Ontological Engineering with Principled Core Ontologies

André Valente
University of Southern California
Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292, USA
valente@isi.edu

Joost Breuker
University of Amsterdam
Department of Computer Science and Law
Kloveniersburgwal 72
1012 CZ Amsterdam, The Netherlands
breuker@lri.jur.uva.nl

Abstract

An important issue in the newborn discipline of ontological engineering is the construction of libraries of ontologies which are designed for maximum reusability. Van Heijst et. al. suggested that a central part of ontology libraries is the definition of what they called a *core ontology*, containing elements that are as generic and method-independent as possible. However, their specification of how these core ontologies should be constructed is highly pragmatical, and leaves many problems unresolved. In this article we propose and discuss a number of specific principles for the construction of core ontologies. We demonstrate the advantages of these principles using as an example a core ontology we have built for the domain of law. Several conclusions about the construction of ontology libraries based on core ontologies are drawn.

Introduction

An important issue in the newborn discipline of ontological engineering is the construction of libraries of ontologies which are designed for maximum reusability (van Heijst, Schreiber, & Wielinga 1996; Gruber 1994). A major challenge in building these libraries is to define how the ontologies are to be constructed, and what the relations should be among them. Several groups have proposed solutions to this problem. Most of these proposals have primarily addressed the organization, or *indexing* problem. That is, they specify how the ontologies should be organized in the library in such a way that they can be meaningfully and efficiently retrieved for reuse. For example, (van Heijst, Schreiber, & Wielinga 1996) propose an organization based on the definition of a "core library". This core library is in fact a very general ontology of a certain application domain, e.g., medicine. However, besides some practical guidelines dependent on the consensus in a domain, they do not make it clear how such a pivotal ontology is to be conceived.

In our view, core ontologies should consist of a clear, theoretical framework for the selection of elements of the domain and principles for their definition. In this paper we propose that these *principled core ontologies* be constructed using four main principles. First, they should be *parsimonious*, i.e., they should contain enough concepts, but only those concepts which are strictly necessary. Second, they should have a *clear theoretical basis.* A core ontology should not be a simple hierarchy of terms, but a theoretical framework that describes what the domain is about. A core ontology is not simply the top part of abstraction hierarchies. Third, and related to the previous point, core ontologies should not aim at the specification of the most common terms, but of *basic categories* of domain knowledge. Fourth, these basic categories should be *coherent.* By coherent we mean more than that the basic categories should be consistent and complete, but also that the frameworks (relations) in which these categories are stored must make sense. This sense is the sense of the domain: e.g., in medicine the diagnosis and treatment of diseases.

In the body of this paper, we will define and discuss these principles in detail. In order to show the advantages of principled core ontologies, we will present in some detail one such ontology we have devised for the domain of law, called the *functional ontology of law* (Valente 1995). This article is structured as follows. In Section we discuss the idea of principled core ontologies. In Section we describe briefly the functional ontology of law. In Section we discuss and illustrate or ideas about core ontologies based on the functional ontology of law. In Section we present our conclusions.

Ontology Libraries and Core Ontologies

A major problem challenge in the construction of libraries of ontologies is the issue of ontology construction and the relations between ontologies. This is critical to obtain reusability. Thus far the focus has been on the indexing (organization) problem.

One strategy to cope with the indexing problem is to use the ontology terms (names) themselves as indexes.

