

# A Legal Case OWL Ontology with an Instantiation of *Popov v. Hayashi*

Adam Wyner

Department of Computer Science, University College London  
Gower Street, London WC1E 6BT, United Kingdom  
adam@wyner.info

Rinke Hoekstra

AI Department, Faculty of Sciences, Vrije Universiteit  
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands  
hoekstra@few.vu.nl  
Leibniz Center for Law, Faculty of Law, Universiteit van Amsterdam  
Kloveniersburgwal 48, 1012 CX Amsterdam, The Netherlands  
hoekstra@uva.nl

## Abstract

The paper provides an OWL ontology for legal cases with an instantiation of the legal case *Popov v. Hayashi*. The ontology makes explicit the conceptual knowledge of the legal case domain, supports reasoning about the domain, and can be used to annotate the text of cases, which in turn can be used to populate the ontology. A populated ontology is a case base which can be used for information retrieval, information extraction, and case based reasoning. The ontology contains not only elements of indexing the case (e.g. the parties, jurisdiction, and date), but as well elements used to reason to a decision such as argument schemes and the components input to the schemes. We use the Protégé ontology editor and knowledge acquisition system, current guidelines for ontology development, and tools for visual and linguistic presentation of the ontology.

## 1 Introduction

In this paper, we provide an OWL ontology for legal cases with an instantiation of a particular legal case – *Popov v. Hayashi*.<sup>1</sup> An ontology makes explicit the conceptual knowledge of the legal case domain, supports reasoning about the domain, and can be used for information retrieval and text annotation. A populated ontology, whether by manual entry or automatic extraction, comprises a case base which can be used for case based reasoning. We use the Protégé ontology editor and knowledge acquisition system, current guidelines for ontology development, and tools for visual and linguistic presentation of the ontology.

<sup>1</sup>©2009 Adam Wyner and Rinke Hoekstra. A previous version of this paper was presented at the Workshop on Modeling Legal Cases, International Conference on Artificial Intelligence and Law, Barcelona, June 8, 2009 and appears in Wyner (2009). The OWL ontology which is discussed here is available upon request from the authors. The case citation is: *Popov v. Hayashi*, 2002 WL 31833731 (Cal.Superior Dec 18, 2002) (NO. 400545).

The elements of the ontology are drawn from a range of sources including typical legal search features, previous theoretical research on legal case ontologies, and research on case based reasoning. For the purposes of this paper, we emphasise the elements of the legal case ontology which are relevant to annotating cases as well as managing, tracking, and checking a case database. To exercise the ontology, we model in the ontology the decision of *Popov v. Hayashi*, which has been chosen because it relates to a body of work concerning property (e.g. the wild animal cases in Berman and Hafner (1993)) and has been the focus of some research on modelling specific legal cases (Wyner et al. (2007); Wyner (2009)). Due to the size of the ontology, the scope of the paper is restricted to the identification of the main ontological elements of cases including classes, object and data properties, class restrictions, and individuals; the whole ontology is more easily examined using the OWL file and an ontology editor such as Protégé. As is discussed in the course of the presentation, a range of issues are not fully addressed such as legal reasoning rules, details of particular legal hierarchies, and full lists of legal subjects or causes of action, among other modules; these are left for future development. Nonetheless, this ontology is designed to be forward compatible with further development of these modules. In addition, we do not incorporate additional components required to apply Description Logic (DL) reasoners to do case based reasoning (but see Wyner (2008)), though the structure of the ontology is intended to support such components. The novelty of the paper is: a fuller, more explicit ontology of legal cases which follow current methodologies of ontology design, the presentation of the ontology in rich graphic and linguistic modes, elaboration of the role of argument schemes, and the instantiation of the case *Popov v. Hayashi* in this ontology.

In section 2, we provide an overview of the representation technique and analytic parameters, outlining general issues of ontology development, our methodology, the graphical and textual modes of representing the ontology, then the scope and context of the ontology. Section 3 presents our ontology for legal cases. This is followed by the ontological representation of *Popov v. Hayashi* in section 4. Related work and future research is discussed in section 5.

## 2 Overview of Representation Technique

In this section, we provide brief overviews about ontologies in the Web Ontology Language (OWL) and the Protégé ontology editor and knowledge acquisition tool, our methodology, the modes of presentation of the ontology, the scope and context of use of the ontology.

