Artificial Intelligence and Law **6:** 27–57, 1998.

A Comparison of Four Ontologies for the Design of Legal Knowledge Systems ⋆

PEPIJN R. S. VISSER and TREVOR J. M. BENCH-CAPON
LIAL – Legal Informatics at Liverpool, Department of Computer Science, University of Liverpool, P.O. Box 147, Liverpool, L69 7ZF, United Kingdom, (+44) 151 - 794 3792, E-mail: {pepijn, tbc}@csc.liv.ac.uk

Abstract. There is a growing interest in how people conceptualise the legal domain for the purpose of legal knowledge systems. In this paper we discuss four such conceptualisations (referred to as ontologies): McCarty's language for legal discourse, Stamper's norma formalism, Valente's functional ontology of law, and the ontology of Van Kralingen and Visser. We present criteria for a comparison of the ontologies and discuss the strengths and weaknesses of the ontologies in relation to these criteria. Moreover, we critically review the criteria.

1. Introduction

In the Oxford dictionary the word ontology is defined as the branch of metaphysics dealing with the nature of being (Allen, 1990). The word ontology is used in AI research as well, although its meaning in the latter field is only remotely related to the original metaphysical meaning of the word†. In the latter field, ontologies are loosely defined as 'explicit conceptualisations of a domain' (Gruber, 1992). Recently the results of gathering explicit conceptualisations of knowledge-system domains has been recognised as a valuable effort in its own right, deserving of widespread attention (e.g., Wiederhold, 1994, p. 7). This trend can also be seen in the legal domain, for instance, Moles and Dayal argue that researchers in the field of AI and Law should study the (implicit) 'assumptions being made *about the nature of law*' (Moles & Dayal, 1992, p. 188). The first International Workshop on Legal Ontologies, LEGONT '97, was held in July 1997 (Visser & Winkels, 1997).

In this article we compare four conceptualisations of the legal domain. These are McCarty's LLD (McCarty, 1989), Stamper's norma formalism (Stamper, 1991, 1996), Valente's law functions (Valente, 1995; Breuker et al., 1997), and Van Kralingen and Visser's frame-based ontology (Van Kralingen, 1995; Visser, 1995; Visser & Bench-Capon, 1996a, b; Van Kralingen, 1997). All four conceptualisations,

⋆ This article extends on the work of the authors presented at JURIX '96, ICAIL '97 (Bench-Capon and Visser, 1996, 1997; Visser and Bench-Capon, 1996b) and LEGONT '97 (Visser and Bench-Capon, 1997).

† For a discussion on the differences, see (e.g., Mommers et al., 1997)

which we will refer to as ontologies, are intended to be used in the creation of legal knowledge systems. Our aim is to compare the four ontologies, to assess their merits and thus, to contribute to future work on legal ontologies. We start by discussing the notion of ontologies (Section 2). Then, we briefly discuss each of the four legal ontologies (Section 3). Thereafter, we introduce a set of ontology-comparison criteria (Section 4) and compare the ontologies using these criteria (section 5). Finally, we discuss the results, draw conclusions and provide suggestions for future research (Section 6).

2. Ontologies

Building knowledge systems involves the creation of a model of a particular domain (e.g., electronic circuits, legislation, plants). Such a model is necessarily an abstraction of the domain being modelled. Models are useful precisely because they abstract from irrelevant details and thereby allow us to focus on the aspects of the domain we are interested in. Building a model of a domain involves deciding what entities in the domain are to be distinguished, and what relations exist between these entities. Moreover, it involves deciding what *types* of entities, and what *types* of relations exist. Often, the latter kind of decisions are straightforward and not always explicitly documented. For instance, in building a model of the blocks world we use predicates such as *block(A)* and *on(A, B)*, thereby implicitly assuming the domain to consist of blocks, tables, and hands, and that blocks, tables and hands have spatial relations. Making such design assumptions requires the domain to be carved up into concepts. Alternatively stated, these assumptions constitute a *conceptualisation* of the domain under consideration. The conceptualisation tells us the types of entities and relations that are considered to exist. In the blocks world it tells us that there are blocks, tables, hands and that blocks and tables have spatial relations. It does however, not tell us what particular blocks, tables and hands there are, nor how they are spatially related.

It is important to note that the creation of a conceptualisation is not an unequivocally defined process. The same blocks world can be conceptualised in different ways. Some entities in the world may not need to be conceptualised at all (for instance, it is not always necessary to conceptualise the hand in the blocks world), other entities could have been specified more abstractly (e.g., recognising only *objects* and *spatial relations*) or less abstractly (e.g., recognising *cubes* and *cylinders*). In conclusion, we remark that making a conceptualisation is a process that is accompanied by a considerable amount of freedom.

2.1. DEFINING ONTOLOGIES

To be able to compare and analyse conceptualisations of a domain, it is useful to make them explicit. This is what the word ontology is used for in AI research: an *ontology* is defined as an explicit specification of a conceptualisation (Gruber,

1992). It establishes a shared understanding of some domain of interest (Uschold Gruninger, 1996) and usually this understanding is obtained by defining a set of terms with their meaning for describing the domain (viz. a vocabulary). An ontology is a knowledge-level description (Newell, 1982) in that it is independent of any representational formalism (Van Heijst, 1995). Also, an ontology is considered to be a meta-level description of the model under construction (viz. the knowledge base) because it abstracts from the particular entities and relations in the model; it only specifies the type of entities, their relations, and constraints on them (Van Heijst, 1995). Typically, an ontology consists of an hierarchically ordered collection of classes, instances, relations, functions and axioms, but conceptual graphs, semantic nets, and database schemes are sometimes referred to as ontologies as well. We could say that an ontology describes the domain knowledge (data) that remains invariant over various knowledge (data) bases in a certain domain (cf. Guarino and Giaretta, 1995). For instance, an ontology could specify that in all knowledge bases of the blocks world an empty block is a block with no block on top of it. It should be stressed here that different interpretations of the word ontology are used by different authors. This is nicely illustrated by Guarino and Giaretta (1995), who discuss seven interpretations of the word 'ontology'.

2.2. FORMALISMS TO SPECIFY ONTOLOGIES

To be able to specify a conceptualisation explicitly we need to have some language. In principle, this can be any language, varying from highly informal to very formal specification languages (Uschold & Gruninger, 1996). Often a frame-oriented representation language is used. Some of the more commonly used languages to specify ontologies are: PROLOG, LOOM (MacGregor, 1990),CML (Schreiber et al., 1994) and KIF/ONTOLINGUA (Genesereth & Fikes, 1992; Gruber, 1992). Most of these languages allow us to define classes, attributes, instances, functions, relations and various constraints (e.g., on the values of the attributes).

2.3. TOWARDS A CLASSIFICATION OF ONTOLOGIES

An ontology describes a set of assumptions about a domain. Explicitly documenting these assumptions is useful, for instance, if different agents have to communicate about the same domain. The ontology then serves as an agreed communication language for the agents. Agreements about the objects and relations being talked about among agents are referred to *ontological commitments* (Gruber, 1993b, p. 201). Here, we use the concept of ontological commitments to classify ontologies*. If we can distinguish different types of ontological commitments we will be able us to distinguish between different types of ontologies, and hence, we will

* For convenience we will also use the term *ontological commitments* to refer to the assumptions underlying an ontology, independent of the existence of an agreement between communicating agents. This allows us to state that an ontology *has* certain ontological commitments, and, that agents

be able us to classify ontologies. As a starting point to define different types of commitments we adopt the commonly made distinction in AI literature between tasks, methods, and domains (e.g., Breuker & Van de Velde, 1994). Hence, we arrive at three different types of ontological commitments: (a) task commitments, (b) method commitments, and (c) domain commitments.

(a) *Task commitments*. An ontology has task commitments if it defines entities and relations that express a task-specific perspective on the domain knowledge (Chandrasekaran & Josephson, 1997) (by a task we mean a specification of a goal together with some input and required output, see also Visser (1995)). Typical task commitments are found in, for instance, an ontology for a diagnosis task, which contains entities such as observations, causes, and hypotheses.
(b) *Method commitments.* An ontology has method commitments if it defines entities and relations that express a method-specific perspective on the domain knowledge (Chandrasekaran & Josephson, 1997) (by a method we mean a specification of how a task can be performed, see also Visser (1995)). Typical method commitments are found in, for instance, an ontology for the propose-and-revise method (within a design task), which contains entities such as proposed solution, constraints, and value-assessment.
(c) *Domain commitments*. An ontology has domain commitments if it defines entities and relations that relate to a particular domain (by domain we here refer to the commonly distinguished fragments of the real world modelled, such as medical, legal, mathematical, financial, or social domains). Typical domain commitments are found in, for instance, an ontology for the legal domain, which contains entities such as norms and acts.

Ontologies can be classified according to the type of commitments they make. Using the three types of commitments mentioned above, we could say that a *task ontology* is an ontology that makes (substantial) commitments towards a certain (group of) task(s), a *method ontology* is an ontology that makes (substantial) commitments towards a particular (group of) method(s), and a *domain ontology* is an ontology that makes (substantial) commitments towards a particular (group of) domain(s) (Chandrasekaran & Josephson, 1997). We note that the number of commitments (of a certain type) made in an ontology may vary, there can be a small amount or a large number of commitments. Hence, we can define a *generic domain ontology*, an ontology that makes commitments towards a particular domain but is generic in that it can be refined for many subdomains. As an example, we mention a generic legal ontology that distinguishes between norms, acts and concept definitions but does not make commitments towards any legal subdomain (e.g., Visser,

have to commit to the ontological commitments of the ontology if they want to use the ontology as a communication standard.