This was roughly the solution adopted by the ARPA Knowledge Sharing Initiative (Neches *et al.* 1991; Patil *et al.* 1992) and implemented in the Ontolingua Repository (Gruber 1994). The problem, of course, is that there is no guarantee that the same term is being used with the same meaning in different ontologies. Most of the work in searching through the library is left to the user, and this strategy assumes that the user's understanding of the terms is similar to the one employed in the ontologies. In fact, the ontologies are not really organized as a library, but only "stored" in the repository. A second approach tries to correct this problem by defining a single meaning to each term, and again using the terms (and thus meanings) as the index. This is done by transforming the library into a single, large ontology, integrated by a very general, top-level ontology that is supposed to be coherent and complete (that is, to cover more or less all knowledge relevant in the given context). The basis proposed for building such generic ontology varies. One group of researchers proposes that it should be based on natural language — that is, the top level terms should be natural language categories or roles of terms (Knight & Luk 1994). A second group, of which the CYC builders are and example, prefers to see the whole enterprise as an encyclopedia, and suggests the careful construction of what amounts to an ontology of commonsense knowledge (Lenat & Guha 1990). While these two groups share the same basic indexing strategy, their results are quite different. CYC-like ontologies are microtheories, which can be used for reasoning purposes by logical inferencing engines. In contrast, the concepts in natural language ontologies are usually almost empty, and most of the meaning is given by their place in the subsumption hierarchies in which they are stored.

A main difficulty in the above approaches is that they want to be both general and independent of possible applications. (van Heijst, Schreiber, & Wielinga 1996) tried to give a different focus by restricting their attention to the use of ontologies in constructing knowledge-based systems. Thus far, their proposal for organizing libraries of ontologies is the most specific and explicit one. Their proposed library has multiple indexing and organization characteristics. First, ontologies are indexed by their level of abstraction and by how their terms were used in in some application. We will not discuss here the problems involved in the latter way of indexing: it does not affect the overall organization of the library. With regard to the levels of abstraction, the highest level of abstraction contains concepts like subsumption, inclusion, cause, etc. that seem more likely to be candidates for representational services than for terms in a knowledge base. In fact they may be presupposed rather than being made part of the library. Therefore, in practice the highest organizing level of the library is the "core ontology" (theory): a very general ontology of a certain application domain, e.g., medicine. This core ontology should contain a number of generic concepts and method-independent definitions, characteristics that would give high reusability to the elements of this library.

While we agree that such core ontologies are extremely useful for reuse purposes, there are several problems in the criteria developed by van Heijst et al. for defining which elements it should contain. First, they propose that these core ontologies should have a very pragmatic, engineering-like character. For instance, they propose that the core ontology should minimize the number of inclusions — which is just another way to define the engineering principle of *modularity* in this specific context. Further, the only clue they provide as to how to recognize the elements to be put in the core library is that they should be centered around "natural categories". By natural they mean that the categories should reflect a "social consensus that exists in the [application domain] community", and thus should allow communication between members of that community (van Heijst, Schreiber, & Wielinga 1996, pg.15). The flaw in this argument is that the fact that a term is used widely across a certain community does not mean that it is used with the same meaning. For instance, despite what one would expect, the meanings of terms like "law" and "right" are highly debated and disputed in the legal domain. Indeed, these terms are used in many different and sometimes contradictory meanings, to such extent that this fact is acknowledged by most legal theorists as unavoidable. Finally, (van Heijst, Schreiber, & Wielinga 1996) take 'theories' as the principle for modularization of the library. Theories are viewed as parts with a high internal cohesion and a relatively low level of coupling with other parts of a library of ontologies.

Principled Core Ontologies

Good engineering practice and pragmatism are of course important for building ontologies, but core ontologies require more than that. They should be based on a clear theoretical view of the elements of the domain, that can provide *principles* for their definition. In this paper we propose that these *principled core ontologies* are characterized by the following:

