

Ontological Modelling for Designing Educational Systems

Joost Breuker & Antoinette Muntjewerff
Dept. of Computer Science & Law
University of Amsterdam
PO Box 1030, 1000 BA Amsterdam
Netherlands
e-mail: {breuker,muntjewerff} @lri.jur.uva.nl

Bert Bredeweg
Social Science Informatics
University of Amsterdam
Roeterstraat 15, 1018 WB Amsterdam
Netherlands
e-mail: bert@swi.psy.uva.nl

June 17, 1999

Abstract

Constructing ontologies in educational design is not really new. The specification of educational goals is what is called now a days an ontology. Although content has always been considered a crucial factor in education, the emphasis in educational research has been on form, as is also pointed out by [Mizoguchi *et al.*, 1997]. Ontological engineering for constructing educational systems may look like putting the same old wine in new barrels, but we should be aware that that these new barrels may give a new flavour to this wine. As an example we discuss a core ontology about law, used in the development of educational systems.¹ A core ontology mediates a top ontology, that reflects our common sense understanding of the world, and an ontology that defines the concepts and structures in a domain. A core ontology tells us what a domain is about. The core ontology discussed is FOLaw [Valente, 1995], a functional ontology of law, as applied in PROSA, a system that trains students to solve problems (cases) in administrative law. A major conclusion is that emphasis has shifted from skill acquisition to obtaining insight and understanding. Another benefit of this ‘ontological view’ is that types of knowledge distinguished in core ontologies can make up categories that provide the similar decompositions as task analyses, but apparently in a more ‘natural’ way.

1 Introduction

Constructing ontologies in educational design is not really new. In Bloom’s taxonomy of educational goals, the category ‘content’ has the role to specify the concepts to be taught in a course [Bloom, 1957]. However, this does not mean that in educational research the specification of content – and the consequences that follow from it for the design – have been an important issue.

¹It was intended to discuss three ontologies used in educational settings. Time and space has constrained us to leave out the presentation of ontologies of ecology and of information systems.

Educational research has been concerned rather with *methods*, or ‘form’ [Mizoguchi *et al.*, 1997] than with content and its structure. Methods of instruction, training and evaluation have been the major topic of research, and whenever domain knowledge has been in focus, it was rather the method of reasoning or problem solving rather than the structure of the domain knowledge and its assumptions that has been object of research.

Educational design is not the only discipline (that should be) concerned with the description of content, *i.e.* knowledge and information. Cognitive science in general, but also more applied disciplines like information sciences, knowledge engineering and knowledge management meet one another at the crossroads of a new, emerging discipline called ontological engineering. After describing what is in ontological engineering, we will discuss the roles that the specification of various kinds of ontologies may play in designing educational systems. We will exemplify these roles by drawing from experiences in which explicit ontologies have been (re)used in three domains: law, ecology and information processing applications (intelligent help systems).

2 Ontology and ontological engineering

In the new-born discipline of ontological engineering the term “ontology” has another meaning than as originally conceived in philosophy. Wolff (1679 – 1754) elaborated on Leibniz’s systematization and introduced the term ‘ontology’ to distinguish the essential from the incidental, the real thing from its appearance.² Ontological engineering is hardly interested in the question of existence (or its representation, see *e.g.* [Hirst, 1991] for an exception), but it is concerned with the specification of content and structure of knowledge.

The ‘engineering’ denotes the fact that there is a focus on the construction of methods and tools, in particular computational ones for all kinds of application areas. Therefore, many (applied) disciplines are involved in and supported by ontological engineering: information (library) science, (computational) linguistics, corporate knowledge management, data base technology (schemas), professional terminological standardization (*e.g.* medicine), and last but not least: knowledge engineering. In this paper we will focus on the last, as the construction of (intelligent) educational systems can be viewed as a branche of knowledge engineering. Each of these application areas bring different views and requirements to the ontology arena. Linguistic ontologies like WordNet [Miller, 1990] or SENSUS [Knight *et al.*, 1995] contain more than 50.000 concept definitions that can be used *e.g.* for translation. At the other extreme are the specialised domain ontologies (re)used in knowledge engineering as *e.g.* deposited at the Ontolingua server [Gruber *et al.*, 1996], Webonto [Domingue & Motta, 1999]. For knowledge engineering the definitions of the terms in an ontology are much more detailed than for linguistic or document indexing applications. However, in all applications there is the assumption that more abstract terms should constrain the interpretation of these definitions. Terms which re-occur in every domain and are part of our common sense understanding of the world, like time, space, cause etc., are part of a *top* ontology *e.g.* [Hobbs, 1995], [Sowa, 1995], [Guarino, 1995, Guarino, 1998]. The ontologies of law, ecology and information processin systems that we present here are neither top ontologies, nor domain ontologies, but they are intermediates between the completely general top ontologies and ontologies of legal, ecological, etc. domains. After [vanHeijst *et al.*, 1997] who developed a similar kind of ontology for medical domains we call such an ontology a *core* ontology. Before

²[Kant, 1787, B.874] who was interested in Leibniz work, learned about Leibniz via Wolff (“the Leibniz/Wolff system”), and classified ontology as one of the four components of metaphysics.

presenting our core ontologies and their application in designing educational systems, we will first discuss the notion of ontology in general.

2.1 A Notion of Ontology

It is impossible to represent the world in its full richness of detail. In order to represent a certain phenomenon or a part of the world (which is called a *domain*), it is necessary to restrict the attention to a small number of concepts which are meaningful and sufficient to interpret the world and provide a representation adequate to a certain task or goal at hand. As a consequence, a central part of knowledge representation consists of elaborating a *conceptualisation*: a set of abstract objects, concepts and other entities which are assumed to exist in a certain domain, as well as the relations that may hold between them [Genesereth & Nilsson, 1987]. The commitments which are implied by the choice of one set of concepts instead of another to describe a certain phenomenon are called *ontological commitments*. A conceptualisation thus also carries a set of ontological commitments. These commitments can be made explicit by ‘importing’ or including the more abstract ontologies that reflect the assumptions or point-of-views that have been made (taken) in conceptualizing a domain. A core ontology, therefore, reflects the views on the world taken by a particular field or discipline.³

In this way, top ontologies also play a crucial role in the construction or selection of a representation formalism to express knowledge. Elaborating conceptualisations, and thus selecting ontological commitments, is an essential component of the task of representing knowledge, because conceptualisations select which things are relevant to be represented and which are not [Davis *et al.*, 1993]. Representing knowledge involves both the design of a *knowledge representation language* and the formulation of a specific set of sentences in this language which describe certain things in the world — such set of sentences is usually called a *knowledge base*. Ontological commitments precede the elaboration of both the knowledge representation language and the knowledge base: the distinction being that the representation language offers a minimal set where the consequences of the commitments emerge as organization and inference services. A good example is Allen’s ontology of time [Allen, 1984].