### 2.1 OWL Ontologies and Protégé

We outline some of the main reasons for providing an ontology in OWL 2 (henceforth just OWL) (Uschold and Gruninger (1996); Motik et al. (2009))<sup>2</sup> and using the ontology development tool Protégé 4 (henceforth just Protégé).<sup>3</sup>

An ontology explicitly and formally defines the concepts and relations that may exist in a given domain (Gruber (1993); Antoniou and van Harmelen (2004); Antoniou et al. (2007); Uschold and Gruninger (1996)). A common vocabulary and framework is provided in an ontology so that researchers can share, test, and modify the conceptualisation. Broadly speaking, an OWL 2 ontology consists of *classes*, the subsumption

<sup>2</sup>See <http://www.w3.org/TR/owl-features/>

<sup>3</sup>See <http://protege.stanford.edu>

relations between them, object properties that relate to instances of classes, and restrictions on what properties may hold of these instances. In a well-designed ontology, one can make *inferences* about classes and the types of individuals (i.e. the classes to which they belong); for example if Jane Doe is the solicitor for the plaintiff in *Popov v. Hayashi* and Patrick Hayashi is the plaintiff in *Popov v. Hayashi*, then we should infer rather than explicitly state that Jane Doe is the solicitor for Patrick Hayashi in *Popov v. Hayashi*. While we give sample inferences for *Popov v. Hayashi*, we do not discuss inferences in depth as the focus of the paper is the ontology itself and a particular instantiation. Given instances of the classes, we have a *knowledge base*. Ontologies have proven to be useful in the legal domain (Gruber (1993); Bench-Capon and Visser (1996); Kralingen et al. (1999); Breuker et al. (2004); Valente (2005); Sartor (2006); Hoekstra et al. (2009)).

We have developed our case ontology in OWL using Protégé. OWL provides a machine readable ontology which can then be processed by *Semantic Web* applications. OWL provides a range of *flavours*, each associated with a degree of logical expressiveness and associated computational properties; for our purposes, we have used OWL 2 DL (which is the DL language  $\mathcal{SROIQ}(\mathcal{D})$ ) which uses qualified cardinality restrictions, complex role inclusion axioms (property chains), equivalent classes, nominals (enumerated classes), existential and universal restrictions, functional properties, and functional data type properties, among other aspects. The Protégé ontology editing and knowledge acquisition tool supports systematic development of an ontology along with structured instantiations. Protégé enables users to *query* the knowledge base, test an ontology for *consistency*, draw inferences, and apply rules to elements of the knowledge base. However, as our focus is on formally modeling a legal case, we primarily present the ontology; moreover, reasoning about cases presupposes the ontology.

## 2.2 Methodology of Ontology Design

The methodology of designing an ontology is a current, open topic. Neither the formal language of DL on which OWL 2 DL is based nor the development tools prescribe the structure of an ontology. For our purposes, we have followed *best practices* from Hoekstra (2009).

One common approach to ontology design is to *minimise* reasoning by highly specifying the ontology by (among other things): asserting the classes and subclasses as well as the properties of individuals, asserting disjointness of classes, and fixing the class membership of individuals. Moreover, the class structure is largely stipulated rather than associated with properties or inferred. In contrast, the approach we follow here attempts to *maximise* reasoning by (among other things): deriving classes and subclasses with respect to class restrictions, deriving disjointness, and inferring the class membership of individuals. Furthermore, the class structure is tied tightly to underlying properties, using the properties to *describe* and *restrict* classes. Maximising reasoning allows individuals to vary their properties depending on the context in which they appear; we infer the class type of an individual relative to the properties they bear with respect to the ontology. In this regard, classes describe *roles* played by individuals relative to the context in which they appear. A particular individual may be a plaintiff in one case, but a defendant or witness in another; individuals can cease to be a plaintiff in a case without ceasing to exist (anti-rigidity, according to Guarino and Welty (2002, 2004)). We briefly illustrate these points.

Asserting class hierarchies leaves information about class relationships implicit. For example, suppose we have a class of *animals* which has two subclasses *dogs* and *cats*. Each of these classes has two subclasses *Spaniel* and *Jack Terrier* for dogs and *Persian* and *Tabby* for cats. However, this class hierarchy tells us little about *what* defines the classes, that is, the properties which distinguish dogs from cats, one breed from another. Moreover, we cannot draw further inferences from known information; that is, given an animal and asserting that it is a member of the *Tabby* class does not imply anything about its properties or whether it is or is not also a member of the *Jack Terrier* class. On the other hand, if we specified the classes with respect to the properties that define them, then we may make inferences. For example, suppose that Spaniels make good family pets but are not intelligent, while Jack Terriers make good shepherding dogs and are very intelligent. In addition, we have the axiom that good family pets are not good shepherding pets. Then, from simply knowing that one has a good shepherding dog, we can infer that the dog is an intelligent dog, not a good family pet, and a Jack Terrier, while knowing one has a Spaniel, we can infer that it makes a good family pet. The instantiated ontology is inconsistent were a dog is both intelligent and a good family pet. In some other context such as import or trade restrictions on animals, these properties may be replaced by others. Thus, using the properties to define the classes makes the knowledge explicit, supports inference concerning the properties of individuals, untethers the individuals from context, and allows easier detection of inconsistencies.

