological knowledge is generic knowledge while assertional knowledge describes the (actual) state of some world: situations and events. However, these asserted states can become generalised into typical patterns related to particular situations. To be sure, if experiences re-occur and have a justifiable structure, it might evidently pay to store these structures as generic descriptions, because they deliver a predictable course of events for free. Eating in a restaurant is a typical example and it served in the Seventies to illustrate the notion of knowledge represented by scripts (Schank and Abelson, 1977) or ‘frames’ (Minsky, 1975). This kind of generic knowledge is indeed rooted in terminological knowledge, but is structured differently. Where ontologies have a taxonomic structure, frames are dominated by mereological and dependency relationships.

Finally, an important reason to distinguish frameworks from ontology proper is that frameworks often imply epistemic roles which require reasoning architectures that go beyond the services provided by OWL-DL reasoners (e.g. meta-level reasoning). It should be noted that frameworks are generic, i.e. they act as pre-specified patterns that get instantiated for particular situations. We have distinguished the following types of frameworks:

Situational frameworks Situational frameworks are stereotypical structures of plans for achieving some goal in a recurrent context. Making coffee may be such a plan. However, the plans may involve transactions in which more than one actor participates. For instance, the definition of **Eating-in-a-restaurant**³ shows the dependencies between actions of clients (ordering, paying) and service personnel (noting, serving) as its major structure. This is the internal structure of the concept, but it usually does not make sense to create class-subclass relations between such frame-like concepts. The **Eating-in-a-restaurant** is not some *natural* sub-class of **Eating**. It refers to some typical model of how eating is put in the context of a restaurant. We can introduce a proliferation of all contexts of eating, such as **Eating-at-home**, **Eating-with-family**, etc. but these contexts do not fundamentally differ, cf. (Bodenreider et al.,

³ In the following all concepts will start with a capital, properties and relations will not

2004; Breuker and Hoekstra, 2004a). In the legal world, such situational frameworks may be pre-scribed in articles of procedural (‘formal’) law. Although *stereotypical* plans (‘customs’) and *prescribed* plans may differ in their justification – rationality vs. authority – their representation is largely analogous. Similarly, legal norms combine generic situation descriptions with some specific state or action. The description is qualified by a deontic term. For instance, the norm that “vehicles should keep to the right of the road” states that the situation in which a vehicle keeps to the right is obliged.

Mereological frameworks Many entities, both objects and processes often have parts: they are *composites*. It is tempting to include a mereological (part-of) view in the definition of a concept. For instance, defining a car as having at least three, and usually four wheels, and at least one motor. However, a full *structural* description of all its parts and connections goes beyond what a car *essentially* is. Mereological frameworks appear under a large diversity of names: structural models, configurations, designs, etc. Arguably, the distinction between a mereological framework and a defining description of a term (ontology) is sometimes be very thin. For instance, if we want to describe a bicycle as distinct from a tricycle, it is necessary to use the cardinality of the wheels as defining properties as these are *central* to the nature of the bicycle. On the other hand, the number of branches a tree might have hardly provides any information as to what a tree *is*.

Epistemological frameworks Inference structures are often represented as epistemological frameworks of interdependencies between reasoning steps. Typical examples are the problem solving methods (PSM) found in libraries of problem solving components (Breuker and Van de Velde, 1994; Motta, 1999; Schreiber et al., 2000)⁴ A problem solving method is not only a break-down of a problem, but also provides control over the making of inferences by assessing success and failure in arriving at the (sub)goals. PSMs have two major components: some method for selecting or generating potential solutions (hypotheses), and some methods for testing whether the solutions hold. Whether they hold may be due to the fact that they satisfy all the specified requirements (constraints) or whether they correspond with (‘explain’) empirical data.

This focus on the *use* of knowledge, its epistemological *status* (e.g. hypothesis vs. conclusion) and the dependencies between distinct steps in a methodology is characteristic for epistemological frameworks. Epis-

⁴ Although the terms ‘reasoning’ and ‘inference’ are often used as more or less synonymous, we want to reserve the term inference for making explicit what is implicit in a knowledge base, given some inference engine.

temological frameworks can be more abstract than PSMs. For instance, the Functional Ontology of Law, which is presented as a core ontology, is an epistemological framework that describes the role of law as a control system in society (Valente, 1995; Breuker et al., 2004).

3. Methodology

The construction of LKIF followed a combination of methodologies for ontology engineering. Already in the mid-nineties, the need for a well-founded methodology was recognised, most notably by (Gruber, 1994; Grüninger and Fox, 1995; Uschold and King, 1995; Uschold and Grüninger, 1996) and later (Fernández et al., 1997). These methodologies follow in the footsteps of earlier experiences in knowledge acquisition, such as the CommonKADS approach (Schreiber et al., 2000) and others, but also considerations from naive physics and cognitive science, such as (Hayes, 1985) and (Lakoff, 1987), respectively.

(Hayes, 1985) describes an approach to the development of a large-scale knowledge base of naive physics. Instead of rather metaphysical top-down construction, his approach starts with the identification of relatively independent *clusters* of closely related concepts. These clusters can be integrated at a later stage, or used in varying combinations allowing for greater flexibility than monolithic ontologies. Furthermore, by constraining (initial) development to clusters, the various – often competing – requirements for the ontology are easier to manage.

Whereas the domain of (Hayes, 1985)’s proposal concerns the relatively well-structured domain of physics, the combination of commonsense and law does not readily provide an obvious starting point for the identification of clusters. In other words, for LKIFcore, we cannot carve-up clusters from a pre-established middle ground of commonsense and legal terms. Furthermore, the field does not provide a relatively stable top level from which top-down development could originate.