A conceptualisation is an abstract entity which is only implied in a knowledge representation. An *ontology* is a specification of a conceptualisation [Gruber, 1994]. It comprises a description (e.g. through *definitions*) of the concepts, objects, relations and so forth which make up a conceptualisation. One of the basic roles of an ontology is to enable the study of conceptualisations and ontological commitments in their own right, i.e. dissociated from the knowledge representations they may yield. Another important role is to support knowledge sharing and reusability [Neches *et al.*, 1991, Patil *et al.*, 1992, Gruber, 1994].

2.2 Roles of Ontologies in KE

Three kinds or levels of ontology can be distinguished in knowledge engineering: top, core and domain.⁴ Each of these levels has its specific roles or use. The different purposes of (re)use of

³One can also see this as a tautology, where a set of domains that has a set of common ontological commitments, forms a discipline.

⁴[vanHeijst *et al.*, 1997, p 193] proposes four levels: application, domain, generic, representation. We believe that there is no real distinction between application and domain ontology. The distinction is made because [vanHeijst *et al.*, 1997] believe that an application ontology fully matches a knowledge base and a domain ontology may only have a partial match with the knowledge base of some domain application. We do not hold such a view and

ontologies put different requirements on ontologies. The qualitative demands increase with the order presented below, *i.e.* the need for careful and valid analysis and in particular for a *formal* base becomes greater (see [Guarino, 1995, Guarino, 1998] on the notion of formal ontology).

Domain ontologies contain the concepts of some domain of application:

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

Core ontologies contain the categories that define *what a field is about*. A field is some discipline, industry or area of practice that unifies many application domains, as *e.g.* (some subfield of) medicine, law, engineering, etc. The categories are not some common denominator of a set of application domains, but constrain what is *relevant* in these domains. As hypothesized by [Valente *et al.*, 1999] (see also [Valente, 1995]), core ontologies have (1) a functional character and (2) reflect the major structure of reasoning or argument in a field.

The functional perspective may be understandable by the fact that fields are typically fields of practice, often concerned with artifacts. As a consequence, types of knowledge can be distinguished by their roles. That these roles may also reflect the predominant structure

argue that real application ontologies are impossible to construct (see also [Motta *et al.*, 1996]). The generic ontology is the same as a top ontology. However, the term 'core' is borrowed from [vanHeijst *et al.*, 1997] where they define the categories that index a library of medical domain ontologies: these categories are generic for medicine. We do not distinguish between generic and representation ontologies because they have the same nature. Whether a generic (= top) ontology is the basis for developing a KR formalism and related inference services or not is in the use of an ontology, in the same way as a domain ontology may be used to construct a knowledge base.

of reasoning is more speculative, but may be conceived as that domain knowledge is a “a model of a system in the world”, and that reasoning means some operation on this simulated system, or the construction of such a system [Clancey, 1992]. For instance, in medicine the major categories that make up the core ontology are diseases, tests and therapies, and the various disciplines that make up our understanding of (disturbed) biological processes play the role of supportive explanation [vanHeijst *et al.*, 1997]. Similarly, in electro-mechanical engineering, three major types of knowledge can be distinguished: structural (a configuration of components), from which, given the behavioural or functional description of the components, a behavioural account can be derived, which can be the input for quantification (mathematical modelling) [Borst & Akkermans, 1997]. Also the core ontology of law, developed by [Valente, 1995], reflects the major structure of argument in this field, as will be explained in Section 3.1.

As intermediates between top and domain ontologies, core ontologies have a double-faced role:

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

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

2.3 Ontologies in designing educational systems

The overview of nature, types and roles of ontologies suggests, that in designing and building educational systems, core ontologies may play a prominent role. Top ontologies reflect basic common sense, and are therefore not typical candidates for playing roles in education. However, as top ontologies describe our common sense views on the world, they contain the ingredients of our “naive” physics (*e.g.* causality), biology (*e.g.* autonomy, agent), psychology (*e.g.* intention, belief) and sociology (*e.g.* communication). These naive conceptions may interact with the acquisition of the more scientific versions and elaborations taught at school. Also in the use of metaphors in teaching new subject matter, top ontologies may make us more precisely aware where metaphors may facilitate and where they may break down. For instance, the hydraulic metaphor for electricity does not explain the behaviour of coils (induction). Domain ontologies may have an important role in reuse, like in knowledge engineering in general. The content of education may be largely constant in a culture, but the educational context in a context may easily vary in requirements and

preferences about didactic methods. However, as we will attempt to show in the rest of this paper, we expect a more important role for core ontologies than for top- and domain ontologies in the design of educational systems.

Core ontologies may have specific roles for educational design as will be illustrated in the next sections.

- By distinguishing the major categories of knowledge core ontologies allow a “divide & conquer” approach in designing educational systems for various reasons:
 - They may provide a top-down knowledge acquisition framework. For instance, in a medical domain (*e.g.* internal medicine, bacterial infections, etc.) the description of diseases and of therapies provide the major division. Even if the particular domain is only concerned with one of these two categories, *e.g.* in surgery, the other category provides the required context in which the first operates.
 - The typical “divide & conquer” principle applied in knowledge engineering and educational design has been task decomposition. However, task decompositions have a strongly arbitrary and idiomatic flavour. Even if work on problem solving methods (PSMs) has created some unified and justified view on task decompositions in knowledge engineering, the mapping to domain knowledge is still in a fuzzy state, while in practical applications domain specific terminology complicates the mapping to PSMs. As hypothesized earlier, a core ontology provides a functional articulation of the domain knowledge, and this articulation may very well reflect the major reasoning structure in a field. It therefore may turn out to be more natural and simpler to start from domain knowledge distinctions to arrive at a task decomposition than the other way around. So, instead of constructing libraries of PSMs, indexed by task- or problem types ⁵ one may construct libraries consisting of core ontologies that bottom out in PSMs that may turn the initial decompositions in well controlled task structures.
 - This approach may not only work for knowledge acquisition in specifying intelligent (educational) systems, but it also reflects the way knowledge and skills are acquired in educational processes. It is well known that instructing methods, even systematic and rational ones, has little effect on acquisition rate. Students hardly ever complain that they do not know *how* to proceed but to find their way in the knowledge to apply, *i.e.* the problem is to come to grips with content rather than method. In skill acquisition, articulate task structures appear to be an emergent property of practice in problem solving, rather than drivers of effective practice and learning.
- Core ontologies provide the initial structure of domain knowledge. In this way they can be used to define the interface between a ‘knowledge level shell’ and its to-be-acquired knowledge base, as we will explain in detail in the EUROHELP example below.