Parsimony A child is usually amazed to discover that solid and sturdy-looking cars with doors like vaults have such very thin metal skins, only a few millimeters thick. Compared to houses, cars look like a cheat. That is also the way many people outside (and even inside) AI think about the knowledge in and understanding of a knowledge system when they look under its hood. There is a general 'Eliza effect' even in model-based, deep reasoning systems (see, e.g., (Hofstadter 1995)). However, as in cars, this "thinness" is the consequence of good engineering. Being parsimonious with the material has the advantage of easy processing and total economy. However, when it comes to reuse, this economy is no longer of value. Parts of a Volvo cannot be reused to a new Renault. Similarly, parsimoniously constructed knowledge bases cannot be reused directly for other applications. Therefore the parsimony we mean here is of a different kind. An old Volvo may be reused for a new Renault by decomposing it to original materials. This relationship between construct and material is analogous to the relation between knowledge base and ontology. Parsimony in an ontology does not mean "thin", but rather "pure", i.e. that the *distinctions* (types in the definitions) are pivotal ones which have no spill or redundancy with other terms in the ontology. Neat taxonomies are a good example. However, many taxonomically defined terms in even very high level ontologies are mixtures of pure types and composite terms. For instance, "stuff" is not a "physical object", but a physical object may consist of stuff (cf. (Dölling 1995)). Medical pharmacology is not a kind of medicine, but a part of medicine (van Heijst, Schreiber, & Wielinga 1996). [1] This purity, or the essence of terms is the old problem in ontology, but therefore not less serious. As it may be out of reach, it is not necessarily without practical implications as degrees of pureness are in chemistry. To put it more directly: ontologies constructed from or in knowledge engineering in specific domains may still have contaminations. Distilling these in more pure, parsimonious categories may enhance clear material reuse, and is a necessary requirement for building and maintaining libraries of ontologies.

Clear theoretical basis A core ontology will embed in one way or another some basic view of what this domain is about, what its components are, and how they interrelate. In order to have principled core ontologies, it is necessary to make this theoretical basis as clear as possible. It provides part of the rationale behind the choices made in selecting the elements of the ontology. However, it is not always easy to select a theoretical background for building an ontology. Knowledge engineering is concerned with domains of practice. Different from domains of science, domains of practice do not have a fixed level of detail. Knowledge used in medical diagnosis and therapy (preventive or curative) — that is what medicine is about — may refer to macroscopic anatomical structures and its mechanics, but may involve as well the behavior of some ions in a biochemical soup. However, the various contributing disciplines have their specific level of aggregation, and therefore enable (cross-)indexing. However, here the point of view is not distinction, but composition and its abstractions (levels), i.e. a rather mereological view. Grouping practice domain terms back to their disciplinary levels will clarify and 'purify' a core ontology. However, not every domain of practice has such distinctive levels of aggregation. For instance, in law the aggregation implied in the domination hierarchies of laws (from constitution to contract) has no effect on the argument structure of legal practice: at every level the same legal ontological commitments are implied. In those cases, engineering a core ontology is far more complicated, because one cannot rely on the "natural" hierarchies of the domain as a theoretical principle for organizing knowledge

Categories instead of terms The choice of the terms to be put in the core ontology cannot be arbitrary or purely pragmatic. (van Heijst 1995) proposes that the common use of a term is a good basis for ontological definition. However, terms which are commonly used may have many contradictory meanings, and sometimes they are much more related to communication than to the essence of the knowledge that is used in the domain. Consequently, we propose that core ontologies should attempt to define *basic categories* of domain knowledge. [2] Categories are not top level terms in an abstraction hierarchy, but rather knowledge types. For instance, the *category* "disease" refers to knowledge about diseases, which may include a taxonomy of "diseases" (where disease is now used as a term). Categories have a meta-term flavor.

Coherence By coherent we mean more than that the basic categories should be consistent and complete for every level of detail. They also should be part

[1] (van Heijst, Schreiber, & Wielinga 1996) evade in fact the proper relation and talk about "hierarchy". However, this vagueness is typical of the 'impurity' meant, because we do not know whether medical pharmacology excludes 'internal medicine' (another branch of medicine), or may have things in common, etc.

[2] Categories are meant in an Aristotelian/Kantian sense: "Kant's categories are the ways in which the propositional structures extracted in logical theory function as the controlling principles of natural knowledge" (Bri 1973, Entry for "category").

of a framework that by itself makes sense. *They should describe what the domain is about.* For example, medicine is about the diagnosis and treatment of diseases. The domain of law is concerned with maintaining social norms whose (non-)observation can be attributed to the responsibility of individual humans. These are highly functional views, which fit these domains of practice very well. Domains of practice are the domains for which applications are built.