In (Uschold and King, 1995), who are the first to use the term ‘middle-out’ in the context of ontology development, it is stressed that the most ‘basic’ terms in each cluster should be defined before moving on to more abstract and more specific terms within a cluster. The notion of this basic level is taken from (Lakoff, 1987), who describes a theory of categorisation in human cognition. Most relevant within the context of ontology engineering (Uschold and King, 1995; Lakoff, 1987, p. 12 and 13) are *basic-level categorisation*, *basic-level primacy* and *functional embodiment*. Categories are organised so that the categories that are cognitively basic are ‘in the middle’ of a taxonomy, generalisation proceeds ‘upwards’ from this basic level and specialisation

proceeds ‘downwards’. Furthermore, these categories are functionally and epistemologically primary with respect to (amongst others) knowledge organisation, ease of cognitive processing and ease of linguistic expression. Basic level concepts are used automatically, unconsciously, and without noticeable effort as part of normal functioning. They have a different, and more important psychological status than those that are only thought about consciously.

For the purpose of the LKIF ontology, we have made slight adjustments to the methodology of (Hayes, 1985; Uschold and Grü-nin-ger, 1996). We established design criteria for the development of the LKIF ontology based on (Gruber, 1993; Uschold and Grü-nin-ger, 1996). These criteria were implemented throughout the following phases: identify *purpose and scope*, *ontology capture* and *coding, integration* with existing ontologies and *evaluation*. The following section describes how these phases have materialised in the context of LKIF Core. Furthermore, an example in which the ontology is put to use is described in section 5.

4. Modules & Outline

This section describes how the methodology described in the previous section was applied to the development of LKIF Core. We first describe the building and clustering phase, followed by a discussion of the existing ontologies we considered for inclusion, and a description of the concepts defined in the different modules of the ontology.

4.1. ONTOLOGY CAPTURE

The LKIF Core ontology should contain ‘basic concepts of law’. It is dependent on the (potential) users what kind of vocabulary is aimed at. We have identified three main groups of users: *citizens*, *legal professionals* and *legal scholars*. Although legal professionals use the legal vocabulary in a far more precise and careful way than laymen, it appears that for most of these terms there is still a sufficient common understanding to treat them more or less as similar (Lame, 2006). Nonetheless, a number of basic terms have a specific legal-technical meaning, such as ‘liability’ and ‘legal fact’. We included these technical terms because they might capture the ‘essential’, abstract meaning of terms in law, but also because these terms might be used to organise more generally understood legal terms.

The Estrella consortium includes representatives of the three kinds of experts. Each partner was asked to supply their ‘top-20’ of legal

concepts. Combined with terms we collected from literature (jurisprudence and legal text-books) we obtained a list of about 250 terms. As such a number is unmanageable as a basic set for modelling, we asked partners to assess each term from this list on five scales: level of *abstraction*, *relevance* for the legal domain, the degree to which a term is *legal* rather than *common-sense*, the degree to which a term is a *common legal term* (as opposed to a term that is specific for some sub-domain of law), and the degree to which the expert thinks this term should be *included* in the ontology. The resulting scores were used to select an initial set of 50 terms plus those re-used from other ontologies (see section 4.2), and formed the basis for the identification of clusters and the development of the LKIF Core ontology.

4.2. OTHER ONTOLOGIES

We expected to be able to reuse terms and definitions from existing core or upper ontologies that contain legal terms, as e.g. listed in (Casanovas et al., 2006). Unfortunately, it turned out that the amount of re-use and inspiration was rather limited. The following core ontologies for law were consulted, both for their potential contribution for creating a coherent top for LKIF Core, and specifically for legal terms already represented.

The intentional nature of the core concepts for the LKIF ontology (see e.g. sections 4.3.2,4.3.3) emphasises the distinction with other more (meta)physically inclined top ontologies such as SUMO⁵, Sowa's upper ontology (Sowa, 2000) and DOLCE⁶ (Gangemi et al., 2002)), but shows similarities with the distinction between *intentional*, *design* and *physical* stances described in (Dennett, 1987). As some of these top- or upper ontologies (SUMO, Sowa) do not have a common-sense basis – e.g. mental and social entities are poorly represented – they could neither be used as a top for LKIF Core, nor as a source of descriptions of legal terms. The upper part of the CYC⁷ ontology and DOLCE (Gangemi et al., 2003; Massolo et al., 2002) are claimed to have a common-sense view, but this common-sense view is rather based upon personal intuition than on empirical evidence. LRI-Core on the other hand is to a large extent based upon empirical studies in cognitive science, and is intended as a core ontology for law. However, the number of typical legal concepts in this legal core ontology is disappointingly small. Nonetheless, its top structure appeared to be valuable in constructing LKIF as is further described in Section 4. The Language for Legal Discourse (McCarty,

⁵ Suggested Upper Merged Ontology; <http://ontology.teknowledge.com>

⁶ Descriptive Ontology for Linguistic and Cognitive Engineering; <http://www.10a-cnr.it/DOLCE.html>

⁷ www.cyc.com

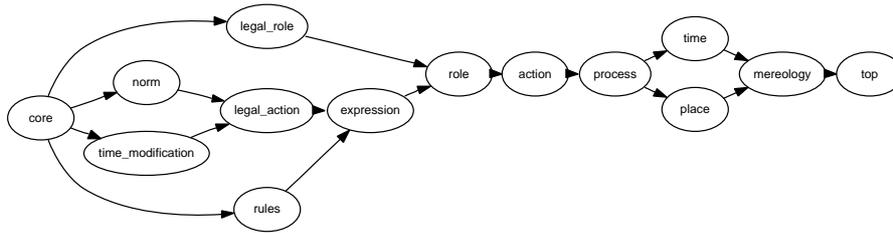


Figure 1. Dependencies between LKIFCore modules.