Thus far we have assumed that ontologies refer to domain knowledge, *i.e.* the subject matter in educational systems. However, an ontology is in the first place a terminology, so we can also make ontologies about education and educational systems themselves. This is the idea behind [Mizoguchi *et al.*, 1997]’s “task ontology” for intelligent educational systems (IES). This ontology, called CLEPE, forms the primary access to authoring tools for IES. Although described as

⁵It can be shown that the indexing of these PSMs by problem types is not self-evident, if appropriate at all [Breuker, 1999]

a ‘task ontology’ the major categories refer to the complete educational (design) process: educational goals, learner state, system functionality and teaching material knowledge; the latter refers to subject matter, *i.e.* real ‘content’. In fact, CLEPE can also be read as an instructional design theory in a more axiomatic form, similar to *e.g.* component display theory [Merrill, 1987b], which is also turned in a an authoring (advice) system [Merrill, 1987a]. CLEPE hides a theory about educational design, and reflects not necessarily the shared understanding of the educational research community, even if many of the concepts specified are undisputed. However, the advantage of CLEPE is that it is more axiomatized and operationalized than most educational modelling frameworks available, and in this respect CLEPE may trigger a more focussed discussion on basic distinctions in educational engineering than by publication of research papers, etc. Field ontologies may neither be undisputed (as our example for law may reveal), but they often describe well established consensus.⁶

3 Designing educational systems for legal domains

In this section we will first explain a core ontology developed for law. This ontology, FOLaw, was originally developed by [Valente, 1995] (see also [Valente *et al.*, 1999]). It is a good illustration of what we mean by core ontology. After this explanation, we present the design of PROSA, a system that trains students in administrative law [Muntjewerff & Groothuismink, 1999].

3.1 A functional ontology of law: FOLaw

The ontological commitments (view) that have lead to the core ontology are the following:

The legal system as a system The legal system is viewed as an entity with a certain internal structure, behaving in an environment.

The legal system as an artifact As the legal system is an artifact it has presupposed functions. Although the legal system is not the result of some ‘grand’ design, but emerged as the result of many successive local design decisions, it definitely serves a purpose: control over social behavior. Therefore, the output of the legal system is an *intended* output.

The social functions of the legal system The main function of the legal system is to prescribe and to react to social behavior. In this sense, the legal system can be regarded as a kind of *social device*. The legal system is a subsystem of the political-power system. It is semi-independent in the sense that it accepts goals and constraints from this governing system, but further works according to its own ‘rules’. Note that this is not some sociological theory; it rather reflects minimal commitments as to comply with *common* sense views on the roles of the legal system.

Functions of law are supported by legal sources Social systems are in general described by the agents (people ,institutions) that interact with one another in semi-fixed patterns. However, this “agent” perspective is probably not the right ontology to describe the legal system. Although many fixed roles of agents in the legal system can be easily described, such as judges, prosecutors and lawyers, and may include such mundane types as defendants, parties and

⁶Domains that are explicit object of semi-automatic support of learning are almost always part of a formal curriculum and are therefore ‘well-established’.

civilians, their roles are only of secondary importance in relation to the rules that constitute the system itself: the law. The law is defined by its *legal sources*, such as legislation and precedent law. Legal sources contain the (codified) knowledge which specifies how the legal system works or should work: not only internally, but in particular in reacting to social behaviour in a society.

These commitments may not look very surprising, but the outcome is somewhat surprising as it cuts across traditional views in law and in sociology on the legal system. Law is traditionally categorized by the social subsystems a particular legal source refers to, such as crime (poenal law), trade, family, property, etc. However, such a categorization is not a categorization of the legal functions but of social systems (as defined by the politico-legal system). Neither do we adopt a sociological view in which the legal institutions and agents are the first class ontological citizens. Agents and institutions have secondary roles.

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 primary functions of legal sources and corresponding categories are proposed and described: *normative knowledge*, *world knowledge*, *responsibility knowledge*, *reactive knowledge*, *creative knowledge* and *meta-legal knowledge*.

3.1.1 Normative Knowledge

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

The basic conception of norm used in the ontology is largely derived from the legal philosopher [Kelsen, 1991]. A norm expresses an idealisation: what ought to be the case (or to happen). This idealisation is expressed by reference to a description of the reality (the world) in which some configurations of facts and behaviour are 'cut out' (prohibited) to make it an ideal world. Since they express an ideal world, norms can be either *observed* or *violated*. A norm is observed when the behaviour in the real world does not conflict with its specification in the ideal world, and violated otherwise. To *apply* a norm means to verify or compare the reality with the ideal world defined in the norm, classifying the reality as either compliant or non-compliant with the norm. This classification is the *normative status* of the behaviour with respect to the norm.

3.1.2 Meta-legal Knowledge

Legal sources are made up of individual norms. These norms may give conflicting normative statuses for the same situation. This is already evident by the fact that regulations invariably contain exceptions. Meta-legal knowledge provides the rules by which these conflicts are resolved: in general the conflicts inherent in exceptions are resolved by *lex specialis*, i.e. the more specific norm should be applied. Another function of meta-legal knowledge is to specify which legal knowledge is *valid*. Validity is a concept which can be used both for specifying the dynamics of the legal system and its limits. A valid norm is the one which belongs to the legal system, and vice-versa.

3.1.3 World Knowledge

By its very nature, law deals with behaviour in the world. Therefore, it must contain some description of this behaviour. For instance, in order to describe how the world should (ought to) be, primary norms must describe how things are or can be. This description is not directly available from the legislation, but is usually implicit. However, this type of knowledge is distinct from (albeit connected to) normative knowledge — that is, primary norms describe an ideal world *based* on the description of reality.

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*. Thus, the term *legal abstract model* or LAM is used as a synonym for world knowledge when its model character is to be stressed. Also, this term stresses that this model is a ‘double abstraction’: as discussed below, the LAM gives a more restricted meaning to commonsense concepts which are already abstract.