1995). Most ontologies will have a combination of the commitments mentioned. For instance, an ontology that distinguishes 'hypothesis-is-angina-pectoris', 'oppressing pain' and 'EG scan results' commits to both a (diagnosis) task and a (heart diseases) domain.

2.4. THE MERITS OF ONTOLOGIES

Recently the results of gathering vocabularies and structuring domains been recognised as a valuable effort in its own right, deserving of attention (Wiederhold, 1994, p. 7). In general, we can say that ontologies may contribute to the following five areas (cf. Uschold & Gruninger, 1996).

(1) *Domain-theory development.* Because an ontology explicitly states the building blocks of particular domains, it can be used for the analysis, comparison, and development of domain theories. An example of this kind of ontology use can be found in Sim and Rennels (1995).

(2) *Knowledge acquisition.* Ontologies describe and structure the entities and relations that need to be acquired for the domain under consideration. Examples of this kind of ontology use are CUE (Van Heijst & Schreiber, 1994; Van Heijst, 1995), and MOBAL (Morik et al., 1993).

(3) *System design.* Ontologies are reusable constructs in the design of knowledge systems because they can be used to represent the invariant assumptions underlying different knowledge bases in the same domain. As such, they can be considered as initial building blocks of the knowledge base under construction. An example of this kind of ontology use can be found in the GAMES methodology (Van Heijst, 1995).

(4) *System documentation.* Ontologies provide a meta-level view (vocabulary, structure) on their application domain which facilitates adequate system documentation for end-users. An example of this kind of ontology use is found the Cyc project (Lenat and Guha, 1990).

(5) *Knowledge exchange.* Ontologies can be used to define assumptions that enable knowledge exchange between different agents. This can be done either by taking one ontology as a standard (e.g., Kuokka et al., 1993; Fox & Gruninger, 1994) or by mapping between individual agent ontologies (e.g., Visser et al. 1997).

So far, in the legal domain ontologies have mainly been used for knowledge acquisition and system design, and to a lesser extent, for domain-theory development.

3. Legal Ontologies

Conceptualising a domain is inherent to making a knowledge system. Most research in the field of Artificial Intelligence and Law uses some conceptualisation of (a fragment of) the legal domain. Although much work in conceptualising the legal domain has been done in legal theory (e.g., Hart, 1961; Kelsen, 1991), very few legal ontologies have been reported today (cf. Visser & Winkels, 1997). That is, few authors have explicitly specified their conceptualisation of the legal domains in a (semi-) formal language. In this section we describe four conceptualisations that have been made explicit in such a language: McCarty's Language for Legal Discourse (section 3.1), Stamper's norma formalism (Section 3.2), Valente's functional ontology of law (Section 3.3), and the ontology of Van Kralingen and Visser (Section 3.4). In the remainder of this article we use the following abbreviations;

- LLD: for the ontology underlying McCarty's Language for Legal Discourse (McCarty, 1989),
- NOR: for the ontology underlying Stamper's norma formalism (Stamper, 1991, 1996),
- LFU: for Valente's Functional Ontology of Law (Valente, 1995; Breuker et al., 1997),
- FBO: for the Frame-Based Ontology of Van Kralingen and Visser (Van Kralingen, 1995; Visser, 1995; Visser & Bench-Capon, 1996a, b; Van Kralingen, 1997).

We note that only the last two conceptualisations have been proposed as ontologies of the legal domain (and formally described in the ONTOLINGUA language), McCarty's work is considered an early attempt to conceptualise the legal domain using a (semi-) formal language, Stamper's discussion of NORMA is regarded as a discussion of ontological assumptions specified in a semi-formal language.

The selection of the four ontologies discussed in this paper is arbitrary to a certain extent. It can be argued that almost all researchers in the field of AI and Law have their own conceptualisations of the legal domain. Therefore, the list could easily be extended with other work in AI and Law (as long as the conceptualisations are intended for covering all legal sub domains and legal tasks and they are designed for the purpose of knowledge-system design). As far as our list is concerned, at the time of this research only two of the four ontologies were actually proposed as ontologies and are described in a dedicated ontology language (viz. LFU and FBO). The other two proposals (viz. LLD and NOR) are representational formalisms from which we (viz. the authors) have derived some of their underlying ontological assumptions. This research is a first attempt to systematically compare different conceptualisations of the legal domain for the purpose of legal-knowledge systems; we do not claim the ontologies to be representative for all conceptualisations that can be found in the field of AI and Law. They are, however, the most extensive and available treatments in the literature. LFU and FBO have been described both at

book length and LLD and NOR have been described in a series of articles. For some preliminary descriptions of other current legal ontologies see (Visser & Winkels, 1997).

3.1. MCCARTY'S LLD

McCarty (1989) has proposed a language for legal discourse (LLD). He considered the language to be a first step towards a general applicable representation language for legal knowledge. Although LLD itself is a representational language and not an ontology it clearly reveals a generic conceptualisation of the legal domain. We confine ourselves to this conceptualisation.

The basic components of LLD are atomic formulae and rules. Together they allow the creation first-order expressions. Modalities, such as time and permissions, are stated as second-order expressions. Below we discuss atomic formulae, rules and modalities in turn. For more details on LLD we refer to McCarty (1989, 1993), and Schlobohm and McCarty (1989).

3.1.1. *Atomic formulae*

Atomic formulae are merely predicate relations used to express factual assertions, such as 'O1 is the ownership of actor A having property P', and, 'company C has issued stocks S'. Terms in these predicates are ordered sorted, that is, variables and constants belong to a sort and sorts are ordered hierarchically. Every instantiated predicate relation is treated like an individual object (viz. a constant, or a variable). This allows to refer to individual predicate instantiations in a convenient manner (see also: modal operators). A distinction is made between *count terms* (to express tangible objects, such as houses, and persons) and *mass terms* (to express intangible objects, such as cash, and stock) (cf. tangible objects; Lenat and Guha, (1990)). Mass terms are expressed by attaching quantitative measures to them, such as value, and volume.

3.1.2. *Rules*

Rules are formed by connecting atomic formulae with logical connectives. They have a left-hand side which is an atomic formula, and a right-hand side which is a compound expression. The compound expression determines the type of rule involved. There are five types of rules: (1) horn clauses, (2) 'horn clauses' with embedded implications, (3) 'horn clauses' with embedded negations, (4) default rules, and (5) prototype-and-deformations. Below, we briefly discuss the kind of legal statements that can be expressed by these rules. The first – and most important – type of rule is the horn clause (in which the left-hand side of the rule is an atomic formula, and the right-hand side is a conjunction of atomic formulae). The horn clause allows the expression of ordinary definitions, such as 'an employee is a natural person, under the age of 65, with a labour relationship'. The second

type of rule is the horn clause with embedded implications (in their right-hand side) allow expressions such as ‘C is a sterile container if we know that if there is a bug B inside C, B is dead’. The third and last type of rule is the horn clause with embedded negations (in their right-hand side only) which are used to express statements such as ‘P is an unowned property if for every agent A, it is not the case that A owns P’. (4) Default rules allow expressions such as ‘A new-born infant who is found abandoned in the United Kingdom shall, unless the contrary is shown, be deemed to have been born in the United Kingdom’. (5) To allow disjunctive (or existential) definitions in one rule, such as ‘A Christmas block is either a red block or a green block’, LLD facilitates the expression of so-called prototypes and deformations. This is a proof-procedural technique to avoid arbitrary disjunctive proofs. Basically, one disjunction is marked prototypical, and the others are treated as deformations of the prototype.

3.1.3. *Modalities*

Modalities are stated as second-order expressions. Currently, the following modalities are supported: time, events and actions, and deontic expressions (McCarty, 1989). To express temporal statements LLD recognises states. A state essentially is the (temporal) reification of a predicate relation. Predicate relations can be reified both with points of time, as well as with intervals (two points of time). Changes in states are realised by events. Events are either elementary (viz. a state-change) or complex (viz. elementary events connected by the operations of disjunction, sequential and parallel composition, and universal and existential quantification applied to the elementary events). An action is the relation between an actor and an event. With regards to deontic statements, LLD supports four modal operators: permitted (P), forbidden (F), obligatory (O), and enabled (E). Deontic statements are formed by the combination of a name, a (possibly negated) modal operator, a condition, and an action. We note that, because there are four deontic operators which can all be negated, there are, in principle, eight different deontic modalities of an action (negating the actions is possible as well, but – to the best of our knowledge – McCarty has not elaborated on this possibility).

3.2. STAMPER’S NORMA FORMALISM

Stamper has criticised the use of traditional logics for the representation of (legal) knowledge because they suffer from some important semantic problems (Stamper, 1991). Briefly stated, traditional logics rely on symbolic representations that have only a very weak connection to the real-world concepts they intend to denote. In particular, symbolic representations rely (according to Stamper – invalidly) on notions such as truth, individuality, and identity. Accordingly, expressing legal knowledge in the form of rules is an over simplification of what legal knowledge is about.