There are several advantages to using principled core ontologies. First, they provide a common basis for *comparison* and *translation*. There is no such thing as the "right" or "best" ontology in general. The quality of an ontology is usually measured by its usefulness, or its reusability. However, since there is no way to define the best ontology, it is essential that we be able to compare them and translate ideas formulated in one ontology to another. We will show that principled core ontologies improve the chances of being able to compare and translate (core) ontologies. Second, a principled core ontology divides the domain into highly modular parts, greatly simplifying the problems of analyzing the specific reasoning mechanisms used in the domain and thus helping to design or classify problem-solving methods which can be used for that domain.

A Principled Core Ontology for Law

In order to illustrate how we think a principled core ontology would look like, and what its benefits can be, in this section we will discuss an ontology we developed for the domain of law, called the *functional ontology of law* (Valente 1995). This ontology is based on a relatively complex theory of law, with foundations in legal theory, and centered on the idea of defining functional roles of legal knowledge. It is outside the scope of this paper to describe the ontology in detail. The brief description below is intended to explain the main aspects of the ontology. In section we use the functional ontology of law to illustrate principled core ontologies.

Basic View

The functional ontology of law, adopts a *functional* perspective that can be summarized as follows. First, it is assumed that the legal system as a whole exists to accomplish a certain *function*, in order to obtain certain *social goals*. 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 behavior.[3] Second, 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 *function* or *role* distinguished by the legal system in the operationalization of its functions and tasks. These functions 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.

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.

Normative Knowledge Normative knowledge consists of *primary norms* (or only *norms*. A norm expresses an idealization: what ought to be the case (or to happen), according to the will of the agent that created the norm. This idealization is expressed by reference to a description of the reality (the world) in which some configurations of facts and behavior 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 behavior 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 behavior with respect to the norm. Primary norms are entities that refer to human behavior, 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 behavior, in the sense that it can provide a status only when it is applied to certain types of cases. For the remaining types of cases,

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

the norm is said to be *silent*.

Meta-legal Knowledge Normative systems are defined on the basis of individual norms in the sense that the standard defined by a normative system is defined in terms of the standards defined by the individual norms. There may be a difference between the normative statuses given by a single norm and by the normative system. 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 statuses 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 *valid*.

World Knowledge (Breuker 1990) called *world knowledge* the knowledge in a legal system that describes the world being regulated and the possible behaviors in that world. In addition to adopt the distinction of a category of legal knowledge which describes the world, we propose that this knowledge constitutes a structured *model* that we call the *legal abstract model* or LAM. 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 LAM consists of definitions of *concepts* that represent entities and relations in the world. Primitive concepts are supposed to be interpreted by people, and thus assumed to be commonsense. Thus, the legal abstract model is in fact a layer of definitions of concepts and relations built on top of a large layer of commonsense 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 commonsense knowledge.

The legal abstract model is not only a static description of salient features of society, but a model of social behavior. A LAM should ideally allow predictions about behavior in the world. We propose that this description of possible and relevant behaviors is built around the concept of *cause*, in order to allow the assignment of responsibility of an agent for a certain case. Causal knowledge is used by the responsibility knowledge to describe who or what have caused a given state of affairs, and can thus be considered responsible for it.

Responsibility Knowledge Cause and responsibility are important concepts in 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 behavior, 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.

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 if compared to the other ones. In this case, the law not only wants to classify or to react over 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.

The Big Picture

In the previous sections, we proposed a number of categories of legal knowledge. In order to show that they are coherent — in the sense that they form a coherent whole — we need to show that together they perform the function the whole legal system performs. This is pictured in Fig. 1. The rounded boxes represent functions (or, alternatively, bodies of knowledge which perform the function), and the solid arrows indicate functional dependencies (inputs and outputs). 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 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 analyzed 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 'disallowed'). In another path, the situation is analyzed 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.

Principled Core Ontologies

We have presented in the previous section an ontology that follows the principles we propose in this paper:

1. It is parsimonious in that, despite trying to cover the whole of legal knowledge it departs from only six primitive concepts, which have at most a few subconcepts 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 instance, there is a large amount of legal sources in the US legal system pertaining 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.[4]

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

Several advantages came from using the functional ontology of law using these principles. They are loosely related to the fact that principled core ontologies can be excellent tools to implement a divide-and-conquer strategy in knowledge acquisition, and that is how we used the functional ontology of law. 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 commonsense 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 domain-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 dif-

[4]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.