The LAM 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. It expresses the legislators view on some domain, *e.g.* crime, traffic, etc. This view is necessarily abstract and constrained to legal functions. It is here where common sense accounts of cases meet a legal interpretation and a selection of “legally relevant facts”. Besides an identification of relevant facts, the law also needs to establish *causal* accounts between the these facts in order to establish which agents can or cannot be held responsible for violations of norms. The initial attribution of causal accounts of events is in the LAM left to common sense reasoning. However, the law itself may have additional or more specific views on responsibility (see below).

3.1.4 Responsibility Knowledge

Responsibility is more than simple, physical causation. Two more ontological ingredients are required: intention and belief. One may cause *e.g.* physical harm, but may not have intended so, or may have good reasons to belief that some action should not harm. In this sense, responsibility reasoning is still in the realm of (deep) common sense. But this causal connection is not always necessary or always sufficient for establishing responsibility in a legal context. The role of legal responsibility knowledge is exactly to ‘interfere’ with this *prima facie* connection between causing and being responsible. This interference is made so that legal systems “extend responsibility [or] cut it off in ways which diverge from the simpler principles of moral blame” [Hart & Honore, 1985, pag. 67]. This mechanism has rather practical motives. Given the innumerable problems in establishing, proving and reasoning with causal connections, and the attribution of intention and belief to agents, the assignment of legal responsibilities may bypass these connections, and give reasons. For instance, one of the reasons may be that a priori risks in actions should put the burden of proof to the agent who took the risk. For instance, in many legal systems a car driver has to proof that she has no responsibility when she was involved in an accident with a pedestrian. In other words, the pedestrian is not required to present a causal account as long as the car driver cannot be absolved.

3.1.5 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. Rewards are *e.g.* (financial) benefits, or rights, and they are contingent upon *no* violations of norms.

3.1.6 Creative Knowledge

Law is created by some bootstrapping procedure in which power becomes distributed over agents and institutions in the form of authority and rights. This distribution is itself object of law (constitutional and administrative law) in which the creation of *legal* agents is a special case. This type of knowledge appears at a first glance to be almost negligible, but it has in practice a very important role. The legal system must regulate itself as just another social organisation. There is an important difference, though: the law can design the structure of the legal system as an organisation, in much the same way companies design their structure by their internal regulations. The creative knowledge performs this function.

3.2 Full FOLaw: dependencies between categories

So far, a set of categories that divide legal knowledge like pieces of a puzzle was presented. These categories are distinguished by their *function* in the legal system, and together realize the main function of the legal system: social control. Figure 1 shows how the functions/categories identified compose together this main function. Of course, it is not meant that there is a one to one correspondence between these abstract processes and the actual social processes and procedures in a legal system. Instead, they are functional dependencies which describe how the main function is decomposed in sub-functions which together perform it. In Figure 1, the rounded boxes represent functions (or, alternatively, bodies of knowledge which perform the function), and the solid arrows indicate functional dependencies, *i.e.* the inputs and outputs of these functions. The dependencies which correspond to actual interactions with the society are indicated in the figure in non-solid arrows. The entities in the society are specific social agents, *e.g.* the University of Amsterdam (a school), Joost Breuker (a private person), the Ministry of Education (a government agency), Philips BV (a company), etc.

A cycle starts with a *real world situation*, which is interpreted in order to generate an abstract description of the case in the terms that the legal sources use. This abstract case description is called a *legal situation*, and the knowledge used to produce this step is the *world knowledge*, which forms the *legal abstract model*. Then, the legal situation is analysed against the *normative knowledge* to verify whether it violates any norm, thus producing what is called a *classified situation* (a situation classified as either ‘allowed’ or ‘disallowed’). In another path, the situation is analysed using again *world knowledge* (but here particularly its *causal* component) in order to find out which agents in the world (if any) have caused the situation. This information is then used as input to the *responsibility knowledge* which determines which agents (if any) are to be held responsible for the situation. The results obtained in these two paths (the classified situation and the responsible agents) are then used as inputs for a function that defines a possible *legal reaction* using *reactive knowledge*. Further, outside this cycle, the law may also create an abstract entity (part of the legal system) using *creative knowledge*; this entity is also added to the *legal abstract model*. Finally, *meta-legal knowledge* refers to all these entities.

Another way to see the interdependencies shown in Figure 1 is that they provide the connections between the (sub)functions from a reasoning point of view. That is, legal reasoning can be