To overcome these problems Stamper argues that there is need to escape from the frame of reference within which the classical logic is created (Stamper, 1991, p. 229). Building on his LEGOL work (see Stamper, 1980), he proposed the NORMA formalism (Stamper, 1991). NORMA, which means 'logic of norms and affordances', is based on two main philosophical assumptions: (1) there is no knowledge without a knower, and (2) the knowledge of a knower depends on his behaviour (Stamper, 1996). Using NORMA (henceforth: NOR) the entities in the world are described by their behaviour rather than by assigning them an individuality and truth values. The main ontological concepts are (a) *agents*, (b) *behavioural invariants*, and (c) *realisations*. Below, we discuss them in turn (for a more extensive discussion we refer to Stamper (1991, 1996)).*

3.2.1. *Agents*

An agent is an organism standing at the centre of reality. It gains knowledge, regulates, and modifies the world by means of actions. For its actions the agent takes responsibility. The concept of an agent can be extended to include groups, teams, companies, social agents or even nation states.

3.2.2. *Behavioural invariants*

One of the underlying ideas of NOR is that entities in the world are described by features that remain invariant over some time. Also, it is assumed that these features are found in the behavioural characteristics of these entities. For instance, a cup is described by the ability to hold liquids, the noise it makes in hitting various surfaces, the visual shape it displays etc. To capture this, NOR uses the construct of a behavioural invariant. A behavioural invariant is a description (e.g., using verbs, nouns, or adjectives) of a 'situation' whose features remain invariant. Here, a situation loosely denotes some knowledge of the world, such as an object (e.g., a cup, a piano) or a state of affairs (e.g., walking, paying).

3.2.3. *Realisations*

Agents realise situations by performing actions. The realisation of a situation – a realisation – is specified as the combination of (1) an agent and (2) a behavioural invariant, shortly written as Ax (the situation, denoted by behavioural invariant x, that is realised by agent A). An example of a realisation Ax is *John walks*. Different kinds of realisations are recognised, for instance, Ax^*(denoting the ability of A to realise x), $Ax@$ (denoting the authority of A to realise x), $Ax+$ (denoting A starts to realise x), $Ax-$ (denoting A finishes the realisation of x), and $Ax\#$ (denotes that x can be divided into individuals, cf. classes and objects).

* A full appreciation of Stamper's theory requires a more extensive discussion than the – necessarily very brief – description presented in this article.

We have attempted to compile Stamper's 1991 and his 1996 article although there are some notable differences between both articles. When confusion could arise, we have used his 1996 article.

By combining behavioural invariants composite realisations can be made. We here mention the most important composite realisations: Axy (denoting that A cannot realise y without first realising x), $A.x.y$ (denoting that x is a part of A and y is a part of x), $A(x \textit{ while } y)$, $A(x \textit{ orwhile } y)$, $A(x \textit{ whilenot } y)$, $A(x \textit{ whenever } y)$ (denotes that x is realised whenever y is realised); $A(x \text{ then } y)$ (denotes that if x is realised then y is realised), A "Bx" (denoting that an agent A can tell another agent B to bring about x, for instance by commanding or suggesting), and $A(a : b : c) \rightarrow d)$ denoting that a, b, and c are instances of d.

3.3. VALENTE'S FUNCTIONAL ONTOLOGY OF LAW

Valente's ontology of law (1995) is based on a functional perspective of the legal system. The legal system is considered an instrument to change or influence society in specific directions, determined by social goals. Its main function is reacting to social behaviour. This main function can be decomposed into six primitive functions, each corresponding with a category of primitive legal knowledge in LFU. Accordingly, LFU distinguishes six categories of legal knowledge: (a) normative knowledge, (b) world knowledge, (c) responsibility knowledge, (d) reactive knowledge, (e) meta-legal knowledge, and (f) creative knowledge. Below, we discuss each of these categories (an ONTOLINGUA specification of the ontology is given in Valente, (1995)).

(a) *Normative Knowledge*

Normative knowledge is characterised as knowledge that defines a standard of social behaviour. It thereby prescribes behaviour of the people in society. The standard is defined by issuing individual norms, expressing what ought to be the case. Because the norms contained in this category closely correspond to Hart's primary norms (Hart, 1961), Valente adopts this name for the norms in this category. It should be noted however, that not all legal norms are contained in this category, some legal norms – for instance Hart's secondary norms – are contained in the category of meta-legal knowledge (see below). Primary norms prescribe approved and disapproved behaviour by assigning a normative status to different situations. A primary norm either tags a situation according to the normative status allowed, and disallowed, or it does not give a normative status to the situation in which case the norm is said to be silent.

(b) *World Knowledge (Legal Abstract Model)*

In LFU world knowledge is legal knowledge that describes the world that is being regulated. It delineates the possible behaviour of (people, and institutions) in society, and thereby it provides a framework to define what behaviour ought (and ought not) to be performed. An example of world knowledge for traffic-law regulations is that a car, and a bicycle are driven by drivers, and that a bicycle is usually slower

than a car. Because the legislation implicitly assumes a model of the world, World Knowledge is also referred to as the Legal Abstract Model (in the remainder we refer to World Knowledge as Legal Abstract Model – or LAM – to avoid confusion between legal and non-legal world knowledge). The LAM can be considered an interface between the commonsense knowledge of people in society (cf. consensus reality knowledge, Lenat and Guha, 1990) and the normative knowledge. Within the LAM, two separate knowledge categories are distinguished: (b.1) definitional knowledge, and (b.2) causal knowledge. The definitional knowledge (b.1) is the static part, it consists of definitions of (b.1.1) legal concepts (*e.g.*, agents, objects), (b.1.2) legal relations (*e.g.*, legal qualifications of actions), (b.1.3) a case (viz. the problem case under investigation), (b.1.4) circumstances (viz. the grounded facts, or, building blocks of a case), (b.1.5) generic cases (viz. typical generic legal cases), and (b.1.6) conditions (viz. the building blocks of the generic legal cases). Together these constructs provide a vocabulary which can be used to describe the relevant aspects of the world under a specific perspective taken by the legislator. The causal knowledge (b2) is the dynamic part, describing the behaviour of people in society in terms of the definitional knowledge. This part of LFU has not been elaborated thus far.

(c) *Responsibility Knowledge*

Responsibility knowledge is legal knowledge that either extends (assigns), or restricts the responsibility of an agent for its behaviour. Its function is to provide the legal means to reject the common idea that someone is only responsible for what one causes. Restricting or assigning responsibility is done by (dis)establishing a link between the violation of a norm and an agent which is to be considered responsible for this violation.

(d) *Reactive Knowledge*

Reactive knowledge is legal knowledge that specifies which reaction should be taken (and how) if an agent violates a primary norm. Usually, this reaction is a sanction but it can be a reward as well.

(e) *Meta-Legal Knowledge*

Meta-legal knowledge is legal knowledge about legal knowledge, or, legal knowledge that refers to other legal knowledge. This category of legal knowledge is roughly equivalent to Hart's secondary rules; it includes norms that refer to primary norms. Two basic functions are supported by meta-legal knowledge. First, it regulates the dynamics of the legal system, for instance, by prescribing how to make amendments, and how to issue new primary norms). Second, it provides mechanisms to solve conflicts between instances of legal knowledge. These two functions lead to the distinction of four sub categories of meta-legal knowledge: (e1) norm data, (e2) ordering norms, (e3) normative default, and (e4) validity

knowledge. Norm data (e1) includes information about norms, such as their scope of application, their type, their place in the norm hierarchy, their power origin, their promulgation, and the norm goal. Ordering norms (e2) are norms that determine how to solve conflicts. Examples are the well-known meta-norms: *Lex Specialis derogat Legi Generali*, *Lex Superiori derogat Legi Inferiori*, and *Lex Posteriori derogat Priori*. Normative default (e3) determines what the legal status is of behaviour in case all primary norms are silent (see: normative knowledge). Finally, validity knowledge (e4) determines whether legal knowledge is valid. It comprises two types of norms: empowering norms (determine conditions under which norms can be promulgated and considered valid), and derogating norms (determine conditions under which existing norms can be rendered invalid).

(f) *Creative Knowledge*

Creative knowledge is legal knowledge that allows the creation of previously non-existent legal entities. It usually is stated in imperative terms, designating an entity (e.g. a governmental committee, or, a contract between two parties) that previously did not exist to come into being from a certain point of time.

3.4. VAN KRALINGEN'S AND VISSER'S ONTOLOGY

Van Kralingen (1995) and Visser (1995) have studied development techniques for legal knowledge systems. One of the ideas underlying their work is that robust (conceptual and formal) ontologies of the legal domain are necessities for reducing the task-dependency of legal knowledge specifications. Although there are some minor differences between the (conceptual) ontology as defined by Van Kralingen, and the (formal) ontology as defined by Visser, their similarities allow us to treat them as one ontology.

The main ontological distinction in fbo concerns the *legal ontology* and the *statute-specific ontology*. The distinction is based on the observation that some parts of an ontology are reusable across different legal subdomains. Terms such as norm, procedure and definition are likely to be found in any legal (sub)domain. Other terms, such as job-termination reason, and knife are less likely to be found in all different legal subdomains. Roughly stated, the legal ontology is defined such that all its terms can potentially occur in any legal subdomain. The statute-specific ontology contains the terminology which does not meet that criteria. Below, we discuss these two ontologies in turn, a more elaborate discussion can be found in Van Kralingen (1995) and Visser (1995). An ONTOLINGUA specification of the legal ontology is given by Visser and Bench-Capon (1996).

3.4.1. *Legal ontology*

The legal ontology, in contrast to the statute-specific ontology, is the generic and reusable part of the ontology. It divides legal knowledge over three distinct entities:

norms, acts and concept descriptions. For each of these entities the ontology defines a frame structure that lists all attributes relevant for the entity. Note that this frame structure is not an implementation commitment but simply a useful notational device for expressing the types of entities in the ontology and the attributes they possess. Norm, act, and concepts are briefly discussed below.

