Legal Ontologies: A Functional View

Joost Breuker[1] and **Andre Valente**[2] and **Radboud Winkels**[1]
[1]Dept. of Computer Science & Law, University of Amsterdam
PO Box 1030, 1000 BA Amsterdam, the Netherlands
e-mail: {breuker,winkels}@lri.jur.uva.nl
[2]Information Sciences Institute, University of Southern California
4676 Admiralty Way, Marina del Rey, CA 90292, USA
e-mail: valente@isi.edu

Abstract

The ontology of law that we present here is one 'in between' a completely general top ontology and ontologies of legal domains. It adopts a *functional* perspective — that is, the Law is analysed and interpreted through a functional point of view. We propose a number of primitive functions of legal sources and corresponding categories of legal knowledge: *normative knowledge*, *world knowledge*, *responsibility knowledge*, *reactive knowledge*, *creative knowledge* and *meta-legal knowledge*. Together, these functions realize the main function of the legal system: to regulate social behaviour. The ontology is illustrated with examples from Dutch labour law.

Ontology and ontological engineering

In the new-born discipline of ontological engineering, the term "ontology" has another meaning than as originally conceived in philosophy. Wolff (1679 – 1754) elaborated on Leibniz's systematization and introduced the term 'ontology' to distinguish the essential from the incidental, the real thing from its appearance.[1] Ontological engineering is hardly interested in the question of existence (or its representation, see e.g. (Hirst, 1991) for an exception), but it is concerned with the specification of content of knowledge. Epistemological or semantic engineering may have been more appropriate terms.

The 'engineering' denotes the fact that there is a focus on the construction of methods and tools, in particular computational ones for all kinds of application areas. Therefore, many (applied) disciplines are involved in and supported by ontological engineering: information (library) science, (computational) linguistics, corporate knowledge management, data base technology (schemas), professional terminological standardization (e.g. medicine), and last but not least: knowledge engineering. In this paper we will focus on the last, in particular in the context of artificial *legal* reasoning and information serving. Each of these application areas bring different views and requirements to the ontology arena. Linguistic ontologies like WordNet (Miller, 1990) or SENSUS (Knight *et al.*, 1995) contain more than 50.000 concept definitions that can be used e.g. for translation. At the other extreme are the specialised domain ontologies (re)used in knowledge engineering as e.g. deposited at the Ontolingua server (Gruber *et al.*, 1996). In the latter, the definitions of the terms are much more detailed than in the former. However, in both cases there is the assumption that more abstract terms should constrain the interpretation of these definitions. Terms which re-occur in every domain and in our common sense understanding of the world, like time, space, cause etc. are part of a *top* ontology e.g. (Hobbs, 1995), (Sowa, 1995), (Guarino, 1995). The ontology of law that we present here is neither a top ontology, nor a domain ontology, but it is 'in between': it intermediates between the completely general top ontologies and ontologies of legal domains. After (vanHeijst *et al.*, 1997) who developed a similar kind of ontology for medical domains we call such an ontology a *core* ontology. Before presenting our core ontology of law in section we will first discuss the notion of ontology in general and the role ontologies can play in knowledge engineering.

A Notion of Ontology

It is impossible to represent the world in its full richness of detail. In order to represent a certain phenomenon or a part of the world (which is called a *domain*), it is necessary to restrict the attention to a small number of concepts which are meaningful and sufficient to interpret the world and provide a representation adequate to a certain task or goal at hand. As a consequence, a central part of knowledge representation consists of elaborating a *conceptualisation*: a set of abstract ob-

[1](Kant, 1787, B.874) who was interested in Leibniz work, learned about Leibniz via Wolff ("the Leibniz/Wolff system") classified ontology as one of the four components of metaphysics.

jects, concepts and other entities which are assumed to exist in a certain domain, as well as the relations that may hold between them (Genesereth & Nilsson, 1987). The commitments which are implied by the choice of one set of concepts instead of another to describe a certain phenomenon are called *ontological commitments*. A conceptualisation thus also carries a set of ontological commitments.

Elaborating conceptualisations (and thus selecting ontological commitments) is an essential component of the task of representing knowledge, because conceptualisations select which things are relevant to be represented and which are not (Davis *et al.*, 1993). Representing knowledge involves both the design of a *knowledge representation language* and the formulation of a specific set of sentences in this language which describe certain things in the world — such set of sentences is usually called a *knowledge base*. Ontological commitments precede the elaboration of both the knowledge representation language and the knowledge base: the distribution being that the representation language offers a minimal set where the consequences of the commitments emerge as organization and inference services.

A conceptualisation is an abstract entity which is only implied in a knowledge representation. An *ontology* is a specification of a conceptualisation (Gruber, 1994). It comprises a description (e.g. through *definitions*) of the concepts, objects, relations and so forth which make up a conceptualisation. One of the basic roles of an ontology is to enable the study of conceptualisations and ontological commitments in their own right, i.e. dissociated from the knowledge representations they may yield. Another important role is to support knowledge sharing and re-usability (Neches *et al.*, 1991; Patil *et al.*, 1992; Gruber, 1994). Particularly relevant are *domain ontologies*, i.e. ontologies which describe a part of the world or a human activity such as medicine, law or engineering.

The Role of Ontologies

Three kinds or levels of ontology can be distinguished in knowledge engineering: top, core and domain.[2] Each of these levels has its specific roles or use. Most of these roles may be subsumed under the term 'knowledge sharing' (Gruber, 1994), but the different purposes of (re)use of ontologies also put different requirements on ontologies. The qualitative demands increase with the order presented below, i.e. the need for careful and valid analysis and in particular for a *formal* base becomes greater.

Domain ontologies contain the concepts of some domain of application:

- Domain ontologies may be used not directly related to building knowledge systems, but as repositories for (organizing) knowledge and information. Ontologies may refer to distributed knowledge in organizations ('corporate knowledge'),or to common, standardized terminology in professional or scientific communities.
- A second role, related to the previous one is in knowledge acquisition where teams have to work together and an ontology becomes a common, agreed upon understanding of the terms in a domain, that can be read by team members with different background knowledge (Gruber, 1994).
- Ontologies make explicit to what conceptualization of terms a particular knowledge base(d system) is committed. In constructing a knowledge base one has to make commitments anyway, so making them explicit in an ontology enables more controlled development and also maintenance of a KBS. An (application) ontology is more than a simple documentation. It has a strong justification and quality assurance flavour, because an ontology supports consistent use of terms. This role is the standard one in the context of knowledge engineering methodology (see e.g. (Wielinga *et al.*, 1992), (vanHeijst *et al.*, 1997))
- The most often cited role is in enabling reuse of knowledge for building (new) applications. Although the previous role is probably the most frequently used nowadays in knowledge engineering, the major future benefit is to be expected from consulting repositories (libraries) of ontologies aggregated from earlier experiences or projects. The

[2](vanHeijst *et al.*, 1997, p. 193) proposes four levels: application, domain, generic, representation. We believe that there is no real distinction between application and domain ontology. The distinction is made because (vanHeijst *et al.*, 1997) believe that an application ontology fully matches a knowledge base and a domain ontology may only have a partial match with the knowledge base of some domain application. We do not hold such a view and argue that real application ontologies are impossible to construct (see also (Motta *et al.*, 1996)). The generic ontology is the same as a top ontology. However, the term 'core' is borrowed from (vanHeijst *et al.*, 1997) where they define the categories that index a library of medical domain ontologies: these categories are generic for medicine. We do not distinguish between generic and representation ontologies because they have the same nature. Whether a generic (= top) ontology is the basis for developing a KR formalism and related inference services or not is in the use of an ontology, in the same way as a domain ontology may be used to construct a knowledge base.