1989, LLD) is a first attempt to define legal concepts in the context of legal reasoning, using formulae and rules. Properly speaking, LLD is not an ontology but a framework but it is a relatively rich source for legal terms and their definitions. The Core Legal Ontology (CLO) is used to support the construction of legal domain ontologies (Gangemi et al., 2005). CLO organises legal concepts and relations on the basis of formal properties defined in DOLCE+. Although purpose and layers are largely similar to those of LRI-Core, the top structures differ considerably.

4.3. ONTOLOGY MODULES

The list of terms and insights from the requirements-phase resulted in a collection of ontology modules, each of which represents a relatively independent cluster of concepts: *expression*, *norm*, *process*, *action*, *role*, *place*, *time* and *mereology* (Breuker et al., 2006; Breuker et al., 2007). The concepts in these clusters were formalised using OWL-DL in a middle-out fashion: for each cluster the most central concepts were represented first.⁸

Discussions, further literature study and the consideration of existing ontologies, led to an extension of the original set of clusters to 14 modules (see Figure 1), each of which describes a set of closely related concepts from both legal and commonsense domains. Nonetheless, we maintained the original views used to identify the clusters, as the explanations and justifications are still valid and applicable to the current version of the ontology. We can distinguish three layers in the ontology: the *top* level (Section 4.3.1), the *intentional* level (Section 4.3.2) and the *legal* level (Section 4.3.3).

4.3.1. First Things First: The top-level

The description of any legally relevant fact, event or situation requires a basic conceptualisation of the context in which these occur: the backdrop, or canvas, that is the physical world. Fundamental notions such as

⁸ We used both TopBraid Composer (<http://www.topbraidcomposer.com>) and Protégé 3.2/4.0 (<http://protege.stanford.edu>).

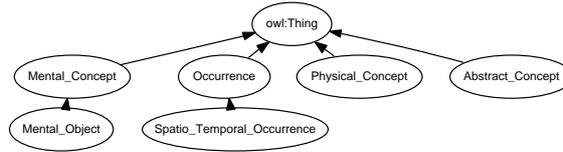


Figure 2. Concepts defined in the Top module.

location, time, parthood and change are indispensable in a description of even the simplest legal account. The top level clusters of the ontology provide (primitive) definitions of these notions, which are consequently used to define more intentional and legal concepts in other modules. The most general classes of the LKIF ontology are borrowed from LRI Core. We distinguish between mental, physical and abstract concepts, and occurrences (Figure 2).

Mereological relations allow us to define parts and wholes, allow for expressing a systems-oriented view on concepts, such as functional decompositions, and containment (Figure 3). Furthermore, they form the basis for definitions of *places* (location) and moments and intervals in *time*.

The ontology for places in LKIF Core is based on the work of (Donnelly, 2005), and adopts a distinction between *relative* places and *absolute* places, which goes back to Isaac Newton. Whereas a relative place is defined by reference to some thing, absolute places are part of absolute space and have fixed spatial relations with other absolute places. See figure 3 for an overview of concepts defined in the place module. A `Location_Complex` is a set of places that share a reference location.

Of the properties defined in this module, `meet` is the most basic as it is used to define many of the other properties such as `abut`, `cover`, `coincide` etc. See (Breuker et al., 2007; Donnelly, 2005) for a more in depth discussion of these and other relations. The current version of the ontology of places does not define concepts and relations that can be used to express direction and orientation.

Closely related to the theory of places of (Donnelly, 2005) is Allen’s theory of time (Allen, 1984; Allen and Ferguson, 1994). We adopt his theory, and distinguish between the basic concepts of `Interval` and `Moment`. Intervals have an extent (duration) and can contain other intervals and moments. Moments are points in time, they are atomic and do not have a duration or contain other temporal occurrences (see figure 4).

The relations between temporal occurrences are what defines time. Like (Donnelly, 2005), (Allen, 1984) adopts the `meet` relation to define two immediately adjacent temporal occurrences. We call this relation

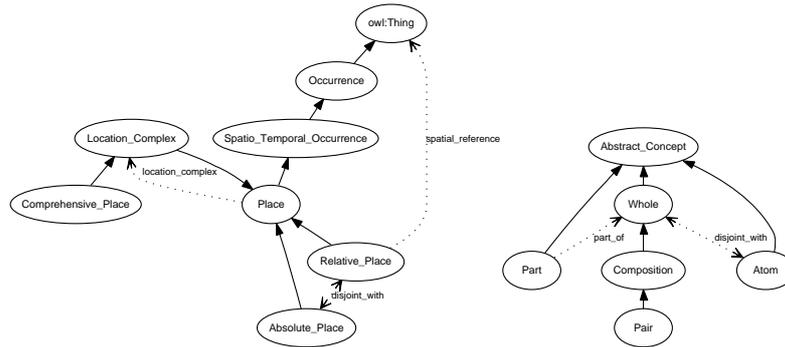


Figure 3. Place and Mereology related concepts.

`immediately_before`, as the temporal meet relation holds only in one direction, and is asymmetric. The property is used to define other temporal relations such as `before`, `after`, `during`, etc.