(a) *Norms* are the general rules, standards and principles of behaviour that subjects of law are enjoined to comply with. In the ontology a norm comprises the following eight elements: (1) a norm identifier (used as a point of reference for the norm), (2) a norm type (either norm of conduct or norm of competence), (3) a promulgation (the source of the norm), (4) the scope (the range of application of the norm), (5) the conditions of application (the circumstances under which the norm is applicable), (6) the norm subject (the person or persons to whom the norm is addressed), (7) the legal modality (either ought, ought not, may or can), and (8) the act identifier (used as a reference to a separate act description).

(b) *Acts* represent the dynamic aspects which effect changes in the state of the world. Within the category of acts we make two distinctions. The first distinction is between *events* and *processes*. Events represent an instantaneous change between two states, while processes have duration. The second distinction is between *institutional acts* and *physical acts*. The former type of acts are considered legal (institutional) versions of the (physical) acts that occur in the real world (more precisely: an institutional act is a legal qualification of a physical act). We note that these two distinctions result in four different types of acts. All acts are assumed to have the following thirteen elements: (1) the act identifier (used as a point of reference for the act), (2) a promulgation (the source of the act description), (3) the scope (the range of application of the act description), (4) the agent (an individual, a set of individuals, an aggregate or a conglomerate), (5) the act type (both basic acts, and acts that have been specified elsewhere can be used), (6) the modality of means (material objects used in the act or sub acts; e.g., a gun), (7) the modality of manner (the way in which objects have been used or sub acts have been performed) (e.g., aggressively), (8) the temporal aspects (an absolute time specification; e.g., on the first of August, on Sundays, at night, etc, but not: during a fire, after the King dies, etc), (9) the spatial aspects (a specification of the location where the act takes place; e.g., in the Netherlands, in Leiden, on a train), (9) the circumstantial aspects (a description of the circumstances under which the act takes place; e.g., during a war), (10) the cause of the action (a specification of the reason(s) to perform the action, e.g., revenge), (11) the aim of the action (the goal visualised by the agent; e.g., with a view to unlawfully appropriate an object), (12) the intentionality of an action (the state of mind of the agent; e.g., voluntary), and (13) the final state (the results and consequences of an action; e.g., the death of the victim).

(c) *Concept descriptions* deal with the meanings of the concepts found in the domain. They may be definitions or deeming provisions and can be used to determine definitively the meaning of a notion, either by, as in the case of the former, providing necessary and sufficient conditions, or, as in the case of the latter, establishing a legal fiction. Another type of concept is the factor, which may either establish a sufficient condition, or indicate some contribution to the applicability of the concept, as discussed above. Finally there are meta concepts which are provisions governing the application of other provisions. Concept descriptions comprise the following seven elements: (1) the concept to be described, (2) the concept type (definition, deeming provision, factor, or meta), (3) the priority (the weight assigned to a factor), (4) the promulgation (the source of the concept description), (5) the scope (the range of application of the concept description), (6) the conditions under which a concept is applicable, and (7) an enumeration of instances of the concept.

3.4.2. *Statute-specific ontology*

The legal ontology of FBO contains constructs that are thought to be generic for the legal domain. That is, norms, acts and concept descriptions are considered to be present in any legal domain. Modelling a legal sub domain also involves deciding upon numerous ontological questions. For instance, is it necessary to distinguish between male and female employers in the Unemployment Benefits Act? This motivates the distinction between the legal and the statute-specific ontology. The statute-specific ontology consists of predicate relations that are used to complement the terminology for norms, acts and concept descriptions. Van Kralingen and Visser argue that the statute-specific ontology cannot be reused for other legal subdomains, and should always be created for each legal sub domain under consideration. The statute-specific ontology should not be confused with the application knowledge base: the statute-specific ontology shows, for example, that the distinction between male and female workers is important in a particular statute. The distinction can then be *used* in the knowledge base. The statute-specific ontology states the vocabulary with which the knowledge base is constructed.

4. Criteria for Ontology Comparison

Ontologies can be written for many different purposes. Even if we consider one particular domain, we are likely to arrive at different ontologies if we build ontologies for different purposes. In other words, ontologies are dependent on the purposes they are made for (Chandrasekaran & Josephson, 1997). For this reason we deem the comparison of ontologies to be fruitful only if they are designed (roughly) for the same purpose. Thus, we have to examine the purposes the ontologies under investigation are designed for. Assessing the adequacy of the ontologies has to be done in the context of this purpose.

All four ontologies have been proposed as a basis for the design of legal knowledge systems. All authors claim – more or less explicitly – to cover what is to be

distinguished as the building blocks of legal knowledge (although Stamper focuses on social domains in general he considers the legal domain to be an instance of a social domain). Clearly, they make several commitments to the legal domain. Regarding task and method commitments the situation is different. All four ontologies are meant to make very few task and method commitments. McCarty's LLD is explicitly aimed to be a framework that supports a broad range of tasks: 'Clearly, if a language of this sort could be developed, it would provide a uniform framework for the construction of a legal analysis/planning/retrieval system' (McCarty, 1989, p. 180). Stamper's NOR is not reported to be tied to certain types of tasks. Moreover, the language is 'potentially capable of capturing the complexities of real social behaviour' (Stamper, 1991, p. 235). The LFU ontology of Valente is thought to be an instrument to make domain knowledge reusable for different tasks and methods (Valente, 1995, pp. 23–24). The FBO ontology of Van Kralingen and Visser is explicitly designed so as to make as few task and method commitments as possible (Visser, 1995; pp. 11, 102–103).

The discussion above suggest that we should compare the ontologies with respect to their suitability in the design of legal knowledge systems for different legal tasks, methods and legal sub domains. This purpose can be used to draw up a set of ontology-comparison criteria. We note that different authors tend to distinguish different ontology criteria (cf. Bench-Capon, 1990; Schreiber, 1992; Gruber, 1993a; Valente, 1995; Visser, 1995; Van Kralingen, 1995; Uschold & Gruninger, 1996). After an analysis of the AI, and AI-and-Law literature, we adopt the following typology of criteria:

(1) *Epistemological adequacy*: The epistemological adequacy of an ontology refers to the degree to which the ontology resembles the cognitive framework of the human problem solver (Schreiber, 1992, p.26). Thus, it is a measure of the extent to which lawyers use the concepts and relations distinguished in the ontology. The epistemological adequacy comprises five sub criteria. We illustrate these criteria by listing some relevant questions that can be used to apply them.
   (a) *epistemological clarity*: Do all concepts and relations in the ontology have a clear and unequivocal meaning? Does the ontology effectively communicate the intended meaning of the defined concepts and relations (cf. clarity; Gruber, 1993a, p. 2)?;
   (b) *epistemological intuitiveness*: Do the ontological concepts and relations provide a vocabulary that matches the intuition of the experts in the domain (cf. notational convenience; Bench-Capon, 1990, p. 17; cf. conceptual distance, Falkenberg, 1989)?;
   (c) *epistemological relevance*: Are all the concepts and relations in the ontology relevant for modelling legal tasks, methods, and domains (cf. relevance, Bench-Capon, 1990, p. 17);
   (d) *epistemological completeness*: Does the ontology cover all legal concepts and relations that may be relevant for any combination of legal task, method

and subdomain? Are there entities that cannot be modelled with the entities distinguished in the ontology?;

(e) *discriminative power*: Does the ontology have enough discriminative power in that it provides distinctions at a sufficiently high granularity level (viz. sufficient detail)?

(2) *Operationality*: The operationality of an ontology refers to the effort required to implement the ontological concepts and relations in a representational language (cf. Schreiber, 1992, p. 122; cf. practical validity, Valente, 1995, p. 43). Therewith, the criterion is a measure for the ease with which the concepts and relations in the ontology can be used as a basis for an operational language given a legal task, method and domain. We distinguish three sub criteria and list some relevant questions to apply them.

(a) *encoding bias*: Does the ontology rely on symbol-level choices? An ontology should be specified at knowledge level, an encoding bias results when a representation choice is made purely for the convenience of notation or implementation (cf. encoding bias; Gruber, 1993a, p. 3);

(b) *coherence*: Is the ontology coherently defined in that it is internally consistent? An ontology is not coherently defined (incoherent) if a sentence can be inferred from the definitions that is inconsistent with another definition or (informally specified) example (cf. coherence; Gruber, 1993a, p. 3).

(c) *computationality*: Does the ontology provide a suitable basis for (computational) representation, and is this representation computationally adequate?

(3) *Reusability*: The reusability of an ontology refers to the degree in which the ontology can be reused to conceptualise new legal tasks, methods and subdomains (this includes the extendibility of the ontology, being the degree in which the ontology can be extended without revising existing elements of the ontology; Gruber 1993a, p. 3). In our comparison of the ontologies (see: section 5), we distinguish two sub criteria.

(a) *task-and-method reusability*: is the ontology dependent on certain (types of) tasks and methods, or alternatively, to what extend is the ontology reusable for various methods and tasks?;

(b) *domain reusability*: is the ontology dependent on certain (types of) legal subdomains, or alternatively, to what extend is the ontology reusable for various legal subdomains?