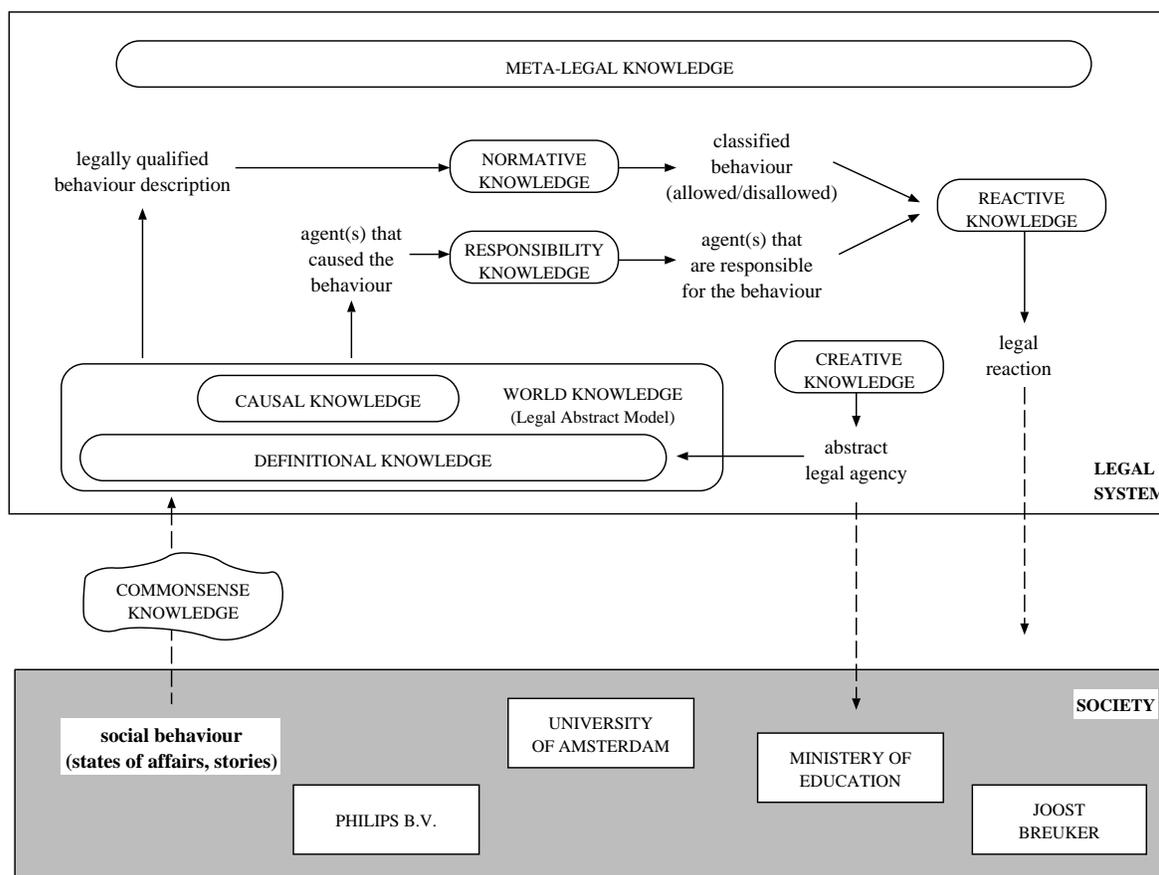


Figure 1: Functional roles of legal knowledge in the operation of the legal system.

made modular, with each function corresponding to a module and the dependencies between the modules being provided by the dependencies between the functions. Such dependencies must of course be detailed. Moreover, the main path in the picture can be also seen as the global structure of legal arguments: starting from the ‘facts of the case’ and going up to sentencing, with each function providing the source for argument steps in specific places within the global argument.

The scheme shown in Figure 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 the figure; for instance, an argument involving world knowledge (say, concluding that a certain person is considered a ‘minor’ according to a certain definition) being used as subsidiary for an argument involving normative knowledge (say, concluding that a situation in which this person was driving a car is disallowed according to a certain norm). Legal reasoning can be thus seen as the production and analysis of arguments involving one or more of these categories.

Ours is not the only core ontology for law. [Visser, 1995] and [vanKralingen, 1995] also proposed a core ontology. Also, the work by McCarty on a “language for legal discourse” can be

viewed as a core ontology for legal domains [McCarty, 1989]. Although the ontologies are structurally very different, there is an important overlap of categories. The fact that competing ontologies have been proposed has led to a reflective debate in the AI & Law community (see *e.g.* [Visser *et al.*, 1997], which also made it obvious that this debate has to be shared with legal theorists and legal philosophers.

3.3 Divide and conquer for educational design

The core ontology has been applied in knowledge acquisition in a number of legal domains, and it has been the basis for a legal reasoning prototype system, called ON-LINE [Valente, 1995]. ON-LINE is also the basis for a legal information server architecture, CLIME, which is tested and to be used in commercial applications [Winkels *et al.*, 1998, Winkels *et al.*, 1999]. The architecture of ON-LINE, respectively CLIME follows the ontology almost literally, *i.e.* it consists of reasoning modules for each category. The world knowledge is handled by the LOOM classifier [MacGregor, 1991], while for the normative reasoning module special inference mechanisms were developed which have a major advantage over (standard) deontic logic that they correspond to normative common sense, *i.e.* do not give rise to bizarre paradoxes, and that they are (more) tractable. The causality/responsibility module is the hardest nut to crack, and is still under investigation [?]. For reactive and creative knowledge very simple reasoning mechanisms are required.

Here we are mainly interested in educational applications. In principle, ON-LINE/CLIME could be re-used as the legal problem solving component of an IES. CLIME is capable of solving legal cases, represented in terms of its world knowledge. Moreover, CLIME is extended by explanation modules and by modules that allow one to pose questions about a regulation [Winkels *et al.*, 1999]. CLIME is still under development and constructing a training system for solving legal cases couldn't wait for practical reasons. Therefore, Antoinette Muntjewerff decided to use our legal core ontology as an analytic tool for the design of a conventional educational system, called PROSA,⁷ written in Authorware. PROSA coaches students in solving cases in Dutch administrative law (see [AWB,] for an English version of this law). This regulation, called AWB, is mainly concerned with procedures to appeal to decisions by authorities. However, these procedures, their conditions and other prescriptions are described by articles that do not constitute a coherent text; they have to be inferred, and, as usual, they are riddled with exceptions. Therefore, there is no simple mapping between a case at hand and the text of the law, and one has to go back and forth in order to identify applicable articles.

Before designing PROSA several empirical studies were undertaken on how students solved typical cases in administrative law. Administrative law is considered to be a relatively difficult area of law (for our students there is an unlikely high failure rate for examinations (approaching 70%); the examination consists of solving a number of cases). The general conclusion was that students are not able to analyse a case systematically. This is to a large extent due to the fact that it takes a long time to acquire insight in the mapping between the textual organization of the law and the conceptual structure that cannot be acquired by simple explanation.⁸ Another important problem is the identification of terms of the law in the case. Law is noteworthy for its vague terms, but an extra problem for students is that the world the law refers is too abstract. For instance, the term "administrative authority" in the AWB (1.1. art. 1) is defined as "an authority of (a) a legal person,

⁷PROSA stands for: PROblem Solving of cases in Adminstrative law

⁸It is possible to represent the backbone of this law as a kind of decision tree, but there are so many hairy side issues, that it provides a dangerous overgeneralization.

which has been established under public law, or (b) another person or body corporate which is invested with any public authority”. The first part of the definition refers to public law in general (and that is a world by itself), but that is not the major problem: the major problem is in the second part which is a catch-all to capture also “non-legal” persons. In these kind of problems one may easily recognize that they are concerned with the mapping of the common sense case descriptions to the legal terminology, *i.e.* the LAM.

The initial diagnosis appears to be that the students do not acquire sufficient insight into the major conceptual structure – the procedures and their dependencies –, and into the world of civil administration. However, the standard course really emphasizes these issues. The problem appears to be rather in the mappings between cases and ‘theory’, which is also a typical problem in *e.g.* solving physics problem, where students often fail because they do not take sufficiently care that the interpretation of the problem situation (case) is complete. Novice physics students map partial situation descriptions to (mathematical) constraints without careful checking whether the assumptions are satisfied: a kind of “jumping to conclusions”. Moreover, their lack of overview does not allow them to mentally keep track of sufficient covering of the problem situation. Similar observations were made about our law students. Therefore, in the design of PROSA a major emphasis was on the mapping between the LAM, *i.e.* the world knowledge implied by the AWB should be explicitly mapped onto the case description in such a concrete fashion that it also should act as an external memory that marks which propositions of the case have been covered by the law, and which have not: the latter may mean that the proposition is not relevant in legal terms or has been overlooked. This is accomplished by a spatial design of the user interface in PROSA (see Figure 2).