Ontolingua server is an example, but one may also think about a company specific library that is organized around the types of industry served, etc.

As (re)use of a domain ontology for constructing knowledge bases is by far the most important role, one should be aware that there is not a simple transfer or transformation that turns an domain ontology into a knowledge base. There are at least two reasons that complicate this process:

1. A representational mismatch. Ontologies are indexed by (the names of) concepts and some terminological (description) logic is used to represent these concepts and their relations. However most knowledge bases are cast in rule formats, where this indexing is absent or of secondary importance. The advantage of rule formats is that they easily express the reasoning, while the terminological version better clarifies what the epistemological – pardon, ontological – commitments are to the domain knowledge, as is well demonstrated by (Clancey, 1985)'s studies. In knowledge representation systems that combine terminological and assertional views on knowledge, such as LOOM, the transfer may be semi-automatic. However, the use of description logic based knowledge representation services is rather limited to academia.
2. A knowledge mismatch. This problem is more serious than the former. Not all knowledge that is used in solving problems in an application domain can be expressed in an ontology. For instance, heuristic, compiled-out knowledge is generally not part of an ontology. Either it can only be expressed in an ontology in a very complex, de-compiled way, or the ontology becomes completely method dependent, i.e. it will look rather like knowledge base and is not reusable for other kinds of tasks or methods (see (Motta *et al.*, 1996) for an illuminating case study).

   Also the reverse is true. Many of the definitions in an ontology are of no relevance in the reasoning of a knowledge system. A medical diagnosis program that uses heuristic classification to identify a kind of disease with a patient does not need to 'understand' that a disease is a process that interferes with normal biochemical processes and that changes over time, etc. (Clancey, 1985)).

   Another kind of knowledge mismatch that is somewhat less structural than the heuristic "contamination" in (efficient) knowledge bases, is the fact that the ontological requirements of problem solving methods — the "method ontology" of a PSM does in general not seamlessly fit a domain ontology (Coelho & Lapalme, 1996). More in general, domain ontologies themselves may overlap and exhibit lots of mismatches, as e.g. classified by (Visser *et al.*, 1997).[3] It should be noted that these kind of mismatches are generally not due to errors in one of the ontologies, but are rather inherent to the fact that the requirements for ontological modelling leave a large "design space" for multiple, correct solutions.

Core ontologies contain the *categories* that define *what a field is about.* A field is some discipline, industry or area of practice that unifies many application domains, as e.g. (some subfield of) medicine, law, engineering etc. The categories are not some common denominator of a set of application domains, but constrain what is *relevant* in these domains. As we will argue in section , core ontologies have a functional perspective. As intermediates between generic and domain ontologies they have a double-faced role:

1. Core ontologies are to be used for *indexing* libraries of domain ontologies. (Valente & Breuker, 1996) describe requirements for a core ontology to be able to play this role.
2. Core ontologies are the ontological basis for constructing special representation language and inference services for a field. A full understanding (and formalisation) of the inferential consequences (calculus) from the relationships between the terms that make up the abstract "top" of a category is required (Davis *et al.*, 1993), (Guarino, 1995).

Top ontologies The role of a top ontology is that it can make explicit what the ontological commitments are of some domain ontology. The most rigourous way to 'import' a top ontology into a domain ontology is when the domain ontology is *expressed by* top ontologies, i.e. when the top ontologies have been operationalized as a representation formalism annex inferential calculus. This role also occurs for core ontologies. A core ontology may entail its own, special formalisms, but may also *modify* or extend the more generic versions of a top ontology. For instance, legal causality is a modified view on the more generic notions of physical causality and intention (see the next section).

A Functional Ontology of Law

A core ontology should reflect what a field is about. In fields concerned with the design and use of physical artifacts, i.e. engineering, this means that the categories

[3] The term "knowledge mismatch" matches more or less their term "conceptual mismatch".

like design descriptions (component-model, process-model, and mathematical-model) constitute the core ontology (Borst & Akkermans, 1997). For instance, in commenting on the ontologies that make up the elevator domain (VT), which is a typical engineering domain, (Gruber *et al.*, 1996, p. 570) state: "Ontologies such as these can be *designed* (italics Gruber *et al.*), with respect to explicit design goals and objective criteria (Gruber, 1995)". The categories used are "component assemblies" and "parametric constraints", which match the component-model and mathematical-model of (Borst & Akkermans, 1997). Their 'process-model' is not present, because in elevator configuration the process descriptions are already completely parameterized. In other words, the VT-configuration does not exhibit the full engineering design categories, which can be summarised by the following dependencies string:

$$component - model \rightarrow process - model \rightarrow mathematical - model$$

This functional perspective on engineering also shows what the "natural" tasks are that are performed in engineering: modelling, design and configuration (Breuker, 1994). In engineering one is no longer directly concerned with the actual production and use (and repair!) of the artifacts that are designed. The manufacturing and maintenance of devices have become separate 'fields', having their own functional perspectives.

The View

To a large extent, law is also an artifact, so a functional perspective is strongly implied already by the objects (concepts) it uses. However, we cannot borrow the same categories as in engineering for several reasons. The law as a social system cannot be 'designed' over and over again: a social system emerges by local "designs", which means in law such activities as drafting laws, justifying legal decisions, legal policies etc. Therefore, the total "design" of the law is rather presupposed and the focus in legal reasoning and practice is on the way the law is applied and how it affects its environment: a society.

The functional ontology of law, as the name indicates, adopts a *functional* perspective — that is, the Law is analysed and interpreted through a functional point of view. This perspective can be summarised as follows:

The legal system as a system It is assumed that the legal system as a whole (and therefore each of its components) exists to accomplish a certain *function*, in order to obtain certain *social goals*. The legal system is thus viewed as an entity with a certain internal structure, behaving in an environment, and that was designed to work in a certain way in order to be able to accomplish specified social goals.

The main function of the legal system The legal system is one instrument to change or influence society in specific directions, determined by certain social goals. The main function of the legal system is to react to social behaviour. In this sense, the legal system can be regarded as a kind of *social device* operating within society and on society, and whose main function is to regulate social behaviour.[4] It is important to note, however, that this 'device-like' view does not imply that the ontology is based on a simulation of the actual behaviour of the legal system.