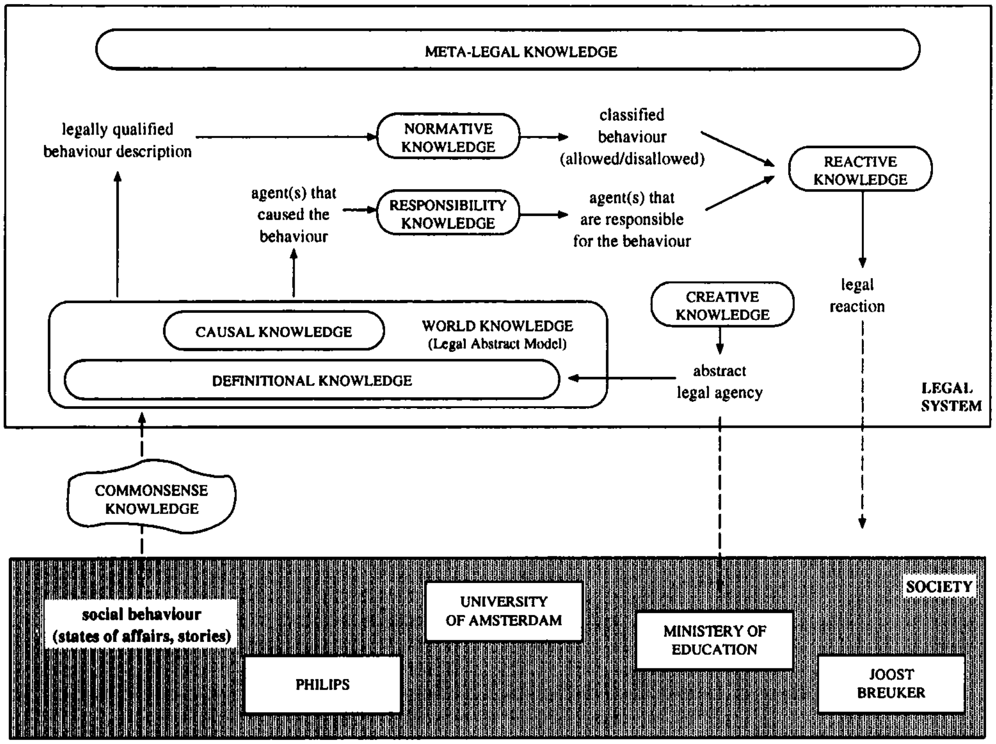

Figure 1: Functional roles of legal knowledge in the operation of the legal system.

fer 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. 1 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 Fig. 1; 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). 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.[5] 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.

Discussion: what is to be reused?

In most works on ontology and reuse it is not clear what is reused more than the literal terms. An ontology consists of concepts (terms), their definitions and some relations with other terms. These relations express in the first place similarities and distinctions between terms via subsumption or inclusion hierarchies. Also causal, temporal or topological relations are used in ontologies. In general, there is no structural correspondence between the relationships in ontology and a knowledge base. [6] Correspondence is only by identity of names of terms, and (parts of) definitions. As the terms in an ontology are defined as concepts, (parts of) the definition may be *translated* directly into knowledge base propositions. It should be noted that the definitions in a useful ontology are in general richer than what is to be used in the knowledge base. The meaning of the concepts get a more specific *sense* when it comes to a particular use in the knowledge base. The famous "interaction effect" demands a specific, thin interpretation of a term. This is what is meant by ontological commitment: the sense/meaning of a term that is assumed, and not explicit in a knowledge base. For example, in the functional ontology of law we could have specified the meaning of cause in more detail, by adding one of the theories of causation proposed in AI. However, there is no guarantee that the assumptions made in one such theory are valid in a specific legal system. Indeed, we found that to some extent each legal system constructs its own theory of causation based on commonsense views held in a particular society, and that frequently have to do with ethics and religion. In other words, while the *role* of causation is the same across legal systems, and we were able to express that in our ontology, it would be a mistake to include a specific interpretation of the term, since it would probably not reusable.