The left window contains the legal rules, *i.e.* the ‘theory’ that should be applied to the case description in the right window. The middle window is where the student constructs his solution by associating selected rules with selected facts; the specific problem posed to the student is put at the top of this window. The construction of a solution is accomplished by associating the facts of the case with the applicable rules. This mapping procedure is in the first place guided by the *terms* (concepts) in the AWB and those in the case. So in Figure 2 the “city council of Amsterdam” has made a decision (“order”) but a precondition for further analysis is whether this council is an “administrative authority”. If so, the city council can be associated with administrative authority, etc. The student’s work space allows him to keep track of his local decisions made, and as there is no prescribed method or order to the way he maps rules to facts, the student may work both ‘theory’ or ‘case’ driven. Therefore, in the end the student is capable to wrap up the evidence and come to a conclusion. As the associations can be evaluated (“assess-button”) by the student by asking PROSA so that the set can be pruned from wrong associations. Moreover, there is also completeness checking. A typical problem for students of law (and legal practitioners in general ⁹) is that the normative consequences may contradict one another, and that these contradictions – exceptions – have to be resolved: here PROSA helps in providing an overview and knows behind the screen about the solutions. However, the AWB has only a few articles that play a normative role, and there are only few exceptions: the RAWB is a typical ‘definitorial’ law. Therefore, we have little experience yet how PROSA may help students to come to grips with the problems of this category of knowledge.

Based upon Merrill’s Component Display Theory [Merrill, 1987b], the screen consists of two horizontal layers. The upper layer represents the “primary representations”, *i.e.* they present the

⁹We have found that this is also a serious problem even for the drafters of legal regulations [denHaan, 1996]

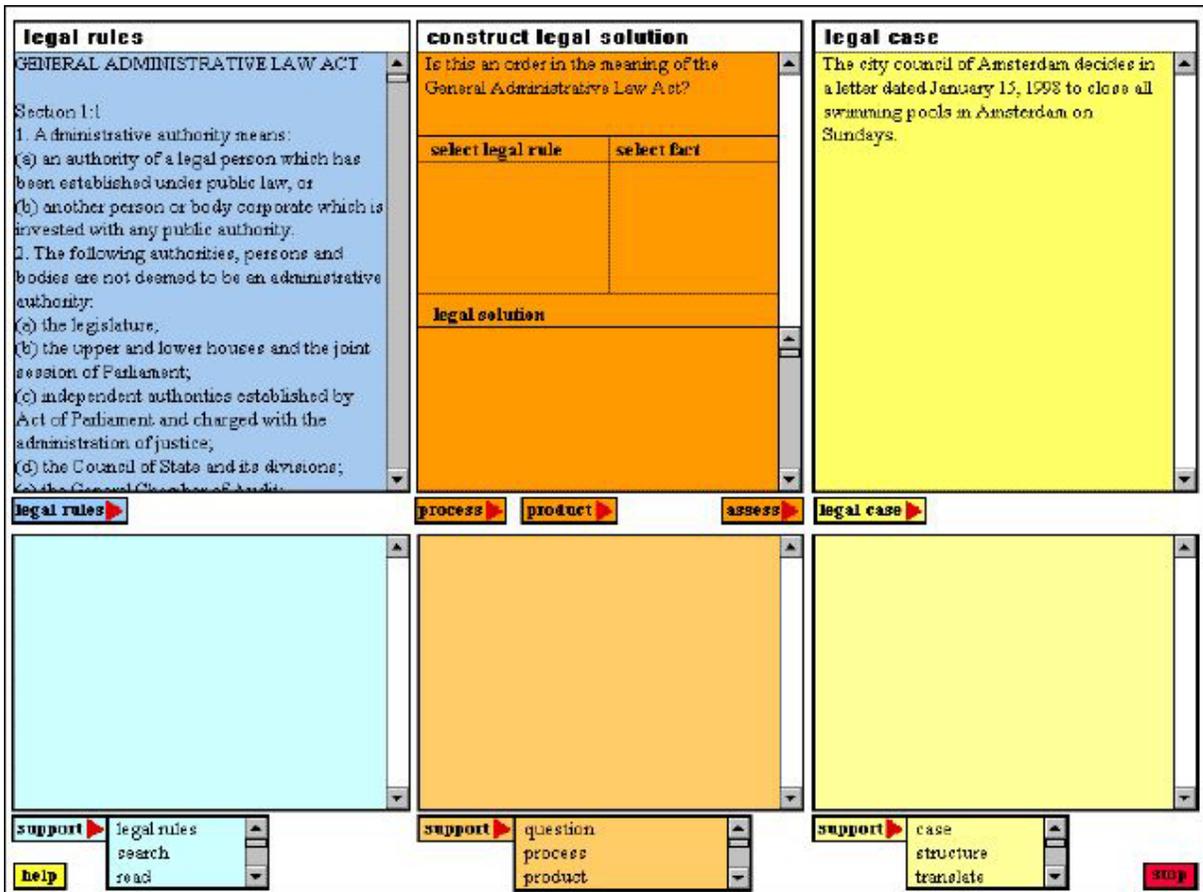


Figure 2: User interface of PROSA

instruction material in a direct way. The bottom layer provides the educational support (instructions, guidance, feed-back, hints, suggestions, etc.) for the primary layer.

Although PROSA does not have explicit knowledge about the domain, its construction behind the screens is based upon FOLaw. The AWB has been analysed in terms of the ontology and this has led both to a(n uneven) distribution of the content of the AWB over the various categories, and to a further detailing of FOLaw. All cases in PROSA (currently about 40) have been in similar terms.