Functions of law = functions of legal sources The *legal sources* (the elements accepted as binding law in a legal system) are the knowledge which specifies how the legal system works. Therefore, the functions that legal knowledge may perform correspond to roles of legal sources. In other words, the categories of legal knowledge provide a direct classification principle for legal sources. Important to note that the proposed classification is in principle independent of (and in fact orthogonal to) the form or structure of these sources — they may be regulations, precedents, statutes, etc. Further, a legal source does not necessarily perform a single function: an elementary legal source (e.g. an article in a regulation) may express (contain) knowledge of more than one type.

From functions to roles Like a knowledge-based system, the legal system executes a number of *tasks*, for which it uses extensive *knowledge*. Consequently, each piece of knowledge used by the legal system has a specific *role* distinguished by the legal system in the operationalisation of its functions and tasks. In this way, the functions which are identified to distinguish types of legal knowledge have a dual character: they point out sub-functions (and sub-tasks) of the legal system, and at the same time divide legal knowledge into a number of basic categories which provide support for each of these functions/tasks. As described above, the main function of the legal system corresponds to a macro-argument from social behaviour to legal consequences. This means

[4] Enacting legislation is not the only way to change social behaviour. One can change the world physically so that the behaviour cannot occur (e.g. building cars that only drive at less than 100 km/h) or simply convince people through education or propaganda that the behaviour is 'bad'.

that these subtasks are connected to the generation of *micro-arguments* which together compose the desired macro-argument, and not necessarily to actual tasks in the legal system.

The Ontology

Given the view described above, an ontology of law can be built by identifying these functions and using them to distinguish categories of legal knowledge. In the following sections, a number of primitive functions of legal sources and corresponding categories are proposed and described. Some of these categories are primitive, while others are defined using the primitive categories. The primitive categories are: *normative knowledge*, *world knowledge*, *responsibility knowledge*, *reactive knowledge*, *creative knowledge* and *meta-legal knowledge*. It was shown in (Valente, 1995) that most of the important non-primitive categories, such as *rights*, can be defined based on the primitive ones. This issue, however, is outside the scope of this paper.

While this ontology has been inspired by the work of many legal theorists, (Hart, 1961) and (Kelsen, 1991) in particular, it diverges from their work in significant ways. In particular, the separation we make between responsibility, causal and normative knowledge is quite non–standard, and to some extent defines a novel legal–theoretical position. However, it is outside the scope of this paper to defend these ideas, or even the legal theoretical soundness of our ontology. The interested reader can find a detailed discussion in (Valente, 1995).

Normative Knowledge Normative knowledge is the most characteristic category of legal knowledge, to such an extent that to many authors 'normative' and 'legal' are practically the same thing. From an ontological perspective, however, it may be interesting to differentiate more types of legal knowledge and thus give to normative knowledge a more specific structure, content and role. Nevertheless, the restricted scope that is adopted for normative knowledge retains its importance as the central knowledge type in law.

Normative knowledge is seen in the legal theory literature as having two functions: prescribing behaviour and defining a standard of comparison for the social reality. These functions are clearly related. One can think that norms prescribe behaviour by defining a standard or, the other way around, set up a standard by prescribing behaviour. Which side is the central one is an issue subject to debate in legal theory. In this ontology we adopt the conceptualization that the definition of a standard is the central characteristic of normative knowledge, and the prescription of behaviour is one of its effects.

A norm expresses an idealisation: what ought to be the case (or to happen), according to the will of the agent that created the norm. This idealisation is expressed by reference to a description of the reality (the world) in which some configurations of facts and behaviour are 'cut out' to make it an ideal world. Since they express an ideal world, norms can be either *observed* or *violated*. A norm is observed when the behaviour in the real world does not conflict with its specification in the ideal world, and violated otherwise. To *apply* a norm means to verify or compare the reality with the ideal world defined in the norm, classifying the reality as either compliant or non-compliant with the norm. This classification is the *normative status* of the behaviour with respect to the norm.

Meta-legal Knowledge Normative systems are defined on the basis of individual norms in the sense that the standard defined by the normative system is defined in terms of the standards defined by the individual norms. The difference between the standard defined by the normative system on the one hand and the standards defined by the norms on the other hand is fundamental to understanding the role of normative knowledge in law, and it is accounted for by knowledge *about* individual norms. This distinction is captured by defining the categories of *primary norms* and *meta-legal knowledge*. Primary norms are entities that refer to human behaviour, and give it a *normative status*. This normative status is, in principle, either *allowed* (legal, desirable, permitted) or *disallowed* (illegal, undesirable, prohibited). However, each norm refers only to a few types of behaviour, in the sense that it is only able to provide a status if applied to some types of cases. For the remaining types of cases, the norm is said to be *silent*.

There may be a difference between the normative status given by a single norm and the one given by the normative system. The normative status with respect to the normative system is based on the normative status with respect to the primary norms. One of the functions of meta-legal knowledge is to specify how this process occurs, i.e. how the normative status with respect to the normative system is built from the normative status with respect to the primary norms. The basic mechanism involved is the solution of *conflicts* between primary norms. Another function of meta-legal knowledge is to specify which legal knowledge is *applicable*. Applicability is a concept which can be used both for specifying the dynamics of the legal system and its limits.

World Knowledge By its very nature, law deals with behaviour in the world. Therefore, it must con-

tain some description of this behaviour. For instance, in order to describe how the world should (ought to) be, primary norms must describe how things are or can be. This description is not directly available from the legislation, but is usually implicit. However, this type of knowledge is distinct from (albeit connected to) normative knowledge — that is, primary norms describe an ideal world *based* on the description of reality. This separation between the knowledge used to describe the world and the normative knowledge was first explicitly proposed by Breuker (Breuker, 1990), who called this category *world knowledge*.

In addition to adopting a category of legal knowledge which describes the world, we propose that this knowledge constitutes a structured *model*. Thus, the term *legal abstract model* or LAM is used as a synonym for world knowledge when its model character is to be stressed. Also, this term stresses that this model is a 'double abstraction': as discussed below, the LAM gives a more restricted meaning to common sense concepts which are already abstract.