Any list of ontology criteria is open to discussion; arguably criteria can be removed from the list or new criteria can be added to the list. However, we believe the above list of criteria to cover most common ontology criteria mentioned in the AI (and Law) literature. It should be noted that criteria such as ‘how well does the ontology communicate ontological commitments between different research groups’ can be derived from the list above. This criterion would, for instance, be covered by the (more primitive) criteria: *epistemological clarity, epistemological convenience,* and *coherence*. Similarly, the criterion ‘how well does the ontology

	Basic knowledge categories
LLD	Atomic formula, Rules, and Modalities*
NOR	Agents, Behavioural invariants, and Realisations
LFU	Normative knowledge, World Knowledge, Responsibility knowledge, Reactive knowledge, Meta-legal knowledge, and Creative knowledge
FBO	Norms, Acts, and Concept Descriptions

Figure 1. The four ontologies and their basic knowledge categories.

support the knowledge acquisition phase?' would be covered by criteria such as: *epistemological clarity*, *epistemological relevance*, *epistemological completeness*, and *discriminative power*.

5. The Four Ontologies Compared

The typology of criteria discussed in the previous section allows us to compare the four legal ontologies described in Section 3. We stress that the opinions given here are the personal and subjective opinions of the authors (when we phrase the opinion of other authors, this is stated explicitly). The purpose of the comparison is not to identify the "best" or most adequate ontology, but rather to draw out the similarities and differences so as to understand what choices are made in committing to an ontology, and the considerations that inform these choices. As can be noted from the discussion in Section 3, the four ontologies differ substantially in the basic categories they distinguish. In Figure 1 we summarize the four legal ontologies with respect to their basic knowledge categories†.

In Sections 5.1 through 5.3 we discuss the ontologies focussing on each of the three main criteria, in turn. The scope of this article does not allow us to apply the criteria to each of the ontologies in their most detailed level. Hence, the discussion below is necessarily abstract and incomplete.

5.1. EPISTEMOLOGICAL ADEQUACY

The epistemological adequacy of an ontology refers to the degree in which the ontology resembles the cognitive framework of the human problem solver. In Section 4 we distinguished between epistemological clarity, epistemological intuitiveness, epistemological relevance, epistemological completeness, and discriminative power.

* LLD has not been defined as an ontology. Arguably, these categories are representational categories rather than basic knowledge categories. However, since no alternative classification of legal knowledge is given we use these categories as the basic knowledge categories in LLD.

† With respect to FBO we confine ourselves to the legal ontology.

5.1.1. *Epistemological clarity*

The clarity of an ontology refers to the extent in which the concepts and relations defined in the ontology have a clear and unequivocal meaning. In principle, this is a measure for the degree in which the author of the ontology succeeds in providing an adequate (possibly textual) definition for the concepts and relations distinguished.

In LLD, the categories (viz. atomic formula, rules, and their modal versions) seem to be rather distinct. Probably, this is because of their reference to the structure of the legal knowledge. We consider the interpretation of some expressions in nor to be less clear, for instance, it is not clear what the expression $A(x, y)$ denotes (defined by Stamper as: '*A* realising both x and y', but used in an example to express that *A has* both x and y). Also, the distinction between certain expressions is not always clearly defined. This applies, for instance, to the distinction between the statements 'A(x whenever y)' and 'A(y then x)'. More clearly defined is LFU although we consider the boundaries of both the categories meta-legal knowledge and creative knowledge to be somewhat vague. The category of meta-legal knowledge seems to be dedicated to all knowledge that is meta with respect to the category of normative knowledge (viz. expressing knowledge about (primary) norms). However, knowledge that is meta with respect to normative knowledge can be located in other categories as well. In particular, this holds for responsibility knowledge, reactive knowledge (arguably both meta knowledge because they express knowledge about the violation of norms) and creative knowledge (arguably meta knowledge because it expresses knowledge about the creation of norms). Consequently, it is not always clear in which of the four categories a fragment of meta-level knowledge should be placed. It seems as if the category of meta-legal knowledge is a sort of catch-all category for knowledge not residing in the categories responsibility knowledge, reactive knowledge, and creative knowledge. FBO appears to suffer from a similar problem as LFU. As with the category of meta-legal knowledge (in LFU) the circumstances-element from the act (in FBO) seems to be a sort of catch-all category. That is, there are eight act elements that can be used to specify different aspects of act circumstances. However, there is also a separate circumstance element in the act. It is not always obvious when a fragment of legal knowledge should be considered as the element circumstances or as one of the other act-circumstances elements, such as the modality of means, the modality of manner, the spatial or temporal aspects, the cause, aim or intentionality.

5.1.2. *Epistemological intuitiveness*

The epistemological intuitiveness of LLD for the legal domain is limited. The main categories of LLD seem to originate from representational constructs (logic) rather than legal theory. NOR adopts a legal theory but its ontological constructs are sometimes counter intuitive in that all entities have to be described by their behaviour (instead of by their features). Expressions in NOR can be rather confusing and one need to be aware of misinterpretation. Consider the following expressions:

John (car, dog) means 'John – perceives that he – has both a car and a dog', and the expression *John (car, burns)* means 'John has both a car and John burns' (nothing is said about his car burning or not). Another example of a counter-intuitive expression is *John book finish* which means 'the period of existence of the situation in which John has a book has finished. LFU can be criticised for the names chosen to denote the categories* In particular, LFU recognises particular norms (viz. the equivalence of Hart's secondary rules) that are considered to be part of the category meta-legal knowledge instead of the (perhaps more obvious) category of normative knowledge. Also, we consider the name world knowledge to be somewhat counter intuitive since it is a part of legal knowledge, and its name suggests that this knowledge is a form of commonsense knowledge†. In contrast to LFU, FBO aims to achieve cognitive validity. The real intuitiveness of an ontology will differ between different practitioners, each having their own preferences.

5.1.3. *Epistemological relevance*

The epistemological relevance of an ontology refers to the degree in which the categories distinguished are relevant for all tasks and sub domains. The distinctions in LLD all seem to make sense in that the constructs all support the expression of certain 'distinguishable' types of legal knowledge. This also holds for LFU and FBO (although the statute-specific ontology of FBO is obviously not relevant for most other legal domains). With respect to nor the distinction between 'A(*x whenever y*)' and 'A(*y then x*)' is not always clear. The relevance of such a distinction can only be assessed given a specific ontology purpose. That is, it should be much more specific than attempting to cover all legal knowledge for the purpose of legal-knowledge systems.

5.1.4. *Epistemological completeness*

Assessing the epistemological completeness of an ontology is problematic because in order to determine whether an ontology facilitates the modelling of some piece of legal knowledge we need to identify this piece of knowledge first. This requires at least some commonly accepted theory about legal knowledge that tells us what pieces of knowledge exist in the legal domain. The problem is that we do not have such a theory. Briefly stated, there is no golden standard for the comparison. For this reason, we confine ourselves to a brief discussion covering two issues: (a) norm types, and (b) legal procedures.

---

* Regarding the intuitiveness of an ontology Valente (p. 77) remarks: '(...) people do not reason with the type of conceptualisation one embeds in an ontology. The model of legal reasoning embedded in this ontology is very unlikely to be cognitively valid, and it is probably counter-intuitive in several respects to both the average reader and the legal theorist'.

† This issue has been discussed by Hage (1996) as well.

(a) *Norm types*. Although all legal ontologies recognise norms, they differ in the types of norms they distinguish. In the ontologies norms are distinguished by their function and their normative modality. In LLD norms are distinguished by their normative modality. In principle LLD supports eight different norms: four normative modalities – (P) permitted, (F) forbidden, (O) obligated, and (E) enabled – all of which can be negated. LFU recognises norms by their function in the legal system. Explicitly recognised are: primary norms, ordering norms, empowering norms, and derogating norms. Modalities are not explicitly specified and can be chosen (Valente implements his ontology in a variant of deontic logic, distinguishing three modalities and their negations). FBO recognises norms both by their function (viz. norm of conduct, or, norm of competence) and by their normative modality (viz. ought, ought not, may, can), in total giving six norm types. NOR has a radically different approach to norms than the other ontologies. Although recognising that every norm has the form *if* ⟨*condition*⟩ *then* ⟨*some agent*⟩ *is permitted/forbidden/obliged to do* ⟨*action*⟩ Stamper argues that such a rule is only a representation of a norm. A norm is something that exists in a community. Norms are 'issued' by the behaviour of agents (cf. speech acts). They can thus be expressed as behavioural invariants which means that, in principle, every behavioural invariant may denote a norm. This, in principle, allows for a wide range of norm types.

(b) *Legal procedures*. None of the ontologies appears to have an adequate solution for norms that describe legal procedures (e.g., procedural norms of competence). Possibly, the researchers have not addressed this kind of legal knowledge in their ontologies because there was no role for such knowledge in the legal sources that were used in the construction of their ontologies. Possible, they did not address legal procedures because there are severe problems with the (declarative) specification of procedural knowledge. One of the difficulties is to find a language to express procedural knowledge in a declarative way. Related to this difficulty is the question whether legal procedures should be regarded as *control* knowledge or as *domain* knowledge (Visser, 1995).

5.1.5. *Discriminative power*

Assessing the discriminative power of an ontology is difficult for the same reason as assessing the *epistemological completeness* of an ontology. Again, we confine ourselves to some general remarks. The ontologies differ considerably in their level of detail. If we were to order the ontologies with respect to their level of detail (from least detailed to most detailed) we would get: NOR, LLD, LFU and FBO. As stressed before in FBO a distinction is made between a statute-specific ontology and a (generic) legal ontology. Therefore, the combination of the legal ontology and the statute-specific ontology reaches the highest discriminative power (at the cost of a limited reusability for the statute-specific ontology).