- The *concepts* (terms) that map onto terms in the AWB, *i.e.* the definition knowledge subdivision of the ‘world’ knowledge (legal abstract model, LAM). As stated above, they are noteworthy troublesome in the AWB.
- The AWB is a typical procedural law in that it prescribes procedures and conditions for appeal. However, this kind of knowledge was not really foreseen in conceptualizing FOLaw [Valente, 1995] and led to long discussions as to whether these procedures were normative – they are prescriptions – or rather ‘definitions’. We opted for the latter, for various reasons: the most important being that there are no normative consequences. If one does not follow the procedures one does not violate a norm, but simply does not fit the definition. This is analogous to having or trying to obtain attributes that make one a specific kind of legal

agent. For instance, if one fits the terms of the British Nationality Act, one can obtain the status of British citizen (with its rights and duties), but for the law it only means that there is mapping between the definition and the case (agent) [Sergot, 1985]. Note that this analogical reasoning is also typical legal reasoning. An extra problem is that the AWB –as all laws do– leaves implicit what the *dependencies* are between the various procedures; they have to be mentally reconstructed when applying the law to a case. These dependencies are not made explicit to the student: this is part of the cursory education, but the solutions to the cases have been wired according to these dependencies.¹⁰

- Responsibility and normative knowledge are present in the AWB, but they do not give rise to the complications one may easily find in other domains of law.

PROSA is currently under evaluation. A large scale performance evaluation is planned for August 1999. An informal study with seven students has shown that PROSA's interface is remarkably easy to handle by the students. The most surprising comment from most of the students is that “they always worked the way PROSA suggests them to do”, while we know from the empirical studies that that is exactly what the students do not do. In other words, PROSA gives a very ‘natural’ look and feel to the students, both in content and in its operation.

4 Conclusions

The conclusions to be drawn here are only partially substantiated by what is presented here: the reason is that this article is not completely finished. It was foreseen that two other core ontologies and their consequences for designing educational or help systems would be discussed. However, time and planning forbid us at the moment to include our experiences in two other fields: information processing systems and ecology. We hope to make up for this omission in a next article on this subject.

First of all we should acknowledge the fact that the ontological engineering in constructing educational systems is not the new revolutionary thing in AI & Education. It is rather a re-appraisal of content issues in specifying educational systems. This basic knowledge acquisition has been performed in the building any educational system with an explicit representation of its subject matter, but the problem is that this work has hardly ever been explicitly reported. Ontological engineering at least provides the means and tools to communicate the results if not also the experiences.

A second lesson we have drawn from reflecting on our own work in ontology and domain representation for educational systems is that the content road may provide a far more natural way of dividing and conquering in the problem solving than a PSM (problem solving method) perspective. In methodologies for knowledge engineering – CommonKADS being one of the typical examples – the default route has been via task analysis which maps onto PSMs, etc [Breuker & Van de Velde, 1994]. However it appears that a functional view on the knowledge applied in a domain of practice provides a decomposition which is much easier to explain and handle, even if the categories and their dependencies may easily map onto acquired views on the task decomposition, as *e.g.* in medicine.

¹⁰It is not impossible to provide this kind of decision tree as well to the students. However, as there was no full consensus among the legal experts, and because the role of PROSA should remain pure application training we have not (yet) installed such a support.

Specifying ontologies – if done the right way – is more than putting a number of terms in some isa hierarchy. It means in the first place a careful definition of what terms mean in the context of use: it is a kind of semantic analysis. There is more: it is not only the meaning of terms – and their major relationships – but it also includes a specification of the point of view taken for its use, *i.e.* it should include a specification of its commitments. Of course, commitments may regress deep, too deep into top ontologies, but at least some justification of point of view is required: the least commitment being top node terms in ontologies. However, for educational purposes these commitments are very important, because they may reflect both assumed prior knowledge of students and a kind of stance what the domain is about. The commitments for FOLaw described in Section 3.1 are a good example of what is meant and can easily part of explicit instruction.

Finally, the recognition that ontological engineering may have a proper place in designing educational systems may have also as a side effect a more content oriented view on the educational process. Ontologies – in particular to and core ones – show us very explicitly what power there is in the ideas that constitute our beliefs about the physical and social world. According to us, these ideas have far more mileage than method. At least, a new emphasis on content aspects in education may give rise to reappraisal of acquiring insight vs acquiring skill, and of understanding and explanation over solutions to problems.

References

- [Allen, 1984] J. Allen. Towards a general theory of action and time. *ai*, 23:123–154, 1984.
- [AWB,] *Dutch General Administrative Law Act*. http://www.minjust.nl/a_beleid/thema/awb/teksten/awbeng/index.htm.
- [Bloom, 1957] B.S. Bloom. *A taxonomy of Educational Objectives*. McKay, 1957.
- [Borst & Akkermans, 1997] P. Borst and H. Akkermans. Engineering ontologies. *International Journal of Human-Computer Studies*, 46:365 – 408, 1997. special issue on using explicit ontologies in kbs development.
- [Breuker & Van de Velde, 1994] Joost Breuker and Walter Van de Velde, editors. *CommonKADS Library for Expertise Modeling: reusable problem solving components*. IOS-Press/Ohmsha, Amsterdam/Tokyo, 1994.
- [Breuker, 1999] J.A. Breuker. Indexing problem solving methods for reuse. In D. Fensel and R. Studer, editors, *Knowledge Acquisition, Modelling and Management*, pages 315 – 322. Springer, 1999.
- [Clancey, 1992] W. J. Clancey. Model construction operators. *Artificial Intelligence*, 53(1):1–115, 1992.

- [Davis *et al.*, 1993] R. Davis, H. Shrobe, and P. Szolovits. What is a knowledge representation? *AI Magazine*, pages 17–33, spring 1993.
- [denHaan, 1996] N. den Haan. *Automated Legal Reasoning*. PhD thesis, University of Amsterdam, 1996.
- [Domingue & Motta, 1999] J. Domingue and E. Motta. A knowledge based news server supported ontology driven story enrichment and knowledge retrieval. In D. Fensel and R. Studer, editors, *Knowledge Acquisition, Modelling and Management*, pages 103 – 120. Springer, 1999.
- [Genesereth & Nilsson, 1987] M. Genesereth and N. Nilsson. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, Los Altos, CA, 1987.
- [Gruber, 1994] T. Gruber. Toward principles for the design of ontologies used for knowledge sharing. In N. Guarino and R. Poli, editors, *Formal Ontology in Conceptual Analysis and Knowledge Representation*. Kluwer, 1994.
- [Gruber *et al.*, 1996] T. Gruber, G. Olsen, and J. Runkel. The configuration design ontologies and the VT elevator domain theory. *International Journal of Human-Computer Studies*, 44:569 – 598, 1996. special issue on Sisyphus-VT.
- [Guarino, 1995] N. Guarino. Formal ontology, conceptual analysis and knowledge representation. *International Journal of Human-Computer Studies*, 43:625 – 6640, 1995. Special issue on formal ontology.
- [Guarino, 1998] N. Guarino. Formal ontology and information systems. In N. Guarino, editor, *Formal ontology in information systems*, pages 3–15. IOS-Press, 1998.
- [Hart & Honore, 1985] H. Hart and T. Honore. *Causation in the Law*. Oxford University Press, New York, second edition edition, 1985.
- [Hirst, 1991] G. Hirst. Ontological assumptions on knowledge representation. In R.J. Brachman, H.J. Levesque, and R. Reiter, editors, *Principles of Knowledge Representation and Reasoning*, pages 157–169. Morgan Kaufmann, San Mateo, CA, 1991.
- [Hobbs, 1995] J.R. Hobbs. Sketch of an ontology underlying the way we talk about the world. *International Journal of Human-Computer Studies*, 43:819 – 830, 1995. Special issue on formal ontology.