The legal abstract model is an interface between the real world and the legal world. Its role is to define a model of the real world which is used as a basis to express normative and other categories of legal knowledge. The bulk of the LAM consists of definitions of *concepts* that represent entities and relations in the world. These concepts are usually organised in class hierarchies, i.e. some concepts are subsumed by others. In defining concepts and relations based on other concepts and relations, it occurs naturally that some concepts are taken as *primitive*, i.e. are not defined. These concepts are supposed to be interpreted by people. Indeed, since the law is supposed to be comprehensible to the people who are addressed by it, every concept used in general legislation is supposed to be comprehensible to *any* person.[5]That means these concepts are supposed to be part of this 'consensus reality' mentioned by (Lenat & Guha, 1990) — that is, these terms are in fact supposed to be interpreted using common sense knowledge. This means that the legal abstract model is not complete or self-contained: the LAM is in fact a layer of definitions of concepts and relations built on top of a large layer of common sense knowledge. The primitive (non-defined) concepts and relations in the LAM are the interface between the world knowledge and the common sense knowledge. For instance, it is possible to define in detail the characteristics of an `intellectual work` to be used in a copyright law. It is also possible to define it referring to other concepts, such as `books`, `sculpture`, etc. However, if the concept `book` is left undefined (primitive) the only way to interpret it is to rely on common sense knowledge. Therefore, a concept within the LAM is defined in two steps: first defining it in terms of other concepts in the LAM until the level of primitive concepts is reached, and second the common sense 'definition' of the primitive concepts. When one retraces the way back — verifying whether or not a certain legal concept (a concept in the LAM) applies to a certain entity in the world — there are again two interpretive steps. First there is the interpretation of the reality in the primitive terms of the legal abstract model, e.g. asserting that a specific thing in the world is a `book`. Second, there is the interpretation of the primitive terms with respect to the defined ones, e.g. asserting that a `book` is a type of `intellectual-work`. The first process involves purely common sense knowledge, while the second is supported by world knowledge. The process from an entity in the world to primitive concepts or vice-versa is called *common sense interpretation*. The process from a primitive (or defined) concept to a defined concept within the LAM or vice-versa is called *legal interpretation*. The view just described is shown in Fig. 1.

The legal abstract model is not only a static description of salient features of society, but a model of social behaviour. For instance, legislation is passed as a form of social control, under the supposition that it will affect social behaviour in a certain way, and this supposition is to some extent embedded in the legal abstract model. Therefore, a LAM should ideally allow predictions about behaviour in the world. Nevertheless, this requires more than definitions, because these are purely static. The definitions alone are a domain ontology which falls short of a model, because it does not enable such predictions. The view imposed in the law about behaviour is thus largely a static description of possible or relevant behaviours, but the full reasoning about them is left to common sense.

Apart from describing the world, what is behavioral reasoning used for in law? We propose that this description of possible and relevant behaviours is built around the concept of *cause*, in order to allow the assignment of responsibility of an agent for a certain case. Since these responsibility assignments actually form a separate category of legal knowledge, they are discussed in more detail below. [6]Causal knowledge, however, refers to or uses a static description of the

[5]Except for the obvious categories of people like infants or mentally retarded people.

[6]This is not a standard way of viewing causality in law. In most legal systems causal knowledge and responsibility knowledge is mixed, often even with normative knowledge, and so is the reasoning with these types of knowledge. We propose to separate it for clarity and efficiency reasons.

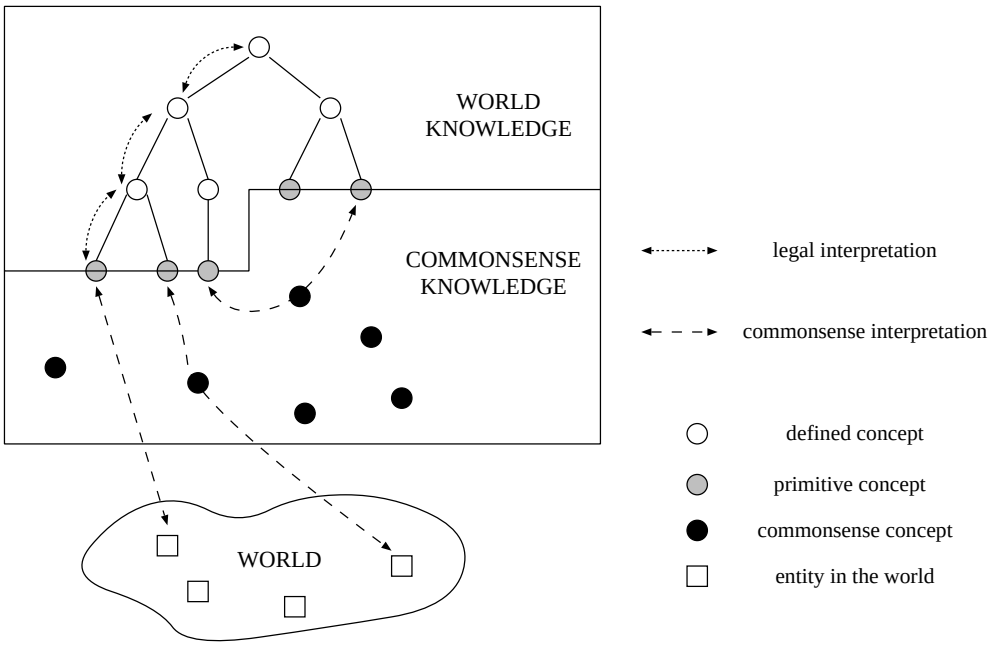

Figure 1: The role of the legal abstract model.

world (e.g. to model world states). Accordingly, we propose that the world model is actually composed of two related types of knowledge: *definitional knowledge* and *causal knowledge*. The definitional knowledge is used by the normative knowledge to describe the ideal world they define. The causal knowledge is used by the responsibility knowledge to describe who or what has caused a given state of affairs, and can thus be considered responsible for it.

Responsibility Knowledge It is commonplace to affirm that cause and responsibility are important concepts in law, but it is not equally simple to provide a more concrete account of how these concepts enter the realm of law. We see *responsibility knowledge* as a category of legal knowledge that has as a function to assign or limit the responsibility of an agent over a given (disallowed) state of affairs — i.e. to (dis)establish a link between the violation of a norm and an agent which is to be considered responsible (accountable, guilty, liable) for this violation. This responsibility link may be established by a causal connection between the agent and the disallowed behaviour, but this is not the only way to establish responsibility.

Responsibility is the intermediary concept between normative and reactive knowledge, since a reaction can only occur if the agent is is held responsible for a certain norm violation. Responsibility knowledge plays the role of linking causal connections with a responsibility connection — i.e. that connection which makes an agent accountable for a norm violation and possibly subject to legal reactions (see also reactive knowledge below). As (Hart & Honore, 1985) point out, however, responsibility does not have "any implication as to the type of factual connection between the person held responsible and the harm" — that is, causal connections are only a "non-tautologous ground or reason for saying that [an agent] is responsible" (Hart & Honore, 1985, pag. 66).

In principle, common sense says that one is only responsible for what one causes. Naturally, there is an *a priori* link between these two relations: in principle, all agents are legally responsible (and *only* responsible) for that which they cause. But this causal connection is not always necessary or always sufficient for establishing responsibility in a legal context. The role of responsibility knowledge is exactly to 'interfere' with this *prima facie* connection between causing and being responsible. This interference is made so that legal systems "extend responsibility [or] cut it off in ways which diverge from the simpler principles of moral blame" (Hart & Honore, 1985, pag. 67). This mechanism has rather practical motives. Given the innumerable problems in establishing, proving and reasoning with causal connections, the assignment of legal responsibilities which bypass these connections to some extent is used to give more precision in situations where the use of the common sense or moral concept of responsibility by the law can lead to inconsistencies or undesired

results, or when there is a practical interest (based on an implicit of explicit policy) that a frontier should be drawn so that it becomes easier to define what are the limits of responsibility under certain circumstances.