In designing the legal case ontology, we have identified the important concepts from the literature on legal case based reasoning, considered keywords in legal case search engines (e.g. Lexis-Nexis), and considered what appeared relevant in the representation of *Popov v. Hayashi*. An emerging methodology is to apply automated text mining to generate the ontology (Maynard et al. (2008) Peters (2009)); however, the results of this method need to be improved, and we leave in depth comparison for future research.

### 2.3 Modes of Presentation

We provide graphical and textual representations of the case, which serve different presentational purposes.


#### 2.3.1 Graphical Representation

Graphics can provide succinct representations of information and are particularly useful where the *linearity* of linguistic information interferes with comprehension. In Figure 1, ovals are classes (C1, C2), dashed lines are inferred relations (C1 favours C3), dotted boxes are 'restrictions' (C3 and C4), open arrows indicate subclass relations (C3 and C4 has two subclasses), closed arrows indicate property relations (C1 favours C3).

However, useful graphics may be, they still must be interpreted in order to relate to the underlying linguistically represented information of legal cases. Moreover, complex graphics can be difficult to read and present. Therefore, we also use an alternative, linguistic representation of the ontology.

#### 2.3.2 ACE View

Since the ontology is meant to be used and understood by legal professionals, not knowledge engineers, a representation that hides some of the logical formulation is a must. We looked primarily at ACE View (Coi et al. (2009)), but also considered the two alternative syntaxes discussed in Schwitter et al. (2008), such as Rabbit and 'Sidney



**Figure 1** Example Graphic

Syntax'. Although OWL 2 provides a standard human-readable syntax in the Manchester Syntax<sup>4</sup>, this is still too close to the logical underpinning of OWL to be comfortable to our prospective users.

*ACE View* is a Protégé 4 plug in which is an ontology and rule editor that uses Attempto Controlled English (ACE) to create, view and edit OWL 2 ontologies and SWRL rulesets.<sup>5</sup> Attempto Controlled English (ACE) is a controlled natural language, where a controlled language is a formal specification of a subset of English with a large vocabulary and an expressive set of grammatical constructions. Once having learned how to author expressions (which is relatively easy), expressions in ACE write and read as standard English. As a formal language, ACE texts are computer-processable and unambiguous. Sentences are parsed and translated into *Prolog* expressions of *Discourse Representation Structures*, which provide a syntactic variant of first-order logic; ACE provides for consistency checks and inferencing. As Prolog expressions, the ontology could be a component of a more extensive Prolog program. *ACE View* supports bidirectional translation of expressions between ACE and OWL 2 DL. With *ACE View*, one can create OWL/SWRL knowledge bases, open OWL 2 ontologies and view them as ACE texts, edit OWL/SWRL knowledge bases, and query the knowledge base in natural language. Relative to our context of use, the linguistic representation in the ontology maintains a closer link to the linguistic representation of the case. It is useful for checking the ontology against semantic intuitions that are linguistically expressed. It is also a more accessible format for legal professionals who may then be involved in developing and maintaining the ontology. We have used *ACE View* as one tool to facilitate the development and presentation of this ontology as in section 4; see Wyner (2008) for further discussion.

#### 2.4 Scoping the Legal Case Ontology

A range of structured lexicons, ontologies, or design patterns for the legal domain have been proposed such as LOIS (Peters et al. (2007); Schweighofer and Liebwald (2007)), LKIF (Hoekstra et al. (2007, 2009)), CLO (Gangemi et al. (2005)), or design patterns (Gangemi (2007)). While these are very significant contributions to our understanding of ontologies in general, and legal ontologies in specific, they are too abstract and general for our purposes, for they relate to higher level ontologies concerning intentionality,

<sup>4</sup>See <http://www.w3.org/TR/owl-manchester-syntax/>.

<sup>5</sup>See <http://attempto.ifi.uzh.ch/>

processes, time, a wide range of entities and classes, or generic design issues. Ashley (2009) proposes ontological requirements for legal reasoning about cases. He sets out three roles for legal case ontologies: to support case-based reasoning, to distinguish between deep and shallow analogies, and to induce and test hypotheses. Indeed, the broader scope is to have an ontological representation of all those components case-based reasoning arguments which a law professor and students might make use of in discussing a legal case. While this range of roles is suggested, it is unclear whether it is intended that one ontology suit all these roles, or whether there may be ontological modules which can be developed which then integrate to support all the roles. Moreover, it does not clearly distinguish between an ontology which represents individuals, classes, and properties as distinct from the rules which apply to them or reason about them.