- [Kant, 1787] I. Kant. *Kritik der Reine Vernunft*. Everyman, second edition, 1787. translation by J. Meiklejohn, revision by V. Politis for the Everyman's Library edition, 1993.
- [Kelsen, 1991] Hans Kelsen. *General Theory of Norms*. Clarendon Press, Oxford, 1991.
- [Knight *et al.*, 1995] K. Knight, I. Chander, M. Haines, V. Hatzivassiglou, E. H. Hovy, M. Iida, S.K. Luk, R.A. Whitney, and K. Yamahada. Filling knowledge gaps. In *Proceedings of IJCAI-95*, 1995.
- [MacGregor, 1991] R. MacGregor. Inside the LOOM classifier. *SIGART Bulletin*, 2(3):70–76, 1991.
- [McCarty, 1989] L.T. McCarty. A language for legal discourse i. basic structures. In *Proceedings of the 2nd International Conference on AI and Law*, Vancouver, 1989. ACM.
- [Merrill, 1987a] M.D. Merrill. An expert system for instructional design. *IEEE Expert*, 2:25–37, 1987.
- [Merrill, 1987b] M.D. Merrill. The new component design theory: instructional design for courseware authoring. *Instructional Science*, 16:19–34, 1987.
- [Miller, 1990] G. Miller. Wordnet: an on-line lexical database. *International Journal of Lexicography*, 3:special issue, 1990.
- [Mizoguchi *et al.*, 1997] R. Mizoguchi, M. Ikeda, and K. Sinita. Roles of shared ontology in AI-ED research: intelligence, conceptualization, standardization, and reusability. In B. du Boulay and R. Mizoguchi, editors, *Artificial Intelligence in Education, Proceedings of AI-ED 97*, pages 537 – 544. IOS-Press, 1997.
- [Motta *et al.*, 1996] E. Motta, A. Stutt, Z. Zdrahal, K. O'Hara, and N. Shadbolt. Solving vt in vital. *International Journal of Human-Computer Studies*, 44:333 – 372, 1996.
- [Muntjewerff & Groothuismink, 1999] A.J. Muntjewerff and J. Groothuismink. Supporting the learning of legal case solving using a computer program as the instructional environment. In M.H. Hamza, editor, *Proceedings of the IASTED International Conference Computers and Advanced Technology in Education (CATE'99)*. IASTED/Acta Press, 1999.
- [Neches *et al.*, 1991] R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W. R. Swartout. Enabling technology for knowledge sharing. *AI Magazine*, pages 36–56, Fall 1991.

- [Patil *et al.*, 1992] R. Patil, R. Fikes, P. Patel-Schneider, D. McKay, T. Finin, T. Gruber, and R. Neches. The DARPA knowledge sharing effort: Progress report. In C. Rich, B. Nebel, and W. Swartout, editors, *Proc. of KR'92*. Morgan Kaufmann, 1992.
- [Sergot, 1985] M.J. Sergot. Representing legislation as logic programs. In J.E. Hayes, D. Michie, and J. Richards, editors, *Machine Intelligence 11*, pages 209–260. Oxford University Press, 1985.
- [Sowa, 1995] J.F. Sowa. Top-level ontological categories. *International Journal of Human-Computer Studies*, 43:669 – 686, 1995. Special issue on formal ontology.
- [Valente & Breuker, 1996] A. Valente and J. Breuker. Towards principled core ontologies. In B.R. Gaines and M. Mussen, editors, *Proceedings of the KAW-96*, Banff, Ca, 1996.
- [Valente, 1995] A. Valente. *A Modelling Approach to Legal Knowledge Engineering*. IOS Press, Amsterdam, 1995.
- [Valente *et al.*, 1999] A. Valente, J.A. Breuker, and P.W. Brouwer. Legal modelling and automated reasoning with ON-LINE. *International Journal of Human Computer Studies*, 50:xxx–yyy, 1999. to appear.
- [vanHeijst *et al.*, 1997] G. van Heijst, A.Th. Schreiber, and B.J. Wielinga. Using explicit ontologies in kbs development. *IJHCS*, pages 183 – 291, 1997.
- [vanKralingen, 1995] R. W. van Kralingen. *Frame-based Conceptual Models of Statute Law*. PhD thesis, University of Leiden, The Hague, The Netherlands, 1995.
- [Visser, 1995] P.R.S. Visser. *Knowledge Specification for Multiple Legal Tasks; A Case Study of the Interaction Problem in the Legal Domain*. PhD thesis, University of Leiden, The Hague, The Netherlands, June 1995.
- [Visser *et al.*, 1997] P. Visser, D.M. Jones, T. Bench-Capon, and M. Shave. An analysis of ontology mismatches; heterogeneity vs interoperability. In *Proceedings of AAI-Symposium on Ontological Engineering*, pages 164 – 172, Stanford, CA, 1997. AAI.
- [Wielinga *et al.*, 1992] B. J. Wielinga, A. Th. Schreiber, and J. A. Breuker. KADS: A modelling approach to knowledge engineering. *KAJ*, 4(1):5–53, 1992.

[Winkels *et al.*, 1998]

R. Winkels, A. Boer, D. Bosscher, and J.A. Breuker. Assessment based legal information serving and cooperative dialogue in CLIME. In J. Hage, editor, *Proceedings of JURIX-98*, pages 15–24. Vermande, 1998.

[Winkels *et al.*, 1999]

R. Winkels, D. Bosscher, A. Boer, and J.A. Breuker. Generating exception structures for legal information serving. In Th. Gordon, editor, *Proceedings of the International Conference on AI and Law*. ACM, 1999.