To conclude this section we briefly address the relation between the epistemological completeness and the discriminative power of an ontology. On the one hand we might argue that an ontology that makes few ontological distinctions has a higher completeness than an ontology that makes many distinctions. For instance, consider an ontology that only makes one distinction, say legal knowledge consists solely of norms and not-norms. This ontology has a high degree of completeness in that there are presumably no legal entities that are not covered by these two categories. However, the distinction does not provide the means to separate the different types of knowledge that can be recognised in the legal domain, it does not have much discriminative power. On the other hand, an ontology that identifies say, six particular types of norms, is likely to miss out a (possible) seventh type of norm. This will impede the ontology completeness. The paradox is that the more ontological distinctions are made to more likely it is that the ontology becomes incomplete. We conclude that there is a trade off between epistemological completeness and discriminative power of an ontology.

5.2. OPERATIONALITY

The operationality of an ontology refers to the effort required to implement the ontological concepts and relations in a representation language. In Section 4 we distinguished between the encoding bias, the coherence, and the computationality.

5.2.1. *Encoding bias*

An ontology has an encoding bias if it makes commitments to a certain representational formalism. In principle an ontology should be independent of a representational formalism (Gruber, 1992, 1993a). Hence, it should be possible to implement any ontology in any representation formalism. However, ontological commitments may prove to be more easily implementable in some representational formalisms than in others. In LFU, for instance, Valente clearly separates the description of the ontology from its implementation. However, in choosing implementation formalisms (for normative knowledge and world knowledge) the suitability of the available formalisms to model these types of knowledge plays an important role (Valente, 1995, pp. 83–84 and pp. 113–115)*. The FBO of Van Kralingen and Visser at first glance seems to commit to frames. However, the frame structures used are not hierarchically ordered and inheritance is not used. Hence, there is no significant commitment to frames, nor any other representational formalism (Visser has used prolog as an implementation language). Due to its abstraction level, NOR only shows commitments to some predicate logic expressions, not necessarily to a representational formalism. Finally, LLD shows commitments towards a representational formalism in that it makes several commitments to arrive at a computationally

* This does not rule out that these types of knowledge can be implemented in other formalisms than the ones chosen by Valente.

adequate formalism (e.g., the use of 'prototypes and deformations' to deal with disjunctive assertions).

5.2.2. *Coherence*

Regarding the coherence of the ontologies discussed here we have not encountered any inconsistencies in the ontologies. We draw the tentative conclusion that the four ontologies are coherent.

5.2.3. *Computationality*

The computationality of an ontology refers to the degree in which the ontology provides a suitable basis for computationally adequate representations. In this respect, little experience is gathered at this stage. We currently cannot assess whether the approaches are adequate or inadequate. However, we can make some general statements about the (experimental) application of the ontologies in the design of knowledge systems. FBO has been used for the implementation of two operational (prototype) knowledge systems, called FRAMER-P (planning system) and FRAMER-A (assessment system). Although the computational adequacy of both FRAMER systems is not optimal there is no adequate means to assess the computationality of the *ontology*. LFU has been partially implemented in a system called ON-LINE, which Valente describes as a 'Legal Information Server'. Valente only briefly elaborates on how an actual knowledge base is built on the basis of the ontology. LLD has been designed both from a legal theoretical perspective and from a computational perspective. Computationality thus has been a constant issue in its construction. NOR has not been implemented. Stamper does not address implementation nor computational aspects of his ontology, but NOR is not likely to be very efficient from a computational perspective.

5.3. REUSABILITY

In Section 4 we defined the reusability of an ontology as the degree in which the ontology can be reused (possibly by extending it) to conceptualise new legal tasks, methods, and domains. In general, we could say that the higher the level of detail of an ontology, the more commitments are made to tasks, methods and domains. Hence, an abstract ontology will have a greater reusability than a ontology with a high level of detail (at the cost of discriminative power). Below, we discuss this criterion by its two sub criteria.

5.3.1. *Task-and-method reusability*

As we stated in Section 4 all four ontologies are intended to make few task-and-method commitments. This suggests that all ontologies should be reusable for other tasks and methods. To compare the ontologies with respect to tasks and methods we use the taxonomy of legal tasks as described in Visser (1995) and in particular

the assessment and planning branches. We then examine whether the ontologies define entities that could play the role of the entities that are required for the tasks.

We first consider assessment tasks. All ontologies seem to be capable of supporting classical assessment (which is a straightforward matching of a problem case with necessary and sufficient conditions in definitions). The support of probabilistic assessment (collecting and weighing factors pro and contra a certain decision) requires the possibility to express probabilistic information. Whether the ontologies support this task is less obvious. FBO supports at least one form of expressing probabilistic information (viz. factors, as part of the concept descriptions). LFU does not explicitly support the expression of probabilistic information but presumably it would be part of the definitional knowledge (which is part of the Legal Abstract Model). In LLD representing beliefs is recognised as an important feature but it is to the best of our knowledge not yet supported. Despite this, LLD as such seems to be biased towards dealing with 'hard cases', stressing the dynamics of legal concepts and the role played by argument in realising that dynamic nature (see also Bench-Capon and Visser, 1996). Apparently, nor does not support probabilistic information. Exemplar-based assessment (or case-based assessment) is a form of assessment in which a case at hand is assessed according to similar cases in a case base. This type of assessment does not seem to require different ontological distinctions rather than those required for a case description. All ontologies support the definition of a case descriptions (although LLD and LFU explicitly distinguish separate ontological entities for 'exemplars' and 'case descriptions', respectively).

For planning tasks we can provide a similar discussion as for assessment tasks. All planning tasks require the specification of states and state-transitions, and all four ontologies support these constructs to some extent. However, there are considerable differences in the ontologies with respect to the specification of state-changes. LLD and FBO have explicit constructs for state-changes, inspired on the well-known STRIPS planning operators (including actors, times, and conditions before and after the state-change). nor views legal knowledge from a dynamic perspective and is centred around agents and their actions. However, NOR defines a language that allows the specification of relative temporal expressions (such as 'A happens during B'), it is not straightforward to express absolute time references. In LFU state-changes are recognised as part of the Legal Abstract Model but they have not been elaborated in much detail.

5.3.2. *Domain reusability*

The domain reusability of the ontologies is the degree in which the ontologies can be (re)used for different legal subdomains. In principle, none of the ontologies is dependent on any legal subdomain as such. However, it is likely that all ontologies have embedded assumptions that stem from the domain to which the ontology has been tested in. For instance, the prominent appearance of time reference slots in FBO, can probably be related to the analyses of the Dutch Unemployment Benefits

Act that was used in the design and evaluation of the ontology (note that the statute-specific ontology of FBO is not reusable for other legal subdomains).

6. Discussion, Conclusions and Future Research

The ontology comparison can be used to state some general remarks about legal ontologies and their usefulness in the creation of legal knowledge systems. In the end, we believe that ontologies should be evaluated more or less suitable for a particular application than merely good or bad. Below, we provide a short discussion (Section 6.1), we draw conclusions (Section 6.2), and we give some suggestions for future research (Section 6.3).

6.1. DISCUSSION

In this section we discuss four issues that arise from the comparison reported in the previous sections. In order, we discuss (a) the priorities of the knowledge types distinguished in the ontologies, (b) the role of legal theory in the ontologies, (c) the relation between ontologies and knowledge bases, and (d) the abstraction level of the ontologies.

(a) Priorities of knowledge types

The most obvious thing to note about the ontologies described above is that they are very different. At the highest level they diverge immediately. This is because of the different perspectives from which the authors begin their conceptualisation. For instance, Valente seeks a functional decomposition of the legal system considered qua system, whereas Van Kralingen and Visser seek a set of building blocks from which they can construct law qua body of knowledge. This difference in perspective leads to substantially different ontologies.

If we compare the four legal ontologies a bit further we can say that the difference in the ontologies is not so much a difference in types of legal knowledge distinguished in the ontologies but it is more a difference in priorities between these knowledge types. The ontologies have several common features but they differ in the priority attributed to these features. For instance, all four ontologies recognise the importance of actions, agents, norms, and some form of definitional knowledge. However, if we consider actions we find that they are among the three basic knowledge categories in FBO. In NOR actions also play a crucial role in that the ontology is built around agents and their behaviour. In LLD actions play an important role although they are not recognised as a *basic* knowledge category. The status of actions in LFU is different from the status of actions in the other ontologies. Although in LFU actions are described (as part of the world knowledge) they do not constitute a separate knowledge category in the (formal) ontology. Rather, actions are implied by the states they bring about (and hence, no differences are made between two different actions that bring about the same state). Another example is found in the

concept of an agent which plays a central role in NOR but in the other ontologies it is attributed to actions (LLD and FBO), or, for instance, to responsibilities (LFU). Not surprisingly, all four ontologies recognise norms as an important category of legal knowledge. Non-normative definitions play an important role in LFU, FBO, and LLD. In NOR however, their role is less pronounced. Time is addressed in all ontologies but again there is a difference in the priority it is given (as well as the way it is treated). Whereas LLD and FBO explicitly represent temporal references (viz. event calculus), in NOR time is implicitly introduced by describing state-changes (viz. situation calculus). LFU leaves the treatment of time to the knowledge engineer. Spatial aspects are only recognised as such in FBO, although the attention payed to spacial aspects is limited.