Our approach scopes the issues significantly. First, we focus on the ontological commitments required to support case based reasoning, but not the rules which relate and compare cases one to the other, which is the essence of the case based reasoning process, though it is our intention that the ontology be used as a basis for such reasoning. Second, we suppose there are ontological modules which further elaborate ontological elements which are found in our ontology such as ontologies for factors, for legal rules, and for legal theories. There are a range of intermediate legal concepts in cases concerning, for example, *guardianship*, *trusteeship*, and *possession*, where an intermediate legal concept is a concept defined by simpler concepts and which is used to define some higher level concept (Wyner (2008)). The ontology and rules to reason about these concepts ought to be found in a distinct ontology which represents them as they are extracted and abstracted from the contexts in which they appear. That a case uses the term *guardianship* does not imply that our ontology for cases must incorporate a complete ontology for intermediate concepts.

A more general third point is that we have only included that which distinctly is used to retrieve or reason with legal cases. In other words, those higher or extra domain ontologies which are used in legal as well as non-legal contexts are not incorporated. In particular, we have no ontological representation of events/processes, causation, time, space, propositional attitudes, subsorts of cognitive agents, and other elements. Nor do we have ontological representations of specific domains, for example, the ontology of baseball to represent *Popov v Hayashi*. In our view, there should be distinct ontologies for these domains that are imported and used as and where needed in a case.

It is no doubt true that case decisions report reasoning concerning these upper or extra domain elements for they appear as constituents of the factors and facts of the case; it is, therefore, tempting to incorporate them into the legal ontology since they are, after all, what is reasoned about in the course of arriving at a legal decision. Take for example, the case *Popov v Hayashi*, where authorities were brought to clarify matters concerning complete or incomplete actions. While in a legal context, where the clarification mattered, the authorities indulge in informal and unempirical debate about natural language semantics and psychological categories. In other instances, courts are known to engage in informal reasoning about causation, intentionality, and the meanings of words (Solan and Darley (2001); Solan (2003, 2005)). In general, since the law can bear on any matter concerning human affairs, it gives rise to the impression that so too must ontological research on the law. Yet, by the same token, there is no reason to incorporate them since what we know and how we reason about these elements is, often, the same outside the legal scope as within it; whether an action is completed or not can be debated within a court room or at a bar, reaching the same conclusions with the same

evidence and forms of argument. Therefore, for our purpose, a legal ontology of these domains should be developed only where the elements or relations among elements in the domain have a *distinct conceptual structure* from the non-legal conceptual structure.

For example, suppose a case concerning a restaurant, a salad, and a vegetarian who eats no fruit. The vegetarian is served a salad at the restaurant which contains a tomato after being assured by the restaurant there would be no fruit in her food. She refuses the food, but is obliged to pay. She sues, saying the restaurant misrepresented the food. It may be that the biological ontology which bears on tomatoes does not have the right structure or meaning as it is used in the legal context; in a legal context, a case a tomato might be a vegetable which can accompany other vegetables in a salad, while in a biological context, it is a fruit. The vegetarian applied the biological ontology, while the restaurant applied the legal ontology. In the context of a restaurant and a court case where legal definitions and obligations hold sway, a tomato is legally classified as a vegetable, so the vegetarian loses the case. This example shows that the ontology is there for us to capture a certain perspective; much of what we provide in our ontology are roles, subjective or context dependent entities that we can superimpose on a description in reality. Multiple perspectives may exist on the same domain, but these are of no concern to us as we can abstract away from them. This contrasts to an “absolute”, objective, context independent, or omniscient ontological representation of a domain. Where time, space, events, propositional attitudes, and social roles have the same conceptual structures outside of the legal structure as inside of it, then we do not presume to require that a legal ontology represent them, though the ontology should have the means to incorporate them as need be.

By the same token, an ontology does not itself represent the *dynamics* of reasoning (i.e. argumentation) or modeling, but the *static* data over which participants argue. For example, each case has a cause of action which is put forth at the pleadings phase of the legal procedure, for example, *trespass to chattel*. We assume the *outcome* of the pleadings phase and represent that a cause of action such as *trespass to chattel* holds in the case; in other words, we do not represent how it comes to be that the cause of action holds of a case. Similar considerations can be made of the introduction of evidence or of the judgement. At this point of our analysis, the ontology largely relates to the properties as they are given in the case decision. To represent how the values of such elements are assigned over the course of a legal procedure, one would have to introduce rules of legal reasoning (Gordon (1993); Wyner et al. (2007); Gordon et al. (2007a)); procedural issues are beyond the scope of this ontology at this time. Nonetheless, we may, where feasible, derive properties; for example, if a case has precedents, we may want this to be implied by one case citing another case which temporally precedes the citing case.