Another point worth noting is that an ontology may not contain *all knowledge that is in the knowledge base.* Some empirical or compiled-out practical knowledge that is dependent on the particular method may not find an easy, and principled place in an ontology. This is not bug in an ontology, as (Motta *et al.* 1996, p 361) believe, but a feature. An (application) ontology is not a full application documentation, but a conceptualization and consensus of terms. For that reason, the Sisyphus-II reconstructions of the VT-application (not: domain) could not be solely based upon whichever of the ontologies available.

The last point we want to make about the relation between the reuse that can be made of an ontology and the content of a knowledge base concerns the status of *roles.* Knowledge plays roles in problem solving. as is invariably emphasized in knowledge acquisition approaches (e.g. (Clancey 1985; McDermott 1988; Schreiber 1992)). For instance, in a medical domain, a disease may in general play the role of a hypothesis (and at the end, the role of a solution), but it may also play the role of evidence when it is used to explain the likelihood of the presence of some other hypothesized disease However, as roles are dynamically attributed to knowledge in reasoning, they themselves cannot be part of a domain ontology. For example, in the ontology of the VT domain as used in the Sisyphus-II project (see (Schreiber & Birmingham 1996)) 'fixes' is one of the main concepts, and as a consequence it requires a lot of reinterpretation to find out what the meaning is of the rather heterogeneous collection of terms subsumed by this term. Fixes, or hypotheses or evidence are not specific for some domain. This point can be further illustrated by reference to an example of building an ontology for reuse purposes, based on an existing knowledge base of an application. In a recent work by one of the authors, an ontology was constructed from the knowledge base of INSPECT, a

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

[6]Direct, automatic translations between an ontology and a knowledge base may maintain structural similarity, e.g., class concepts in Ontolingua and the T-Box of LOOM (Gruber 1994).

system for critiquing air campaign plans (Valente, Gil, & Swartout 1996). The idea was to make this ontology public within a group of researchers that are building other applications in the air campaign planning domain. In the process of building the ontology, certain concepts that were used in INSPECT had to be removed, because they did not make sense outside the application context. For example, a concept like `objective-with-too-many-parents` (the objectives are arranged in a hierarchy) only makes sense within INSPECT, because the definition of what "too many" means is dependent on the context of critiquing the formulation of the objectives. This concept would either make no sense or have to be totally redefined if it was to be used in e.g., an application that generates air campaign plans.

There is not (yet) a hard criterion to decide when roles are the wrong indexes in a core ontology. The principle is that one should not smuggle in interactions between domain knowledge and general problem solving knowledge. Roles by themselves are not a wrong ontological category. The functional perspective as we propose for the core ontology is also based on roles. However, these roles of e.g. normative and responsibility knowledge are not dynamically attributed to legal knowledge, but make up the legal system itself.

Conclusions

The main theme of this paper is the construction of (more) principled core ontologies to support libraries of ontologies. We largely agree with (van Heijst, Schreiber, & Wielinga 1996) that the organization of such a library should be dependent on its relations with this core library and the major goal of this paper is to take one step further and propose a more principled instead of (or additional to) the rather pragmatic approach of van Heijst *et al.* Principles sound nice. However, principles in engineering are not simply given, but grow from experiences and reflection. As there are no real experiences with libraries of ontologies for knowledge acquisition, we still have to rely on reflection on the results presented so far. Are our results for the law domain different, and 'better' than the results of (van Heijst, Schreiber, & Wielinga 1996) (see also: (van Heijst 1995)) for the medical domain.