There are two basic mechanisms which are used in responsibility knowledge. First, the law may establish a responsibility connection independent of a causal connection — i.e. a *responsibility assignment*. This can be seen in a rule used in e.g. French, German or Brazilian law, by which parents are held responsible for the the damage done by their children even if there is no specific causal link between their attitudes or actions and the damage. That is, the parents are held responsible even though they have not necessarily caused the damage. Second, the law may limit the responsibility of an agent under certain circumstances, disregarding some possible causal connections — i.e. a *responsibility restriction*. For instance, in England a man is not guilty of murder if the victim dies more than one year after the attack, even if the death was a consequence of this attack. Other well-known factors that may influence the establishment of responsibility connections in law are *knowledge* and *intention*.

Reactive Knowledge To reach the conclusion that a certain situation is illegal (based on normative knowledge), and that there is some agent to blame for it (responsibility knowledge) would be probably useless if the legal system could not react towards this agent. That knowledge that specifies which reaction should be taken and how is what we call *reactive knowledge*. Usually this reaction is a sanction, but in some situations it may be a reward. The penal codes, which are usually a fundamental part of legal systems of the Romano-Germanic tradition, contain basically responsibility and reactive knowledge only.

Creative Knowledge A legislator may indirectly create some entity that did not exist before in the world, using what we call *creative knowledge*. It is usually stated in imperative terms, designating an agency that previously did not exist as part (or not) of the reality from a certain point of the time on. The creative function has a somewhat exceptional (or even abnormal) status when compared to the other ones. In this case, the law not only wants to classify or react to certain agents that already exist in the real world, but attempts to create a new agent. A simple example is the creation of a department within the government or a company. Note that the department created is itself an abstract entity which will later be 'filled' with workers, facilities, etc. to become a concrete, physical entity. Creative knowledge is not necessarily restricted to the State: the creation of organisations by contracts and agreements can also be considered creative knowledge. This type of knowledge appears at a first glance to be almost negligible, but it has in practice a very important role. The legal system must regulate itself as any other social organisation. There is an important difference, though: the law can design the structure of the legal system as an organisation, in much the same way as companies design their structure by their internal regulations. The creative knowledge is what is used to perform this function.

The Big Picture

So far, a set of categories that divide legal knowledge like pieces of a puzzle was presented. These categories are distinguished by their *function* or *role* in the legal system, and together realize the main function of the legal system as discussed in section . However, it still remains to be shown how these pieces together make the whole puzzle, i.e. why they are sufficient to represent legal knowledge. In order to do that, this section tries to put these functions together and show that they perform the function the whole legal system performs. That is the *coherence* of the ontology, and to some extent its completeness.

Fig. 2 shows how the functions/categories identified compose together the main function of the legal system. Of course, it is not meant that there is a one to one correspondence between the abstract processes of Fig. 2 and the actual social processes and procedures in a legal system. Instead, they are functional dependencies which describe how the main function is decomposed in sub–functions which together perform it. In Fig. 2, the rounded boxes represent functions (or, alternatively, bodies of knowledge which perform the function), and the solid arrows indicate functional dependencies, i.e. the inputs and outputs of these functions. The dependencies which correspond to actual interactions with the society are indicated in the figure in non-solid arrows. The entities in the society are specific social agents, e.g. the University of Amsterdam (a school), Joost Breuker (a private person), the Ministry of Education (a government agency), Philips BV (a company), etc.

A cycle starts with a *real world situation*, which is interpreted in order to generate an abstract description of the case in the terms that the legal sources use. This abstract case description is called a *legal situation*, and the knowledge used to produce this step is the *world knowledge*, which forms the *legal abstract model*. Then, the legal situation is analysed against the *normative knowledge* to verify whether it violates any norm, thus producing what is called a *classified situation* (a situation classified as either 'allowed' or 'dis-

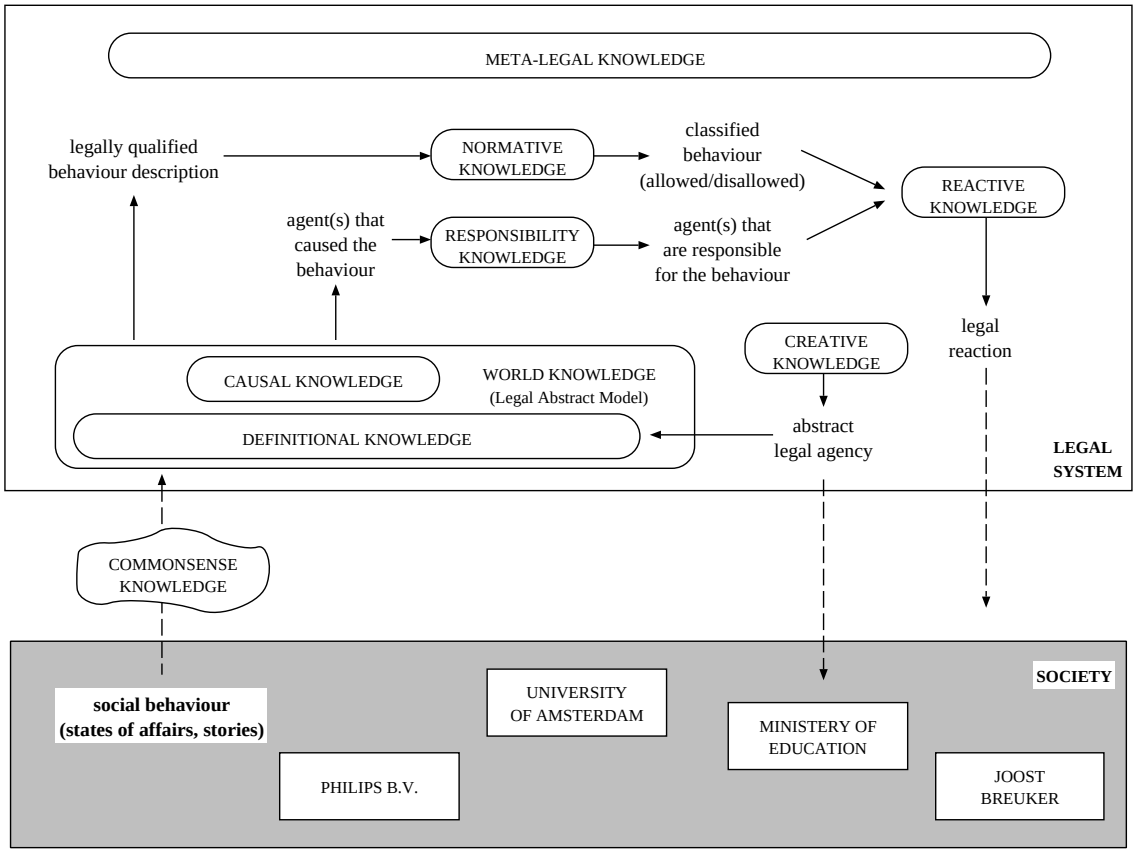

Figure 2: Functional roles of legal knowledge in the operation of the legal system.