Finally, not everything found in the case decision necessarily falls under the scope of the ontology: as judges appear to have wide latitude in reporting case decisions, they may include a range of considerations about legal theories, evidence, witnesses, procedures, and so on, rather than simply the key components of the result. We focus largely on the result, which is particularly relevant to case summaries, case based reasoning, and Shepardization.

We believe abstraction and modularisation are reasonable and practical, making ontological representation of cases feasible. Other modules can be added, thus incrementally and systematically building and integrating an overall legal ontology and legal reasoning system (broadly compatible with Costa et al. (1998)).

## 2.5 Context of Use

Formal, machine-readable ontologies have a range of contexts of use and applications. Ontologies, as representations of knowledge, help us to understand the knowledge under investigation. With a web-oriented tool such as Protégé, ontologies can be used to represent and reason about some domain of interest with a web-based application. Moreover, as a structured representation of knowledge, an ontology can facilitate knowledge acquisition, supporting users to build up large, instantiated databases which can subsequently be queried or reasoned with. Building databases can either be a manual task (users enter values for fields) or automated with information extraction and retrieval techniques (where those techniques can successfully and systematically identify the relevant fields). The latter are usually supported (initially at least) by a *gold standard* corpus of texts which have been annotated relative to the ontology.

In our particular domain of interest, legal case representation and case-based reasoning, these contexts of use focus on the legal domain (see uses in Valente (2005)). The ontology could be used to make explicit the implicit knowledge of legal cases which legal professionals have; it could be used as a tool to build a database of cases; with a database of cases, we could apply automated case-based reasoning rules and query the knowledge base; with a web-based tool, case representation and reasoning could be done over the internet. At this point, building the case base is a manual task, as automated text mining are not yet sufficiently well-developed to bind textual cases to the fields of the ontology (Daniels and Rissland (1997); Brüninghaus and Ashley (1997, 2005b); Jackson et al. (2003); Moens et al. (2007); Sporleder and Lascarides (2006); Schweighofer (1999)). Having a well-developed ontology provides the fields to bind to. Finally, an ontology could be used to support legal training by providing an explicit representation of case knowledge that can be tied to particular instances; students and researchers might find it helpful to apply an analytic tool such as an ontology to cases under study. However, such uses remain for future research.

## 3 Ontological Representation of Legal Cases

In the following, we have all the main classes which relate to the ontological representation of a case. The structure of the ontology is derived from a range of sources and serve different purposes; we refer to these sources in the discussion section.

We have divided the presentation into an abstract representation of a case ontology (this section) and a particular instantiation with *Popov v. Hayashi* (section 4). Some elements are included only schematically while others are represented more fully; future development and instantiation will elaborate these. As there are more elements of the current version of the ontology than can be feasibly represented here, one must consult the ontology for further details. Our presentation touches only on some of the high points.

The presentation follows the key high level divisions of the classes of the ontology, starting with the *Case* class which is the key concept of the ontology and which draws together the other components (the presentation includes a small class *Hearing* as well). We then turn to *Decision* and *Jurisdiction* classes, which are relatively straightforward. The *Participant* class contains a range of participants in a case. We have included an *Argument Scheme* class, which represents patterns of reasoning behind the decision. A rich class labelled *Element* contains a spectrum of information that connects to the argument schemes, are used to construct the decision, and thence links the information