With these classes and properties in hand, we introduce concepts of (involuntary) change. The process ontology relies on descriptions of time and place for the representation of duration and location of changes. A `Change` is essentially a difference between the situation before and after the change. It can be a functionally coherent aggregate of one or more other changes. More specifically, we distinguish between `Initiation`, `Continuation` and `Termination` changes.

Changes that occur according to a certain recipe or procedure, i.e. changes that follow from causal necessity are `Processes`; they introduce causal propagation. Contrary to changes, processes are bound in time and space: they have duration and take place at a time and place. We furthermore distinguish `Physical_Processes` which operate on `Physical_Objects`. Furthermore, at this level we do not commit to a particular theory of causation or causal propagation.

4.3.2. The Intentional Level

Legal reasoning is based on a common sense model of intelligent behaviour, and the prediction and explanation of intelligent behaviour. It is after all only behaviour of rational agents that can be effectively influenced by the law. The modules at the intentional level include concepts and relations necessary for describing this behaviour (i.e. `Actions` undertaken by `Agents` in a particular `Role`) which are governed by law. Furthermore, it introduces concepts for describing the mental state of these agents, e.g. their `Intention` or `Belief`, but also communication between agents by means of `Expressions`.

The class of agents is defined as the set of things which can be the `actor` of an intentional action: they perform the action and are

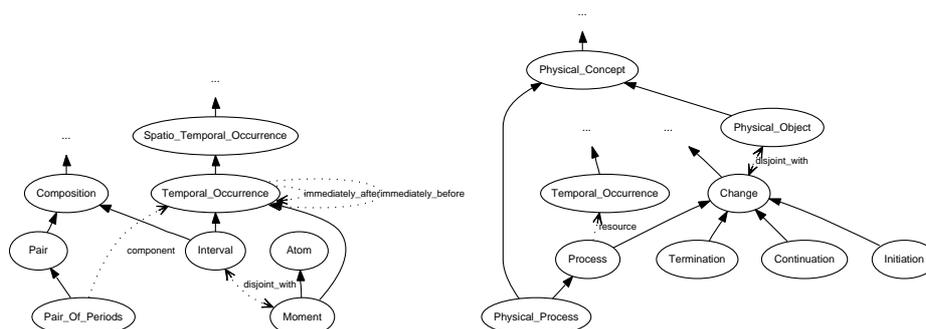


Figure 4. Concepts related to time and change.

potentially liable for any effects caused by the action (see figure 5). Actions are processes, they are the changes performed by some agent who has the intention of bringing about the change. Because actions are processes they can become part of causal propagation, allowing us to reason backwards from effect to agent. Actions can be creative in that they initiate the coming into existence of some thing, or the converse. Also, actions are often a direct *reaction* to some other action (see figure 5).

The agent is the medium of some intended outcome of the action: an action is always intentional. The intention held by the agent, usually bears with it some expectation that the intended outcome will be brought about: the agent believes in this expectation. The actions an agent is expected or allowed to perform are constrained by the *competence* of the agent, sometimes expressed as *roles* assigned to the agent.

We distinguish between **persons**, individual agents such as “Joost Breuker” and “Pope Benedict XVI”, and **Organisations**, aggregates of other organisations or persons which acts ‘as one’, such as the “Dutch Government” and the “Sceptics Society”. **Artefacts** are physical objects designed for a specific purpose, i.e. to perform some **Function** as instrument in a specific set of actions such as “Hammer” and “Atlatl”⁹. Persons are physical objects as well, but are not designed (though some might hold the contrary) and are subsumed under the class of **Natural_Objects**. Note that natural objects can function as tools or weapons as well, the typical example being a stone, but are not designed for that specific purpose.

The notion of roles has played an important part in recent discussions on ontology (Steimann, 2000; Masolo et al., 2004; Guarino and

⁹ An atlatl is a tool that uses leverage to achieve greater velocity in spear-throwing, see <http://en.wikipedia.org/wiki/Atlatl>

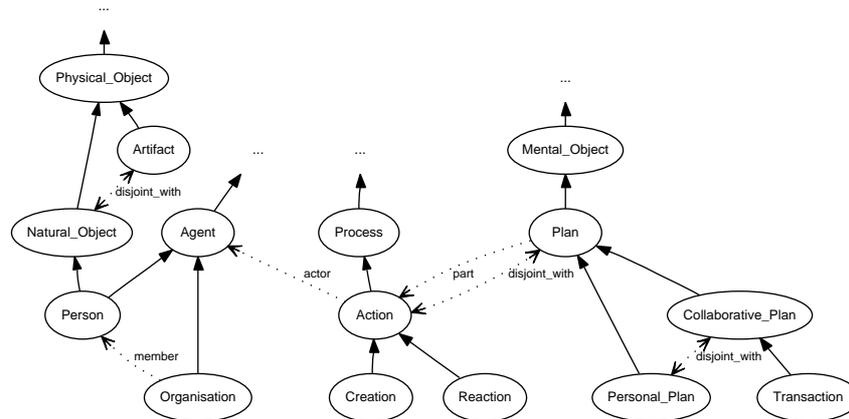


Figure 5. Actions, agents and organisations.

Welty, 2002). Roles not only allow us to categorise objects according to their prototypical use and behaviour, they also provide the means for categorising the behaviour of other agents. They are a necessary part of making sense of the social world and allow for describing social organisation, prescribe behaviour of an agent within a particular context, and recognise deviations from ‘correct’ or normal behaviour. Indeed, roles and actions are closely related concepts: a role defines some set of actions that can be performed by an agent, but is conversely defined by those actions. Roles specify standard or required properties and behaviour (see figure 6). The role module captures the roles and functions that can be played and held by agents and artefacts respectively, and focuses on *social* roles, rather than traditional thematic or relational roles.