allowed'). In another path, the situation is analysed using again *world knowledge* (but here particularly its *causal* component) in order to find out which agents in the world (if any) have caused the situation. This information is then used as input to the *responsibility knowledge* which determines which agents (if any) are to be held responsible for the situation. The results obtained in these two paths (the classified situation and the responsible agents) are then used as inputs for a function that defines a possible *legal reaction* using *reactive knowledge*. Further, outside this cycle, the law may also create an abstract entity (part of the legal system) using *creative knowledge*; this entity is also added to the *legal abstract model*. Finally, *meta-legal knowledge* refers to all these entities.

Another way to see the interdependencies shown in Fig. 2 is that they provide the connections between the (sub)functions from a reasoning point of view. That is, legal reasoning can be made modular, with each function corresponding to a module and the dependencies between the modules being provided by the dependencies between the functions. Such dependencies must of course be detailed. Moreover, the main path in the picture can be also seen as the global structure of legal arguments: starting from the 'facts of the case' and going up to sentencing, with each function providing the source for argument steps in specific places within the global argument.

The scheme shown in Fig. 2 can also be seen as the basic structure of legal arguments. Each category corresponds to a type of argument that has as antecedents the inputs and as conclusions the outputs of each function, and as warrants the knowledge belonging to that category. For normative knowledge, for instance, the conclusion is whether a situation is allowed or disallowed, and the warrants are normative knowledge. Moreover, the conclusions in a legal argument are concatenated as shown by the dependencies in the figure; for instance, an argument involving world knowledge (say, concluding that a certain person is considered a 'minor' according to a certain definition) being used as subsidiary for an argument involving normative knowledge (say, concluding that a situation in which this person was driving a car is disallowed according to a certain norm). Legal reasoning can be thus seen as the production and analysis of arguments involving one or more of these categories.

An Example

We will illustrate the different types of knowledge, and how the are related to legal sources, by an example from Dutch Labour law. Figure 3 presents an excerpt from Dutch civil law concerning labour arrangements.

Art. 1637a
The labour contract is the agreement in which one party, the employee, commits itself to perform labour in the service of the other party, the employer, in return for payment, for a certain period of time.

Art. 1637b
The undertaking of work of a material nature is the agreement in which one party, the undertaker, commits itself to bring about a particular work for the other party, the out-contractor, for a particular price.

Art. 1637c 1.
When an agreement has the characteristics of a labour contract, and of some other type of agreement, both the provisions of the labour contract and of the other agreement will apply; in the case of a conflict between the provisions, those of the labour contract will apply.

Art. 1637z
The provisions of this title are not applicable to persons in the service of the state, province, council [...].

Art. 1639a
- 1. The employee is obliged to perform the labour personally [...]

Art. 1639da
The employee who causes damage to the employer [...] during the realization of the agreement, is not to be held responsible by the employer, unless the damage was caused on purpose, or is a consequence of deliberate rashness.

Art. 1639o 1.
All parties involved can end the agreement without notice [...], but the party that does so without consent of the other party, has to give compensation, unless he does so for a pressing reason of which the other party is informed immediately.

Art. 1639r 1.
The compensation, mentioned in article 1639o, [...] is equal to the amount of wages for the time the agreement had lasted if it were terminated in a legitimate way.

Figure 3: Articles from Dutch civil law (BW 7A) concerning labour arrangements.

To introduce this domain, we will start with discussing the world knowledge in these articles, before we move on to the other categories.

World Knowledge The core issue in the domain of labour law is some sort of 'agreement' between two parties concerning some 'labour' to be performed by one of the parties, in return for a reward (usually money) payed by the other party. In this agreement, written or verbal, things can be settled like working hours, pay, holidays, etc. In some branches of industry collective labour contracts exist, negotiated between the unions and the employers. For individual workers there can be many individual arrangements, but the law only defines one sort of contract: a 'labour contract' (art. 1637a, Fig. 3), and

another type of arrangement called 'undertaking of work' (art. 1637b). These two articles clearly contain world knowledge, and more precisely the *definitional* part of it; they describe concepts the legislator wishes to see in the world of labour arrangements. Apart from the types of arrangements, they mention parties that have the agreement, the work that is subject of the agreement, time, and a price or payment. Terms like 'labour' and 'work' are primitives; their interpretation is left to common sense. The articles also refer to actions that can occur in the world of labour arrangements: Agreements can be ended (art. 1639o), employees can cause damage (art. 1639da), etc. These are also part of the world knowledge, but they represent dynamics in the world and belong to the *causal* knowledge.

Normative knowledge
The working agreements themselves contain norms: they describe all kinds of rights and duties for the parties involved. Moreover, the law contains norms about these types of arrangements. Article 1639a (Fig. refbw), for instance, states that the employee is *obliged* to perform the labour personally. This is a clear example of a primary norm. As we have seen above, an 'employee' is defined in the world knowledge as one of the parties in a labour contract. So they 'ideal' situation represented in this article is one in which the employee in the case of a labour contract performs his labour personally. Should he let some other person perform his labour, then he violates the norm.

Meta-Legal Knowledge The part of Dutch labour law presented above, contains two different examples of meta-legal knowledge. First of all, art. 1637c contains a domain specific version of the 'lex superior' principle. In the case the provisions of a labour contract *conflict* with those of some other type of agreement, the provisions of the labour contract 'win'.[7] The second example is art. 1637z, which limits the applicability of the legislation.

Responsibility Knowledge Art. 1639da contains responsibility knowledge, though the norm that is violated is left implicit here (i.e. there is no norm in this section stating one should not cause damage to someone else). The situation in which the employee *causes* damage to the employer during the working relation, is 'not ideal', does not comply with some norm, but provided it was not done *intentionally*, the employee will not be held responsible for the damage.

Reactive Knowledge Given a violation of a norm, and an agent that is held responsible, a legal reaction can be taken. Article 1639o and 1639r specify that in the situation in which one party ends a working agreement unilaterally without notice (thereby violating a norm specified in art. 1639e ff. not cited here), that party has to give compensation equal to a certain amount of wages. In other words, these articles contain knowledge about a (possible) legal *reaction* to a norm violation.

Creative Knowledge Creative knowledge, finally, is not easy to find in Dutch Labour law. A related law, however, the "Law on labour councils"[8] has creative knowledge as its core. It is a law, which main purpose is to establish the legal entity of 'works council', that did not exist before, and specify how it should be organized, etc.