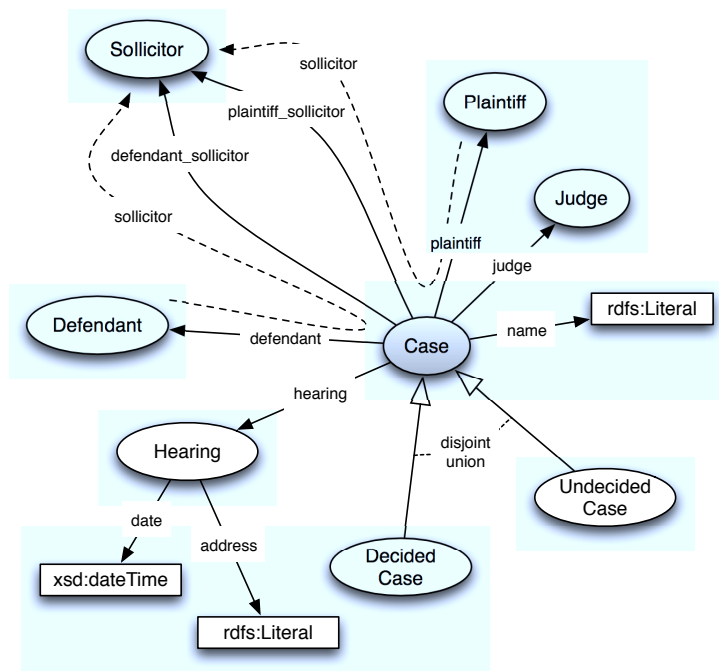


Figure 2 Case

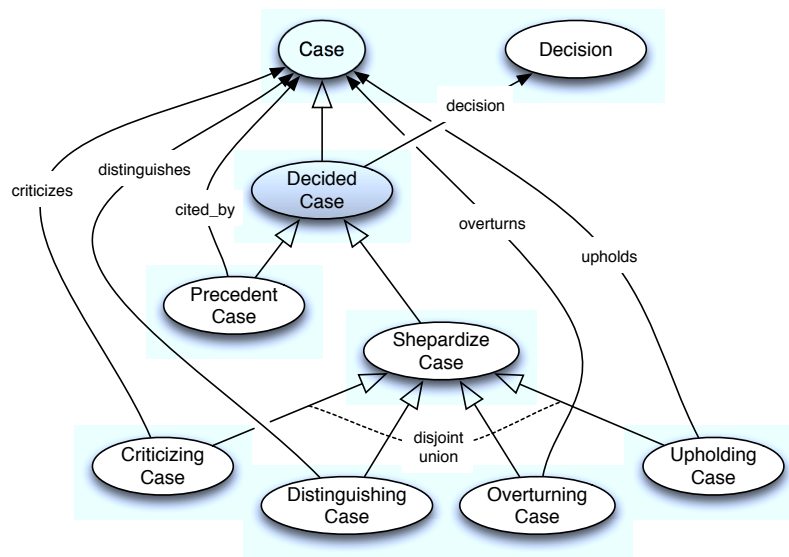
to a particular case. For each of the classes or subclasses, we highlight some of the relevant topics, restrictions, and properties.

### 3.1 Case and Hearing

In Figure 2, we represent the *Case* class, which is equivalent to: the class comprised of individuals which have a defendant, a plaintiff, and a judge; the class comprised of the union of decided and undecided cases. The *Case* class is also a subclass of classes defined by object and data properties such as having solicitors, a hearing, a name, and a jurisdiction. Among the members of the *Case* class, we have cases such as *Popov v. Hayashi* and *Young v. Hitchens*.

We have chosen to *under-define* the class of *Case*; that is, we have not asserted that every thing which is of type *Case* must, for example, have evidence. Rather, we have object properties which would, in a given instance, enable us to assert of a case (or subclass of cases) that it has evidence or a jurisdiction; in this way, specialised classes of *Case* can be constructed from properties. The two subclasses of *Case* are distinguished in that only decided cases have a case decision. In addition to the *Case* class, we have a simple class for *Hearing*, which has as superclass the individuals which have a date and an address.

In Figure 3, we elaborate *Decided Case*, which are those cases which have a *Decision*, which is discussed in the next section. The *Decided Case* class has two subclasses – *Precedent Case* and *Shepardize Case*. Precedent cases are decided cases which are cited by another case; since citation of one case in another is explicit, we do not need to otherwise calculate (say by the date) the precedential relationships. We have included under decided cases a hierarchy to represent *Shepardization*. In searching the case base, one needs to find not only precedent cases with the relevant factors, but also to identify



**Figure 3** Decided Case

*good* law relative to one's purpose. For example, if one finds a case with the relevant factors and a decision in favour of one's client, then the legal researcher wants to know if the decision in that precedent is still being followed and by which level or court in the legal hierarchy. Alternatively, a legal researcher may seek cases associated with the opponent's position which have been overturned. This is standardly referred to as *Shepardization* of a case, and a range of case citation systems exist to identify and relate the relevant cases. Shepardization can refer to the history of a case decision, e.g. whether a later case decision affirms, modifies, reverses, etc an earlier decision; it also relates to the treatment of the decision, e.g. whether the later decision criticises, distinguishes, explains, etc the earlier decision. For our purposes, we have object properties (e.g. criticizes a case) to define the subclasses of cases. Aspects of shepardization in the ontology could be developed further.

### 3.2 Decision

Figure 4 shows that the *Decision* class, which specifies that decisions favour a party, have a rationale related to an argument scheme, and have a date on which they are given. A decision for the plaintiff favours the plaintiff party; a decision for the defendant favours the defendant party; and a split decision favours both. The latter subclass is relevant to the decision of *Popov v. Hayashi*.