A consequence of the prescriptive nature of roles is that agents connect expectations of behaviour to other agents: intentions and expectations can be used as a model for intelligent decision making and planning¹⁰. It is important to note that there is an *internalist* and an *externalist* way to use intentions and expectations. The external observer can only ascribe intentions and expectations to an agent based on his observed actions. The external observer will make assumptions about what is *normal*, or apply a *normative* standard for explaining the actions of the agent.

¹⁰ Regardless of whether it is a psychologically plausible account of decision making. Daniel Dennett’s notion of the *Intentional Stance* is interesting in this context (cf. (Dennett, 1987)). Agents may do no more than occasionally apply the stance they adopt in assessing the actions of others to themselves.

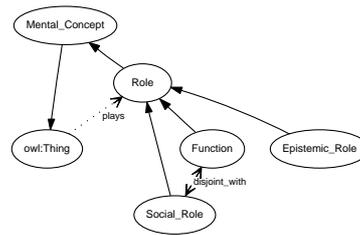


Figure 6. Roles.

The expression module covers a number of representational primitives necessary for dealing with **Propositional_Attitudes** (viz. (Dahllöf, 1995)). Many concepts and processes in legal reasoning and argumentation can only be explained in terms of propositional attitudes: a relational mental state connecting a person to a **Proposition**. However, in many applications of LKIF the attitude of the involved agents towards a proposition will not be relevant at all. For instance, fraud detection applications will only care to distinguish between potentially contradictory observations or expectations relating to the same propositional content. Examples of propositional attitudes are **Belief**, **Intention**, and **Desire**. Each is a component of a mental model, held by an **Agent**.

Communicated attitudes are held towards expressions: propositions which are externalised through some medium. **Statement**, **Declaration**, and **Assertion** are expressions communicated by one agent to one or more other agents. This classification is loosely based on Searle (cf. (Searle and Vanderveken, 1985)). A prototypical example of a medium in a legal setting is e.g. the **Document** as a bearer of legally binding (normative) statements.

When propositions are used in reasoning they have an epistemic role, e.g. as **Assumption**, **Cause**, **Expectation**, **Observation**, **Reason**, **Fact** etc. The role a proposition plays within reasoning is dependent not only on the kind of reasoning, but also the level of trust as to the validity of the proposition, and the position in which it occurs (e.g. hypothesis vs. conclusion). In this aspect, the expression module is intentionally left under-defined. A rigorous definition of propositional attitudes relates them to a theory of reasoning and an argumentation theory. The argumentation theory is supplied by an argumentation ontology. The theory of reasoning depends on the type of reasoning task (assessment, design, planning, diagnosis, etc.) LKIF is used in, and should be filled in (if necessary) by the user of LKIF .

Evaluative_Attitudes express an evaluation of a proposition with respect to one or more other propositions, they express e.g. an evaluation, a value statement, value judgement, evaluative concept, etc. I.e. only

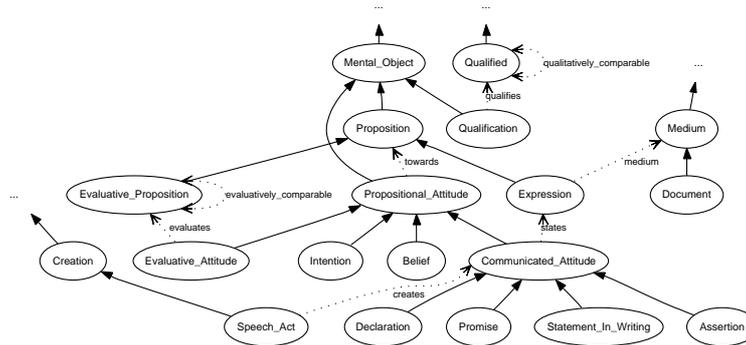


Figure 7. Propositions, Attitudes and Expressions.

the type of qualification which is an attitude towards the thing being evaluated, and not for instance the redness of a rose, as in (Gangemi et al., 2002) and others. Of special interest is the **Qualification**, which is used to define norms based on (Boer et al., 2005). Analogous to the evaluative attitude, a qualification expresses a judgement. However, the subject of this judgement need not be a proposition, but can be any complex description (e.g. a situation).

4.3.3. The Legal Level

Legally relevant statements are created through public acts by both natural and legal persons. The legal status of the statement is dependent on both the kind of agent creating the statement, i.e. **Natural_Person** vs. a **Legislative_Body**, and the rights and powers attributed to the agent through mandates, assignments and delegations. At the legal level, the LKIF ontology introduces a comprehensive set of legal agents and actions, rights and powers (a modified version of (Sartor, 2006; Rubino et al., 2006)), typical legal roles, and concept definitions which allow us to express normative statements as defined in (Boer et al., 2005; Boer, 2006; Boer et al., 2007).

The **Norm** is a statement combining two performative meanings: it is *deontic*, in the sense that it is a qualification of the (moral or legal) acceptability of some thing, and it is *directive* in the sense that it commits the speaker to bringing about that the addressee brings about the more acceptable thing (cf. (Nuyts et al., 2005)), presumably through a sanction. These meanings do not have to occur together. It is perfectly possible to attach a moral qualification to something without directing anyone, and it is equally possible to issue a directive based on another reason than a moral or legal qualification (e.g. a warning).