Another manifestation of similarities can be found when comparing FBO and LFU. In the former ontology we find defined as relations "event-qualification" and "process-qualification". The intention of these relations is to relate a physical event to an institutional event. This is a crucial step in this ontology since it enables physical acts (e.g., A kills B), to be classified in terms of institutional acts (e.g., A murders B, or alternatively, A manslaughters B) which are what tend to be used in norm descriptions. It is precisely these relationships between the physical and the institutional descriptions that are the subject of LFU's category of "world knowledge". Thus we can see that in both conceptualisations the transition from the physical description to the institutional description is of crucial importance. In the case of normative status, both ontologies define it as a function, although for Valente it maps from situations to a normative status, whereas in Van Kralingen and Visser it maps from a(n) (applied) norm to a normative status. There remains, however, a similarity in that we can go in the latter ontology to a situation via the act prescribed in a breached norm which has a slot giving the post-conditions of the act. A difference remains, however, in that in Valente's ontology there are three flavours of normative status, allowed, disallowed and silent, whereas Van Kralingen and Visser subsume both allowed and silent under "not breached".

(b) *The role of legal theory*

There seems to be a tendency over the years in that the influence of legal theory in ontology-related AI and Law research becomes more important. Although the nature of ontologies is such that at least some legal theory is required (recall that a legal ontology is an explicit conceptualisation – theory – of the legal domain), there are substantial differences in the theoretical underpinnings of the ontological assumptions that are made. LLD, which originates from 1989, is based on legal theory but seems to be importantly influenced by representational and proof-theoretical considerations. NOR, originating in 1991, is clearly an implementation of a legal-theoretical view on the domain, although this theory might not be commonly acknowledged in the legal field. Both 1995 approaches, FBO and LFU, both have strong links with the work of legal theorists, such as Hart, Kelsen, Ross, Brouwer, and Von Wright. We could say that the balance between

the legal-theoretical underpinning on the one hand, and the prospects of making a computationally feasible program on the other hand is shifting towards the former. In line with the aims of ontological studies the role of representation and computation becomes less important. It would, of course, be interesting also to examine the ontologies from the perspective of legal theory, as we have done from the standpoint of legal knowledge system developers. Such an examination might well yield different, or additional, criteria. It is, however, outside the scope of this paper.

(c) *Relation between ontologies and knowledge bases*

Although ontologies have become very popular in AI research, the question on what exactly is to be considered an ontology remains open (cf. Guarino & Giaretta, 1995; Uschold and Grunninger, 1996). One of the difficulties we encountered is the relation between an ontology and a knowledge base. On the one hand one can argue that an ontology and a knowledge base are distinct entities because an ontology is a meta-level description of a knowledge base. It captures assumptions underlying a knowledge base, and hence, it can be used to describe shared assumptions underlying a set of different knowledge bases. This makes the ontology a distinct entity (from the knowledge base) and suitable as a reusable component in the design of knowledge systems. On the other hand, one can argue that an ontology and a knowledge base are not distinct entities, but, in contrast, that ontology *is* a knowledge base. An ontology contains knowledge from a particular domain, often allows inference on this knowledge, and should therefore be seen as a knowledge base of that domain. These two arguments do not exclude each other, an ontology is, in fact, both a knowledge base and a meta-level description. Thus, an ontology is a *kind* of knowledge base. However, one might still question what makes an ontology an ontology? Can we, given a certain knowledge base, determine whether it is an ontology or not (and *vice versa*)? In general, we could say that an ontology is a more general description than a knowledge base. However, the VT-ontology (as contained in the ontology library of Stanford University), for instance, contains very detailed information on (instances of) elevators and its level of detail suggests it to be the knowledge base of an actual system. The heart of the problem lies in the definition of an ontology as a *meta-level* description. The problem with defining something as a meta-level description is that the definition is relative to an object-level description. This means that ontologies can only be defined relative to a given object-level knowledge base, and hence, it can only adequately be created if the (object-level) knowledge base is already present. Here, we do not address this issue in any more detail. We opine that there is no fundamental difference between an ontology and a knowledge base, whether a knowledge base is an ontology is to be judged on the basis of its degree of generality, on the completeness of the description as a conceptualisation of a domain (many knowledge bases will be incomplete if considered as an ontology), and possibly, on the existence of other knowledge bases.

(d) *Astraction level*

Related to the question about the exact meaning of an ontology is the desired abstraction level (or: level of detail/granularity level) of an ontology. The ontologies differ substantially in the abstraction level of their concepts and relations. That is, all ontologies start with very abstract concepts after which these concepts are elaborated in more detail but the ontologies differ in the level of detail of their most refined entities. NOR is probably the least elaborated ontology. The FBO, if taken to be both the generic and the statute-specific ontology, probably has the highest degree of detail. What abstraction level should a legal ontology ideally have? Also, if we are to make a library of ontologies what abstraction level should be chosen? We have encountered two competing arguments. On the one hand, it has been shown that a very detailed ontology (viz. one with a high discriminative power) is a useful tool in the acquisition and expression of domain knowledge. On the other hand, the more detailed an ontology, the more commitments are made to particular tasks, methods and (sub)domains. As illustrated in section 5.1 making ontological distinctions may involve that certain types of (legal) knowledge do not fit into the ontology any more. Hence, the more discriminative power, the less likely it is to be reusable for arbitrary tasks, methods and (sub)domains. We opine that there is no generally desirable abstraction level that should be chosen for the expression of an ontology. The approach taken in FBO, using both a (reusable) abstract and a (non-reusable) detailed ontology, may be a suitable compromise for building libraries of ontologies.

6.2. CONCLUSIONS AND FINDINGS

After the comparison of the four legal ontologies we arrive at the following conclusions and findings:

- Different authors create substantially different conceptualisations of the legal domain despite the fact that their purposes are similar (Section 5).
- There is no agreement on the issue as to what are the most elementary building blocks of legal knowledge (Section 5).
- Differences in ontologies are not so much differences in knowledge types distinguished but differences in the priority these knowledge types are given (Section 6.1a).
- Terms defined in an ontology vary in their reusability across legal subdomains. Some terms are reusable across all legal subdomains, others are specific for a particular application (and thus not reusable) (Section 5).
- Epistemological completeness and discriminative power is difficult to assess since there is no golden standard which states which knowledge types should exist (Section 5.1).

- Some criteria to assess the quality of an ontology are subjective (Section 5.1). In particular, this holds for epistemological clarity, epistemological intuitiveness.
- The epistemological adequacy of an ontology cannot adequately be assessed without a precise description of the purpose of the ontology. The common purpose between the ontologies discussed here (their support in the creation of legal-knowledge systems) is too abstract. It is recommended that all ontologies are published with a clear and detailed statement of their purpose and applicability (Section 5.1).
- Non of the ontologies seems to have adequate provisions to specify legal procedures (section 5.1).
- All ontologies seem to have provisions to support certain specific problem-solving tasks. Few ontologies are explicit as to which tasks are supported (Section 5.3).
- In the conceptualisation of legal knowledge there is a shift in balance from representational/computational issues to legal theoretical underpinning (Section 6.1b).
- The most elaborated entities distinguished in each of the four ontologies have substantially different levels of abstraction (Section 6.1d).
- There is a trade-off between reusability and expressive power (Section 5.1, 5.2 and 6.1d).

6.3. FUTURE RESEARCH

The four conceptualisations discussed in this paper, however, have not received widespread attention; for the most part their use has been confined to a single author or group. The result is that while we can examine these conceptualisations, and identify their various attractions, we are left with questions which cannot be resolved until a body of work has been produced, which can form the basis for a reliable evaluation. What is needed is a programme of work which will include the following activities:

(1) The explicit expression of conceptualisations in a complete and rigorous form. Ideally the expression of different conceptualisations should be readily comparable. The trend to use a common ontology language such as ONTOLINGUA to express ontologies is very welcome.
(2) The conceptualisations need to be used across a variety of different applications. In so far as the ontology is meant to be generic, it is important to establish – as was done for the "law as definitions" conceptualisation (Bench-Capon & Visser, 1996) – where it can and cannot be effectively applied.
(3) The use of different conceptualisations of the same area of law. This is important if we are to get a real understanding of the difference that different conceptualisations make.

(4) Discussion of the conceptualisations from a jurisprudential standpoint. At present the conceptualisations are predominately produced by those who main interest is computational. The role of legal theory in this enterprise is to uncover the viability of these conceptualisations from the point of view of legal theory.

This is a programme of work which is too much for a single researcher, or a single group. Moreover the work needs to be done from a plurality of perspectives. Hopes for progress in AI and Law require that the community works in a more synergistic way, explicitly building on and extending the work of others. The obvious means to this end is the creation of libraries of legal ontologies, indexed on task, legal subdomain, applicability, and abstraction level.

References

Allen, R. E. 1990. *The Concise Oxford Dictionary of Current English*, Clarendon Press, Oxford, England.

Bench-Capon, T. J. M. 1990. *Knowledge Representation; An Approach to Artificial Intelligence*, APIC series, No. 32, Academic Press, London, United Kingdom.

Bench-Capon, T. J. M. & Visser, P. R. S. 1996. Deep models, ontologies and legal knowledge based systems, in Van Kralingen et al. (eds.), *Proceedings of the Ninth International Conference on Legal Knowledge-Based Systems*, (JURIX'96), Tilburg, The Netherlands, pp. 1–14.

Bench-Capon, T. J. M. & Visser, P. R. S. 1997. Ontologies in legal information systems; the need for explicit specifications of domain conceptualisations, *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*. (ICAIL '97), Melbourne, Australia, pp. 132–141.