### 3.3 Jurisdiction

Jurisdiction refers to the authoritative body with respect to which the case is adjudicated, which we elaborate as in Figure 5. There are a variety of ways that jurisdiction may be decomposed – national boundary, judicial hierarchy, and subject matter. Subject matter refers to whether the court sits under civil or criminal proceedings, whether it concerns family law, taxation, etc. We only consider two national jurisdictions (US and UK) within which we have a legal hierarchy which is given by the *subordinates* object

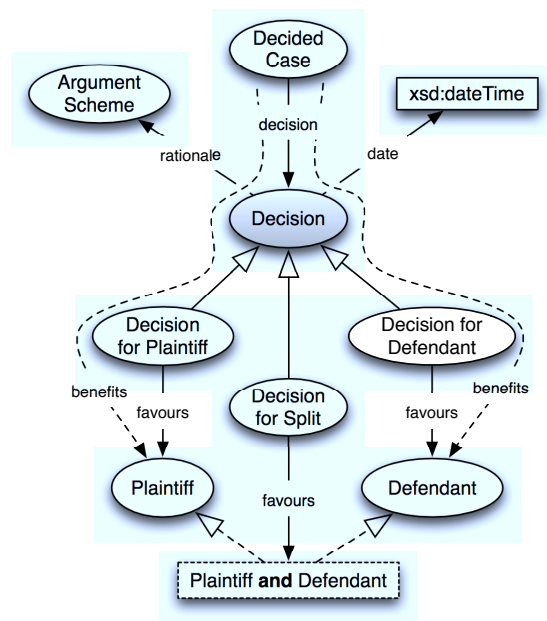


Figure 4 Decision

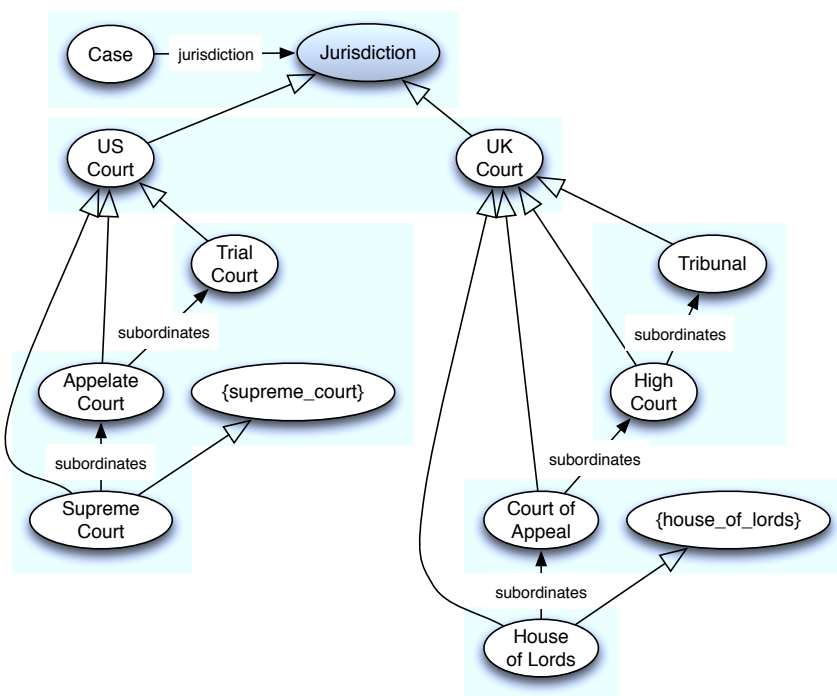
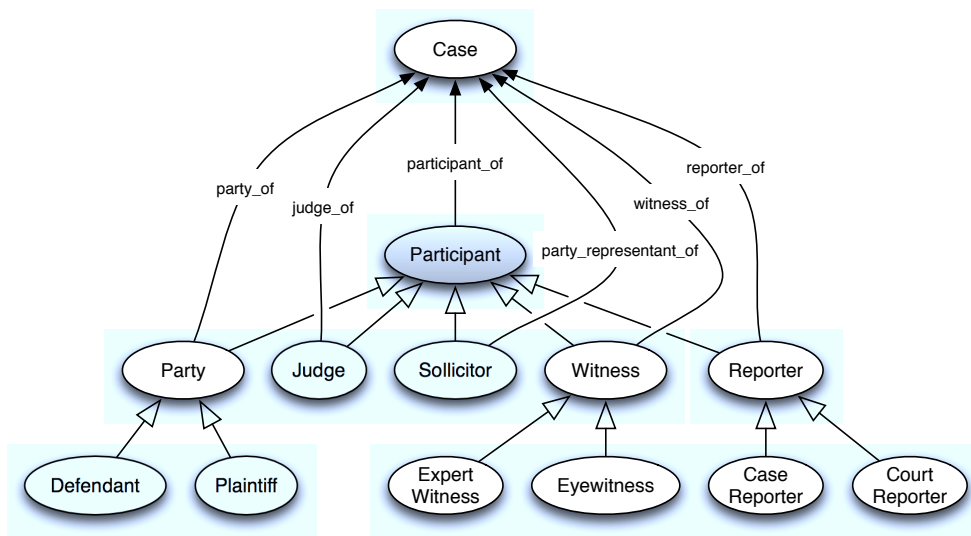


Figure 5 Jurisdiction

property. We can infer that the UK House of Lords subordinates every UK Tribunal. The classes for Supreme Court and House of Lords are restricted to singleton sets.