A norm applies to (or **qualifies**) a certain situation (the **Qualified** situation), allows a certain situation – the **Obliged** situation or **Allowed**

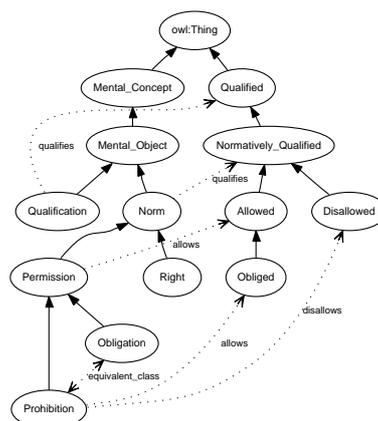


Figure 8. Qualifications and Norms

situation – and disallows a certain situation – the Prohibited or Disallowed situation, see Figure 8. The obliged and prohibited situation are both subsumed by the situation to which the norm applies. Besides that they by definition form a complete partition of the case to which the norm applies, i.e. all situation to which the norm applies are either a mandated case or a prohibited situation. This is true of the obligation and the prohibition: they are simply two different ways to put the same thing into words. The permission is different in that it allows something, but it does not prohibit anything. The logical complement of the mandated situation is here an opposite qualified situation, about which we know only that it cannot be obliged.

5. Putting the ontology to use: the Traffic domain

The LKIF ontology not only provides a theoretical understanding of the legal domain, but its primary use in practice is as a tool to facilitate knowledge acquisition, exchange and representation: i.e. to formalise pieces of existing legislation. We evaluated the use of the ontology by formalising the EU Directive 2006/126 on driving licences,¹¹ a relatively straightforward regulation, in which at least two types of normative statement are recognisable—definitional and deontic.

An example of a *definitional statement* from the EU directive is:

Art. 4(2) **Category AM:** *Two-wheel vehicles or three-wheel vehicles with a maximum design speed of not more than 45 km/h.*

¹¹ The text is available on-line at <http://eur-lex.europa.eu/>.

The mereo module of the ontology along with a qualified cardinality restriction (available with OWL 1.1) allows us to express that AM vehicles have two or three wheels:

$$\text{AM} \sqsubseteq 2\text{composed_of.Wheel} \sqcup 3\text{composed_of.Wheel}.$$

Modelling ‘design speed not more than 45 km/h’ is more challenging as it requires us to represent the rather common sense domain of speeds, distances etc. Of course, one could introduce the datatype property `designSpeed` and require its value be expressed in km/h. This choice, however, would not make justice of the conceptual complexity involved in ‘design speed not more than 45 km/h’, which contains reference to several notions: unit of measurement, number, designed speed, and a no-more relation. In fact, ‘design speed not more than 45 km/h’ can be rendered by imposing an *linear ordering* relation `less-than` on the different (instances of the) subclasses of the class `DesignSpeed`.¹² The ordering allows us to define the class of those `DesignSpeeds` with a value not exceeding some `N45`—i.e., $\forall \text{less-than.DesSpeed-km-h-45}$.

Let us now look at an example of a *deontic statement*:

Art. 4(2) *The minimum age for category AM is fixed at 16 years.*

Art. 4(2) expresses an obligation whose logical form can be rendered by the implication:

*If x is driving a AM vehicle, then x **must** be at least 16 years old.*

To fix some terminology, the antecedent is the *context* to which the obligation applies; the consequent (minus the deontic operator **must**) is the *content* of the obligation itself (what the obligations prescribes it ought to be the case). Consistently with this analysis, the LKIF ontology defines obligations as classes (see Section 4.3.3).

In our case, art. 4(2) allows the situation `DriverAM` \sqcap `DriverOlderThan16` and forbids `DriverAM` \sqcap \neg `DriverOlderThan16`. Suppose that the classes `DriverOlderThan16` and `DriverAM` have already been defined.¹³ To model the obligation that drives of AM vehicles must be at least the 16 years older, we introduced the obligation-type class `MinAgeAM` as follows:

¹² The ordering is linear—i.e., reflexive, antisymmetric, transitive and total—since it mirrors the ordering of the natural numbers. For whenever $n \leq m$, we have that `DesignSpeed-km-h-n(a)` `less-than` `DesignSpeed-km-h-m(b)`, with a, b instances.

¹³ The class `DriverOlderThan16` can be defined by using a `more-than` ordering relation, roughly along the same lines as the class $\forall \text{less-than.DesSpeed-km-h-45}$. The class `DriverAM` can be easily defined.

$\text{MinAgeAM} \sqsubseteq \forall \text{allows.}(\text{DriverAM} \sqcap \text{DriverOlderThan16}).$
 $\text{MinAgeAM} \sqsubseteq \exists \text{allows.}(\text{DriverAM} \sqcap \text{DriverOlderThan16}).$
 $\text{MinAgeAM} \sqsubseteq \forall \text{disallows.}(\text{DriverAM} \sqcap \neg \text{DriverOlderThan16}).$
 $\text{MinAgeAM} \sqsubseteq \exists \text{disallows.}(\text{DriverAM} \sqcap \neg \text{DriverOlderThan16}).$

Other deontic operators, such as permission or prohibition, can be accounted in an alike manner (see (Boer et al., 2007)). Notwithstanding the parsimony of this type of definition, using the LKIF ontology to model normative statements proves to be rather straightforward. Of course, a specialised modelling environment for legislative drafters would need to provide a shorthand for such standard OWL definitions.¹⁴

The representation of art. 4(2) suggests the LKIF ontology be augmented with a module taking care of quantities, units of measurement, numbers, fractions, mathematical operations, and the like. This is crucial not only for the EU Directive 2006/126, in which most definitional statements contain quantitative features of vehicles (e.g., power, cylinder capacity); quantities and calculations play a central role in any legislative text. Note, however, that the LKIF ontology can only provide a *purely terminological* account, without being able to do mathematical computations. This is unavoidable, given that OWL is a purely logical language. We are currently investigating whether we can import an existing OWL ontology dealing with measurements, such as PHYSYS/SUMO or from the Ontolingua server¹⁵.