The core ontology for medicine proposed by van Heijst *et al.* (cf. Fig. 5.2 and Table 5.1 in (van Heijst 1995) contains the following categories in the core library: 0) generic concepts 1) fundamental-medical -concepts (e.g., human-body), 2) anatomy, physiology, 3) findings, drugs, surgeries 4) clinical state abstractions and 5) diseases. The terms are related by in-/exclusions. A closer look reveals three points of view mixed: theories (anatomy, physiology) at the top and clinical environment at the bottom. Between these there is a core consisting of a diagnosis branch (tests, diseases, findings) and a therapy branch (therapy, drugs, surgeries). This core is a highly functional one and we would call this a functional ontology of the medical system. Moreover, the categories look also like lines of argument: clinical state abstraction connect diseases with findings; therapies connect diseases with drugs and surgery. Therefore there are similarities and differences between the two results.

- The medical core contains heterogeneous categories, that should be heterogeneously related. For instance, one may read that surgeries include physiology. There is no way to get the picture right with uniform relations. According to our principles, however, the theories provide a separate perspective and should yield additional indexing. The medical system and the physiological system should not be put into the same pressure cooker. In the domain of car manufacturing, the concept of car may have many views (artifact, commodity, means of transport) but the predominant (functional) one is 'product'.
- What both core ontologies have in common is that they are functional ones. They mark the major roles of practice in the domains.

Indeed, our core ontology looks 'better', if only for the fact that it is less heterogeneous, so that one may expect that it provides a more uniform structure for and access to a large library of ontologies. However, we may have pushed the real heterogeneity of a domain to the background by suggesting that e.g. the scientific theories or the hierarchies of law are subsidiary indexes. Indeed, for a particular sub-domain of law, e.g. copyright law, the regulations have to be broken up into the various legal categories as normative, definitional etc. ones. However, this is not a problem but a feature, because it also *guides* the knowledge acquisition. Still it is true that we need cross indexing to keep all definitions of copyright law together and separate from law about house renting. Moreover, we also have to take care that the concept of property is applicable (identical) to both because they are part of the same higher civil code. Whether in a medical core library secondary indexing by the theoretical perspectives can also be accomplished 'as simple' is hard to say. It would suggest that separate discipline (libraries of) ontologies have to be constructed and maintained as well.

The comparison suggests also another hypothesis: that in domains of practice a functional perspective is the major provider of coherence. We have far too few data to refute this hypothesis. It looks plausible, because the very notion of practice suggests functionality. However, the point is rather that this functionality is also the major argument structure. We do not think this is pure coincidence, but that speculation on this issue requires a separate more reflective study.

Acknowledgments

Andre Valente gratefully acknowledges the support of DARPA with the contract DABT63-95-C-0059 as part of the DARPA/Rome Laboratory Planning Initiative.

References

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

1973. *Encyclopedia Brittanica*, volume 5.

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

Dölling, J. 1995. Ontological domain, semantic sorts and systematic ambiguity. *International Journal of Human-Computer Studies* 43:785–808.

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

Hofstadter, D. 1995. *Fluid Concepts and Creative Analogies*. Basic Books.

Knight, K., and Luk, L. 1994. Building a large-scale knowledge base for machine translation. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*.

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

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

McDermott, J. 1988. Preliminary steps towards a taxonomy of problem-solving methods. In Marcus, S., ed., *Automating Knowledge Acquisition for Expert Systems*. Boston, MA: Kluwer Academic Publishers. 225–255.

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

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

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

Schreiber, A. T., and Birmingham, W. P. 1996. Editorial: the sisyphus-vt initiative. *International Journal of Human-Computer Studies* 44(3/4). Special Issue: the Sisyphus-VT initiative.

Schreiber, A. 1992. The KADS approach to knowledge engineering. *Knowledge Acquisition* 4(1):1–4.

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

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

Valente, A.; Gil, Y.; and Swartout, W. 1996. Inspect: an intelligent system for air campaign plan evaluation based on expect. Technical report, USC Information Sciences Institute.

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

van Heijst, G.; Schreiber, A.; and Wielinga, B. 1996. Using explicit ontologies in kbs development. *International Journal of Human-Computer Studies*. forthcoming.

van Heijst, G. 1995. *The Role of Ontologies in Knowledge Engineering*. Ph.D. Dissertation, University of Amsterdam.