Discussion and conclusions

A 'principled' core ontology

The typical roles for core ontologies as put forward in section are the indexing of libraries of ontologies and the development of representation formalisms. The first role of our ontology is discussed in (Valente & Breuker, 1996), where it is argued that this role can only be performed if the core ontology is well constructed, or "principled" as is summarized as follows:

1. It is *parsimonious* in that, despite trying to cover the whole of legal knowledge it departs from only six basic categories, which cover at most a few concepts each.

2. It has a *theoretical basis*, in our case derived largely from the literature in legal theory. The ontology can also be seen as a contribution to legal theory, and this role has been explored extensively in the longer description the reader can find in (Valente, 1995).

3. It is *complete* in that it provides a self contained set of basic (primitive) categories of legal knowledge such that all types of legal knowledge can be defined based on them. For example, we have no specific definition equivalent to the principle that one can only be judged once for an accusation, but this principle can be found to play the role of meta-legal knowledge. It must be noted that we do not mean that all legal concepts can be defined solely using the concepts defined here, only that their *role* can. For

[7] This looks like a standard meta rule, only in this case the 'object' level are norms in contracts, not in the legislation itself.

[8] "Wet op de ondernemingsraden".

instance, there is a large amount of legal sources in the US legal system pertaining to civil rights. While we do not have a concept `civil-right` in the ontology, we can show (see (Valente, 1995, Chapter 3)) that rights in general can be expressed in terms of composing normative and responsibility knowledge.[9]

4. It was shown that the basic categories in the functional ontology of law are *coherent* by the fact that they are related to each other in a consistent way to perform the main function of the legal system as discussed in Section .

The second role is in developing representation formalisms. For normative knowledge in the legal ontology we have developed formalisms for representing and interpreting the deontic modalities (forbidden, permitted, obliged). It appeared that the ontological basis for handling these modalities by some (extended) form of (classical) deontic logic was rather shaky. Therefore, new logical formalisms were developed and operationalized in the ON-LINE shell (see next section). The formalization is reported in (Valente, 1995, chapter 5).

Some additional advantages

However, some other advantages of taking a principled core ontology for developing artificial legal reasoning systems can be pointed out, based upon the fact that the ontology provides the handle for divide and conquer strategies. Since an ontology defines how one sees the world in terms of primitive knowledge categories and their interrelations, it divides the world into pieces which can be solved separately — provided, of course, that care is taken of their dependencies. For example, the ontological commitment to the six categories proposed in the functional ontology of law naturally divides the study and representation of legal knowledge in six parts. This leads to several interesting consequences.

First, a degree of flexibility is added by the fact that the solutions given to the parts may not be at the same level of detail. One can be formal and principled studying one category and pragmatic and symbol-level studying another. For example, while we were able to propose a formal theory for reasoning with normative and meta-legal knowledge (in fact detailing this part of the ontology), the many challenges in proposing a general solution to the representation of common sense and world knowledge refrained us from doing the same to these categories. Instead, we proposed only that a certain type of representation formalisms (terminological or concept languages) were particularly adequate. This is specially important in the design of field-specific representation formalisms. In general, ontologies are an excellent basis for designing specialized languages or representation formalisms for a certain domain. In principle, different languages may be built based on a certain ontology. These languages may differ for instance in syntax or notation, varying from semi-formal languages to mathematical logic, or in using other (additional) ontological commitments, for instance to using defaults as allowed (and stimulated) by frame-based representations. In the functional ontology of law, we proposed languages varying from formal to semi-formal for representing most of the categories proposed (see (Valente, 1995)).

Because a principled core ontology partitions the domain in reasonably self-contained pieces, one does not need to find a single general formalism to reason with them. For instance, it is commonly proposed in AI & law that the use of rule-based formalisms or some form of deontic logics can be enough to represent legal knowledge. This departs from the assumption that we are looking for a single, unifying reasoning engine. In the operationalization of the functional ontology of law we have proposed different reasoning engines for each category, depending on the specific representation used. Although this was not proven in formal terms, this strategy has the clear potential to make general reasoning with legal knowledge more tractable, since each basic category has specific requirements and restrictions which can be explored in building efficient reasoning mechanisms. In other words, the use of a principled core ontology such as the functional ontology of law can simplify the study of domain specific reasoning, and can potentially improve the tractability of the reasoning engines proposed.

Finally, principled core ontologies naturally constrain the types of arguments and explanations used in the domain. For example, the scheme shown in Fig. 2 can also be seen as the basic structure of legal arguments. Each category corresponds to a type of argument that has as antecedents the inputs and as conclusions the outputs of each function, and as warrants the knowledge belonging to that category. For normative knowledge, for instance, the conclusion is whether a situation is allowed or disallowed, and the warrants are normative knowledge. Moreover, the conclusions in a legal argument are concatenated as shown by the dependencies in the figure; for instance, an argument involving world knowledge (say, concluding that a certain person is considered a 'minor' according to a certain definition) being used as subsidiary for an ar-

[9] The "civil" part defines a type of right based on the type of agent who has it — a definition that can be expressed using world knowledge.

gument involving normative knowledge (say, concluding that a situation in which this person was driving a car is disallowed according to a certain norm). To use Toulmin's terminology (Toulmin, 1969), an ontology defines what types of conclusions, warrants and argument chains are commonly used and/or held valid. This may be an important factor if domain-specific reasoning is viewed as the production and evaluation of arguments (as it is indeed the case in law) .

These advantages were explored in designing the ON-LINE architecture for legal problem solving (Valente & Breuker, 1995). ON-LINE represents the functional ontology in the description logic/system LOOM (MacGregor, 1991), and uses this representation to define specialized representations for objects in each of the categories of the ontology. However, the main advantage of using a principled core ontology as proposed in this paper appears when defining reasoning mechanisms. Different mechanisms are used to reason with each category. For world knowledge, for example, we borrowed LOOM's classification/subsumption reasoning engine.[10] For normative reasoning, on the other hand, we devised a specific algorithm for applying norms to cases (Valente, 1995) (in contrast with e.g. using a deontic reasoning engine). In other words, instead of trying to define an all-encompassing formalism and reasoning mechanism that fits all types of legal knowledge, the functional ontology of law provided a basis for defining specialized representations and algorithms that are easier to handle and potentially more tractable.

Comparing core ontologies

Ours is not the only core ontology. We have already referred to the PhySys ontology developed to cover engineering in electro-mechanical domains (Borst & Akkermans, 1997). For medical domains van Heijst developed a core ontology that is aimed at the indexing of the domain ontologies in the GAMES library. Besides a functional center that views the core business of medicine as the diagnosis and remedying of diseases, other categories refer to the various disciplines that make up the theoretical aspects of medicine. In (Valente & Breuker, 1996) we have argued that this hybrid view on a core ontology makes it less suitable for indexing.