6. Discussion

As LKIF Core was developed by a heterogeneous group of people, we specified a number of conventions to uphold during the representation of terms identified in the previous phases (See (Breuker et al., 2007)). One of these is that classes should be represented using necessary & sufficient conditions as much as possible (i.e. by means of `equivalentClass` statements). Using such ‘complete’ class definitions ensures the ability to infer the type of individuals; this does not hold for partial class definitions (using only necessary conditions).

In retrospect, this convention turned out to pose severe problems for existing OWL-DL and OWL 1.1 reasoners as their performance is significantly affected by the generic concept inclusion axioms (GCI):

¹⁴ See e.g. the SEAL project, <http://www.leibnizcenter.org/project/current-projects/seal>

¹⁵ See <http://www.ksl.stanford.edu/software/ontolingua/>

axioms where the left-hand side of a `subClassOf` statement is a complex class definition. These axioms are abundant when defining classes as equivalent to e.g. `someValuesFrom` restrictions and in combination with lots of inverse property definitions, this creates a large completion graph for DL reasoners¹⁶ As result of these findings, the LKIF ontology has undergone a significant revision since its initial release.

Using LKIF Core in practice, as e.g. in the traffic example, points to the traditional knowledge-acquisition bottle-neck: for any formal representation of any domain, one needs to build formal representations of adjoining domains. As has been said, this can be largely overcome by including specialised foundational or domain ontologies in a representation based on the LKIF ontology provided that the quality of these ontologies is sufficient. Depending on availability we might consider providing a library of ‘compatible’ ontologies to users of LKIF Core. This will be of especial use when the ontology vocabulary will be adopted for expressing the LKIF vendor models that will be developed within ESTRELLA.

With respect to coverage of the legal domain, the purpose of the study outlined in Section 4.1 is more ambitious than only the selection of the most basic terms for describing law, but time and effort constraints make it that we could only consider a small pool of referents. The list of terms will be subjected to a more rigorous empirical study, whereby we will consult a group of legal professionals (taking courses in legal drafting), and law students. These empirical studies are planned in the sideline of ESTRELLA. By applying statistical cluster analysis, we might be able to identify the properties of the scales used (e.g. are they independent?) and whether the statistical clusters have some resemblance to the clusters we have identified based on theoretical considerations. The results of this analysis will be used to evaluate the ontology compared to the requirements we identified in the previous chapters.

The LKIF ontology is available online as separate but interdependent OWL-DL files, and can be obtained from the ESTRELLA website at <http://www.estrellaproject.org/lkif-core>. This website also provides links to online documentation and relevant literature.

References

- Allen, J. (1984). Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154.

¹⁶ Thanks to Taowei David Wang for pointing this out, see <http://lists.owlidl.com/pipermail/pellet-users/2007-February/001263.html>

- Allen, J. F. and Ferguson, G. (1994). Actions and events in interval temporal logic. *Journal of Logic and Computation*, 4(5):531–579.
- Bodenreider, O., Smith, B., and Burgun, A. (2004). The ontology-epistemology divide: a case study in medical terminology. In Varzi, A. and Vieu, L., editors, *Formal ontology in Information Systems*, pages 185–198. IOS-Press, Amsterdam.
- Boer, A. (2006). Note on production rules and the legal knowledge interchange format. Technical report, Leibniz Center for Law, Faculty of Law, University of Amsterdam.
- Boer, A., Gordon, T. F., van den Berg, K., Di Bello, M., Förhécz, A., and Vas, R. (2007). Specification of the legal knowledge interchange format. Deliverable 1.1, Estrella.
- Boer, A., van Engers, T., and Winkels, R. (2005). Mixing legal and non-legal norms. In Moens, M.-F. and Spyns, P., editors, *Jurix 2005: The Eighteenth Annual Conference*, Legal Knowledge and Information Systems, pages 25–36, Amsterdam. IOS Press.
- Breuker, J., Boer, A., Hoekstra, R., and van den Berg, K. (2006). Developing content for lkif: Ontologies and frameworks for legal reasoning. In van Engers, T. M., editor, *Legal Knowledge and Information Systems. JURIX 2006: The Nineteenth Annual Conference*, volume 152 of *Frontiers in Artificial Intelligence and Applications*.
- Breuker, J. and Hoekstra, R. (2004a). Core concepts of law: taking common-sense seriously. In Varzi, A. and Vieu, L., editors, *Proceedings of Formal Ontologies in Information Systems (FOIS-2004)*, pages 210–221. IOS-Press.
- Breuker, J. and Hoekstra, R. (2004b). Epistemology and ontology in core ontologies: FOLaw and LRI-Core, two core ontologies for law. In *Proceedings of EKAW Workshop on Core ontologies*, [Http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/](http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/). Ceur.
- Breuker, J., Hoekstra, R., Boer, A., van den Berg, K., Rubino, R., Sartor, G., Palmirani, M., Wyner, A., and Bench-Capon, T. (2007). OWL ontology of basic legal concepts (LKIF-Core). Deliverable 1.4, Estrella.
- Breuker, J., Valente, A., and Winkels, R. (2004). Use and reuse of legal ontologies in knowledge engineering and information management. *Artificial Intelligence and Law*, (to appear in special issue on Legal Ontologies).
- Breuker, J. and Van de Velde, W., editors (1994). *CommonKADS Library for Expertise Modelling*. IOS Pres.
- Casanovas, P., Casellas, N., Vallbe, J.-J., maria Poblet, Benjamins, R., Blazquez, M., Pena, R., and Contreras, J. (2006). Semantic web: a legal case study. In Davies, J., Studer, R., and Warren, P., editors, *Semantic Web Technologies*. Wiley.
- Dahllöf, M. (1995). On the semantics of propositional attitude reports.
- Dennett, D. (1987). *The Intentional Stance*. MIT-Press.
- Donnelly, M. (2005). Relative places. *Applied Ontology*, 1(1):55–75.
- Fernández, M., Gómez-Pérez, A., and Juristo, N. (1997). Methontology: from ontological art towards ontological engineering. In *Proceedings of the AAAI97 Spring Symposium Series on Ontological Engineering*, pages 33–40, Stanford, USA.
- Gangemi, A., Guarino, N., Masolo, C., and Oltramari, A. (2003). Sweetening WORDNET with DOLCE. *AI Magazine*, 24:13–24.
- Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., and Schneider, L. (2002). Sweetening ontologies with DOLCE. In Gomez-Perez, A. and Benjamins, V., editors, *Proceedings of the EKAW-2002*, pages 166–181. Springer.