Breuker, J. A. & van de Velde, W. 1994. In J. A. Breuker and W. van de Velde (eds.), *Common KADS Library for Expertise Modelling, Reusable Problem Solving Components*, IOS Press, Amsterdam, the Netherlands.

Breuker, J. A., Valente, A., & Winkels, R. G .F. 1997. Legal ontologies: A functional view, P. R. S. Visser and R. G. F. Winkels (eds.), *Proceedings of the First International Workshop on Legal Ontologies*, (LEGONT '97), 4 July 1997, University of Melbourne, Law School, Melbourne, Australia, pp. 23–36.

Chandrasekaran, B. & Josephson, J. R. 1997. The ontology of tasks and methods, *Working Notes of the AAAI Spring symposium on Ontological Engineering*, Stanford University, CA, USA, pp. 9–16.

Falkenberg, E. D. 1989. *Informatiesystemen – aan de kinderschoenen ontstegen?*, University of Nijmegen, Nijmegen, The Netherlands (in Dutch).

Fox, M. S. & Gruninger, M. 1994. Ontologies for enterprise integration, *Proceedings of the 2nd Conference on Cooperative Information Systems*, Toronto, Ont, Canada.

Genesereth, M. R. & Fikes, R. E. 1992. Knowledge interchange format, Version 0.3, Reference Manual: Knowledge System Laboratory, Stanford University, CA, USA.

Gruber, T. R. 1992. *ontolingua: A Mechanism to Support Portable Ontologies*, technical report, Knowledge Systems Laboratory, Stanford University, Palo Alto, CA, USA.

Gruber, T. R. 1993a. *Toward Principles for the Design of Ontologies Used for Knowledge Sharing*, technical report KSL-93-4, Knowledge Systems Laboratory, Stanford University, Palo Alto, CA, USA.

Gruber, T. R. 1993b. A translation approach to portable ontology specifications, *Knowledge Acquisition*, No. 5, pp. 199–220.

Guarino, N. & Giaretta, P. 1995. Ontologies and knowledge bases, towards a terminological clarification, in N. J. I. Mars (ed.), *Towards Very Large Knowledge Bases*, IOS Press, Amsterdam.

Hage, J. 1996. Book review: Legal knowledge engineering; A modelling approach, by A. Valente (in Dutch), *Computerrecht*.

Hart, H. L. A. 1961. *The Concept of Law*, Clarendon Law Series, Oxford University Press, Oxford, England.

Heijst, G. Van & Schreiber, G. 1994. CUE: Ontology-based knowledge acquisition, in L. Steels, A. Th. Schreiber, and W. Van de Velde (eds.), *A Future for Knowledge Acquisition, Proceedings of the 8th European Knowledge Acquisition Workshop EKAW '94*, Vol. 867 of Lecture Notes in Artificial Intelligence, pp. 178–199, Springer-Verlag, Berlin/Heidelberg, Germany.

Heijst, G. Van. 1995. *The Role of Ontologies in Knowledge Engineering*, Doctoral Thesis, University of Amsterdam, Amsterdam, The Netherlands.

Kelsen, H. 1991. *General Theory of Norms*, Translation of "Allgemeine Theorie der Normen", Michael Hartney, Clarendon Press, Oxford, England.

Kralingen, R. W. van. 1995. *Frame-based Conceptual Models of Statute Law*, Computer/Law Series, No. 16, Kluwer Law International, The Hague, The Netherlands.

Kralingen, R. W. van. 1997. A conceptual frame-based ontology for the law, in P. R. S. Visser and R. G. F. Winkels (eds.), *Proceedings of the First International Workshop on Legal Ontologies*, (LEGONT '97), 4 July 1997, University of Melbourne, Law School, Melbourne, Australia, pp. 15–22.

Kuokka, D. R., McGuire, J., Weber, J. C., Tenenbaum, J. M., Gruber, T. R., & Olsen, G. R.: 1993. *SHADE: Knowledge-Based Technology for the Re-Engineering Problem*, Annual Report 1993.

Lenat, D. B. & Guha, R. V. 1990. *Building Large Knowledge-Based Systems; Representation and Inference in the Cyc Project*, Addison-Wesley, Reading, Massachusetts, United States.

MacGregor, R. 1990. The evolving technology of classification-based knowledge representation systems, in J. Sowa (ed.), *Principles of Semantic Networks: Explorations in the Representation of Knowledge*, Morgan Kaufmann.

McCarty, L. T. 1989. A language for legal discourse, I. Basic features, *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pp. 180–189, Vancouver, Canada.

McCarty, L. T. 1993. OWNERSHIP: A case study in the representation of legal concepts, Presented at a Conference in Celebration of the 25th Anniversary of the Istituto Documentazione Giuridica, Florence, Italy, 1993.

Moles, R. N. & Dayal, S. 1992. There is more to life than logic, *Journal of Information Science* (draft version), 3(2): 188–218.

Mommers, L, Schmidt, A. H. J., & Oskamp, E. W. 1997. Controversies in the ontology and law debate, in P. R. S. Visser and R. G. F. Winkels (eds.), *Proceedings of the First International Workshop on Legal Ontologies*, (LEGONT '97), 4 July 1997, University of Melbourne, Law School, Melbourne, Australia, pp. 1–5

Morik, K., Wrobel, S., Kietz, J.-U., & Emde, W. 1993. *Knowledge Acquisition and Machine Learning; Theory, Methods and Applications*, Knowledge-Based Systems, Academic Press Limited, London, United Kingdom.

Newell, A. 1982. The knowledge level, *Artificial Intelligence* 18: 87–127.

Schlobohm, D. A. & McCarty, L. T. 1989. EPS II: estate planning with prototypes, *Second International Conference on Artificial Intelligence and Law*, pp. 1–10, Vancouver, Canada.

Schreiber, G. 1992. *Pragmatics of the Knowledge Level*, Ph.D. Thesis, University of Amsterdam, The Netherlands.

Schreiber, G., Wielinga, B. J., Akkermans, J. M., & van de Velde, W. 1994. CML: The Common KADS conceptual modelling language, *Proceedings of the EKAW '94*, Hoegaarden, Belgium.

Sim, I. & Rennels, G. 1995. *Developing A Clinical Trial Ontology: Comments on Domain Modelling and Ontological Reuse*, Knowledge Systems Laboratory & Stanford Medical Informatics, KSL-95-60, June 1995, Stanford University, Palo Alto, CA, USA.

Stamper, R. K. 1980. LEGOL: Modelling legal rules by computer, in Bryan Niblett (ed.), *Computer Science and Law*, Cambridge University Press, Cambridge, United Kingdom.

Stamper, R. K. 1991. The role of semantics in legal expert systems and legal reasoning, *Ratio Juris* 4(2): 219–244.

Stamper, R. K. 1996. Signs, information, norms and systems, in B. Holmqvist and P. B. Andersen (eds.), *Signs of Work*, De Bruyter, Berlin, Germany.

Uschold, M. & Grunninger, M. 1996. Ontologies; principles, methods and applications, *Knowledge Engineering Review* 11(2).

Valente, A. 1995. *Legal Knowledge Engineering; A Modelling Approach*, University of Amsterdam, The Netherlands, IOS Press, Amsterdam, The Netherlands.

Visser, P. R. S. 1995. *Knowledge Specification for Multiple Legal Tasks; A Case Study of the Interaction Problem in the Legal Domain*, Computer/Law Series, No. 17, Kluwer Law International, The Hague, The Netherlands.

Visser, P. R. S. & Bench-Capon, T. J. M. 1996a. On the reusability of Ontologies in Knowledge System design, *Seventh International Workshop on Database and Expert Systems Applications* (DEXA'96), pp. 256–261, Zurich, Switzerland.

Visser, P. R. S., & Bench-Capon, T. J. M. 1996b. The formal specification of a legal ontology, in R. W. Van Kralingen (ed.), *Proceedings of the Ninth International Conference on Legal Knowledge-Based Systems* (JURIX '96), Tilburg, The Netherlands, pp. 15–24.

Visser, P. R. S. & Bench-Capon, T. J. M. 1997. A comparison of two legal ontologies, in P. R. S. Visser and R. G. F. Winkels (eds.), *Proceedings of the First International Workshop on Legal Ontologies*, (LEGONT '97), 4 July 1997, University of Melbourne, Law School, Melbourne, Australia (also available at: http://www.csc.liv.ac.uk/ pepijn/legont.html), pp. 37–45.

Visser, P. R. S., Jones, D. M., Bench-Capon, T. J. M., & Shave, M. J .R. 1997. An analysis of ontology mismatches; heterogeneity versus interoperability, *Working Notes of the AAAI Spring symposium on Ontological Engineering*, Stanford University, CA, USA, pp. 164–172.

Visser, P. R. S. & Winkels, R. G. F. 1997. *Proceedings of the First International Workshop on Legal Ontologies*, (LEGONT '97), 4 July 1997, University of Melbourne, Law School, Melbourne, Australia (also available at: http://www.csc.liv.ac.uk/ pepijn/legont.html).

Wiederhold, G. 1994). Interoperation, mediation, and ontologies, *Proceedings International Symposium on Fifth Generation Computer Systems*, (FGCS '94), Workshop on Heterogeneous Cooperative Knowledge-Bases, Vol. W3, pp. 33–48, ICOT, Tokyo, Japan.

Wright, G. H. Von 1963. *Practical Reason*, Philosophical Papers, Vol.1, Basil Blackwell, Oxford, England.