Although originally not intended this way, the ontology of *law* as proposed by (Visser, 1995) and (vanKralingen, 1995) can also be viewed as a core ontology. This (K-V) ontology is of the same level of abstraction as ours (B-V) and is intended to cover the same field. Indeed, there is a strong conceptual overlap between both ontologies, i.e. the semantic similarity is remarkable, which would suggest that in the field of law there is a large consensus of what constitute law. Such a conclusion would come as a shock to many a legal theorist. There is more dispute over terms there than in legal practice that is characterized by dispute rather than problem solving. Indeed, it is not difficult to point to subtle differences between similar terms in the K-V and the B-V ontologies, but it is more important to point to structural differences. The first one is that the K-V ontology is not a functional one. It lists a number of categories, or rather slots in a frame with value restrictions, but it does not specify how these categories are dependent on one another. Second, K-V does not distinguish between common sense knowledge and the way the law looks at events in the world. For instance, there is no distinction between causality in common sense terms and the way the law restricts or extends causality in order to define legal causality (responsibility). Similarly, the law often makes its specific definitions that may conflict with common sense descriptions. This is of importance, because, for instance, the case descriptions that have to be legally assessed, are framed in common sense terms, while the interpretation of such a case has to make use of the legal view on the objects, events and states.

[10] This decision is not arbitrary; on (Valente, 1995) we show why description logics are particularly adequate to represent and reason with world knowledge.

References

P. Borst and H. Akkermans. Engineering ontologies. *International Journal of Human Computer Interaction*, 46:365 – 408, 1997. special issue on using explicit ontologies in kbs development.

J. Breuker. Towards a workbench for the legal practitioner. In C. van Noortwijk, A. Schmidt, and R. Winkels, editors, *Legal Knowledge Based Systems: Aims for Research and Development*, pages 25–36. Koninklijke Vermande, 1990.

J.A. Breuker. A suite of problem types. In J.A. Breuker and W. Van de Velde, editors, *CommonKADS Library for Expertise Modelling*, pages 57 – 88. IOS-Press/Ohmsha, Amsterdam/Tokyo, 1994.

W.J. Clancey. Heuristic classification. *Artificial Intelligence*, 27:289–350, 1985.

E. Coelho and G. Lapalme. Describing reusable problem solving methods with a method ontology. In B.R. Gaines and M. Mussen, editors, *Proceedings of the KAW-96*, Banff, Ca, 1996.

R. Davis, H. Shrobe, and P. Szolovits. What is a knowledge representation? *AI Magazine*, pages 17–33, spring 1993.

M. Genesereth and N. Nilsson. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, Los Altos, CA, 1987.

T. Gruber. Toward principles for the design of ontologies used for knowledge sharing. In N. Guarino and R. Poli, editors, *Formal Ontology in Conceptual Analysis and Knowledge Representation*. Kluwer, 1994.

T. Gruber, G. Olsen, and J. Runkel. The configuration design ontologies and the VT elevator domain theory. *International Journal of Human Computer Interaction*, 44:569 – 598, 1996. special issue on Sisyphus-VT.

N. Guarino. Formal ontology, conceptual analysis and knowledge representation. *International Journal of Human Computer Interaction*, 43:625 – 6640, 1995. Special issue on formal ontology.

H. Hart. *The Concept of Law*. Clarendon Press, Oxford, 1961.

H. Hart and T. Honore. *Causation in the Law*. Oxford University Press, New York, second edition edition, 1985.

G. Hirst. Ontological assumptions on knowledge representation. In R.J. Brachman, H.J. Levesque, and R. Reiter, editors, *Principles of Knowledge Representation and Reasoning*, pages 157–169. Morgan Kaufmann, San Mateo, CA, 1991.

J.R. Hobbs. Sketch of an ontology underlying the way we talk about the world. *International Journal of Human Computer Interaction*, 43:819 – 830, 1995. Special issue on formal ontology.

I. Kant. *Kritik der Reine Vernunft*. Everyman, second edition, 1787. translation by J. Meiklejohn, revision by V. Politis for the Everyman's Library edition, 1993.

Hans Kelsen. *General Theory of Norms*. Clarendon Press, Oxford, 1991.

K. Knight, I. Chander, M. Haines, V. Hatzivassiglou, E. H. Hovy, M. Iida, S.K. Luk, R.A. Whitney, and K. Yamahada. Filling knowledge gaps. In *Proceedings of IJCAI-95*, 1995.

D. Lenat and R. Guha. *Building Large Knowledge-based Systems. Representation and Inference in the CYC Project*. Addison-Wesley, Reading, MA, 1990.

R. MacGregor. Inside the LOOM classifier. *SIGART Bulletin*, 2(3):70–76, 1991.

G. Miller. Wordnet: an on-liexical database. *International Journal of Lexicography*, 3:special issue, 1990.

E. Motta, A. Stutt, Z. Zdrahal, K. O'Hara, and N. Shadbolt. Solving vt in vital. *International Journal of Human Computer Interaction*, 44:333 – 372, 1996.

R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W. R. Swartout. Enabling technology for knowledge sharing. *AI Magazine*, pages 36–56, Fall 1991.

R. Patil, R. Fikes, P. Patel-Schneider, D. McKay, T. Finin, T. Gruber, and R. Neches. The DARPA knowledge sharing effort: Progress report. In C. Rich, B. Nebel, and W. Swartout, editors, *Proc. of KR'92*. Morgan Kaufmann, 1992.

J.F. Sowa. Top-level ontological categories. *International Journal of Human Computer Interaction*, 43:669 – 686, 1995. Special issue on formal ontology.

S. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, 1969.

A. Valente and J.A. Breuker. ON-LINE: An architecture for modelling legal information. In T. Bench-Capon, editor, *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 95 — 104, New York, 1995. ACM Press.

A. Valente and J. Breuker. Towards principled core ontologies. In B.R. Gaines and M. Mussen, editors, *Proceedings of the KAW-96*, Banff, Ca, 1996.

A. Valente. *A Modelling Approach to Legal Knowledge Engineering*. IOS Press, Amsterdam, 1995.

G. van Heijst, A.Th. Schreiber, and B.J. Wielinga. Using explicit ontologies in kbs development. *International Journal of Human-Computer Studies*, pages 183 – 291, 1997.

R. W. van Kralingen. *Frame-based Conceptual Models of Statute Law*. PhD thesis, University of Leiden, The Hague, The Netherlands, 1995.

P.R.S. Visser. *Knowledge Specification for Multiple Legal Tasks; A Case Study of the Interaction Problem in the Legal Domain*. PhD thesis, University of Leiden, The Hague, The Netherlands, June 1995.

P. Visser, D.M. Jones, T. Bench-Capon, and M. Shave. An analysis of ontology mismatches;heterogenity vs interoperability. In *Proceedings of AAAI-Symposium on Ontological Engineering*, pages 164 – 172, Stanford, CA, 1997. AAAI.

B. J. Wielinga, A. Th. Schreiber, and J. A. Breuker. KADS: A modelling approach to knowledge engineering. *Knowledge Acquisition*, 4(1):5–53, 1992.