- Gangemi, A., Sagri, M., and Tiscornia, D. (2005). A constructive framework for legal ontologies. In Benjamins, V., Casanovas, P., Breuker, J., and Gangemi, A., editors, *Law and the Semantic Web*, pages 97–124. Springer Verlag.
- Genesereth, M. and Fikes, R. (1992). Knowledge interchange format, version 3.0 reference manual. Technical Report Logic-92-1, Computer Science Department, Stanford University.
- Gruber, T. (1994). Towards principles for the design of ontologies used for knowledge sharing. In Guarino, N. and Poli, R., editors, *Formal Ontology in Conceptual Analysis and Knowledge Representation*, pages –. Kluwer Academic Publishers. also: Technical Report KSL 93-40, Knowledge Systems Laboratory, Stanford University.
- Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5:199–220.
- Grüninger, M. and Fox, M. S. (1995). Methodology for the design and evaluation of ontologies. In *IJCAI'95, Workshop on Basic Ontological Issues in Knowledge Sharing*.
- Guarino, N. and Welty, C. (2002). Evaluating ontological decisions with OntoClean. *Communications of the ACM*, 45(2):61–65.
- Hayes, P. J. (1985). The second naive physics manifesto. In Hobbs, J. R. and Moore, R. C., editors, *Formal Theories of the Common Sense World*, pages 1–36. Ablex Publishing Corporation, Norwood.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things*. University of Chicago Press.
- Lame, G. (2006). Using NLP techniques to identify legal ontology components: Concepts and relations. *Artificial Intelligence and Law*. this issue.
- Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Ferrario, R., Gangemi, A., and Guarino, N. (2004). Social roles and their descriptions. In *Proceedings of Knowledge Representation Workshop*.
- Massolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A., and Schneider, L. (2002). The WonderWeb foundational ontologies: preliminary report. Technical Report Deliverable D17, version 2, ISTC-CNR (Italy).
- McCarty, T. (1989). A language for legal discourse I. basic structures. In *Proc. of the Second International Conference on AI and Law*, pages 180–189, Vancouver. Acm.
- Minsky, M. (1975). A framework for representing knowledge. In Winston, P. H., editor, *The psychology of Computer Vision*, pages 211–277, New York. McGraw-Hill.
- Motta, E. (1999). *Reusable Components for Knowledge Modelling*. FAIA-Series. IOS Press, Amsterdam NL.
- Nuyts, J., Byloo, P., and Diepeveen, J. (2005). On deontic modality, directivity, and mood.
- Rubino, R., Rotolo, A., and Sartor, G. (2006). An owl ontology of fundamental legal concepts. In van Engers, T. M., editor, *Legal Knowledge and Information Systems. JURIX 2006: The Nineteenth Annual Conference*, volume 152 of *Frontiers of Artificial Intelligence and Applications*. IOS Press.
- Sartor, G. (2006). Fundamental legal concepts: A formal and teleological characterisation. Technical report, European University Institute, Florence / Cirsfid, University of Bologna.
- Schank, R. and Abelson, R. (1977). *Scripts, Plans Goals and Understanding*. Lawrence Erlbaum, New Jersey.

- Schreiber, G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N., Van den Velde, W., and Wielinga, B. (2000). *Knowledge Engineering and Management: The CommonKADS Methodology*. MIT Press.
- Searle, J. and Vanderveken, D. (1985). *Foundations of Illocutionary Logic*. Cambridge University Press, Cambridge.
- Sowa, J. F. (2000). *Knowledge Representation: Logical Philosophical, and Computational Foundations*. Brooks Cole Publishing Co, Pacific Grove, CA.
- Steimann, F. (2000). On the representation of roles in object-oriented and conceptual modelling. *Data and Knowledge Engineering*, 35:83–106.
- Uschold, M. and Grüninger, M. (1996). Ontologies: principles, methods, and applications. *Knowledge Engineering Review*, 11(2):93–155.
- Uschold, M. and King, M. (1995). Towards a methodology for building ontologies. In *Workshop on Basic Ontological Issues in Knowledge Sharing, IJCAI-95*, Montreal, Canada.
- Valente, A. (1995). *A Modelling Approach to Legal Knowledge Engineering*. PhD thesis, University of Amsterdam.