Matters of jurisdiction are considerably more complex and need to be articulated further. For instance, we have not represented here the notion of *binding decisions*, where courts of a higher level or even within the same level, must be followed by a



**Figure 6** Participant

court. Nor have we represented restrictions on the *appeals* process, where decisions may be challenged by the parties and passed to another court at a higher level of the judicial hierarchy to be decided. Individual courts should be associated with addresses, sitting judges, and related details. Finally, we have no representation of US circuit courts, nor the distinct UK jurisdictions of England, Scotland, Wales, and Northern Ireland.

### 3.4 Participant

In Figure 6, we represent a range of participants in a case, from the judges, solicitors, and reporters to the parties in conflict and the witnesses called. Some components are more relevant for information extraction (judge, solicitor, witnesses, and reporters) than they might be for case based reasoning. For example, for reasoning it is sufficient to know that a claim was made by a witness, not who the witness is.

### 3.5 Argument Scheme

Argument schemes are prototypical, defeasible reasoning patterns (Walton (2002)). For example, the *Eyewitness Testimony Argument Scheme* is: X claims to be an eyewitness to B, an eyewitness reports what happened, therefore it is plausibly true that B happened. Wyner et al. (2007) analyse the legal reasoning in *Popov v. Hayashi* with a range of argument schemes; Wyner and Bench-Capon (2007) propose legal argument schemes for case based reasoning. We hypothesise that every case includes some such schemes since they are claimed to be the ways in which legal reasoning is done; indeed, the notion of argument scheme could be generalised to cover all of the reasoning patterns found in a case, including legal theories, legal rules, and causes of action, though we have left this open.

In Figure 7, we have include only as *classes* those schemes that are relevant to *Popov v. Hayashi*, and only some of the schemes have more developed structure – the *Eyewitness* and *Expert* schemes. Presumably, in an ontology which incorporates case based reasoning directly, the argument schemes of Wyner and Bench-Capon (2007)

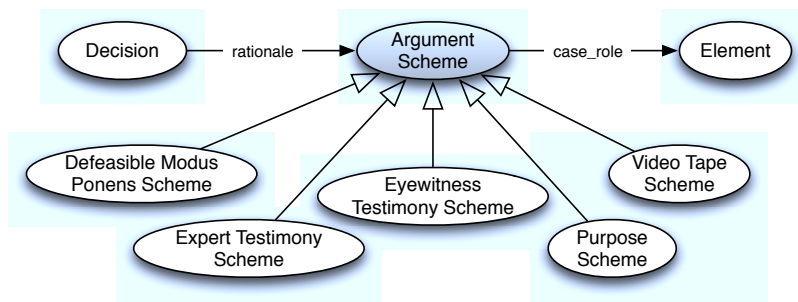


Figure 7 Argument Scheme


would appear as subclasses. The argument schemes are tied to particulars of a case by *case\_role* object properties, which are discussed in the next section on *Elements* and subproperties of *case\_role*. In general, the representation of argument schemes is partial. The schemes are not articulated: various parties may play various roles in the scheme; we leave unspecified who asserts what statement; the defeasible conclusions are not represented. By the same token, the relationships among the schemes are not given such as subsumption relationships (Rahwan and Zablith (2007)). This is a module of the ontology which merits further development.

There are significant issues concerning how much of the semantics of argument schemes can be expressed as OWL DL class restrictions. So as to incorporate them, we only specify unique necessary and sufficient conditions for argument schemes. For instance, the presence of an expert witness testimony indicates that the argument scheme used to provide the rationale for a case was at least of the type *Expert Testimony Scheme*. A further substantive issue is that argument schemes, which may include exception clauses and often express plausible reasoning patterns rather than certain ones, are outside the scope of DL (Gordon et al. (2007b)). Despite these limitations, it is useful to indicate what schemes are used in a case, thereby classifying cases according to the reasoning patterns used to make a determination.

While “standard” argument schemes appear here, we view the argument schemes as the reasoning patterns or structures which are used to reason to a decision. In this light, we could include legal theories, legal rules, and causes of action as argument schemes. For our purposes, we take a simplistic view that a legal theory is another defeasible reasoning pattern, albeit abstract, complex, tied specifically to law, and providing an overarching framework for the decision (though see Bench-Capon and Sartor (2003); Chorley (2007); Ashley (2009)). By the same token, legal rules such as that used to determine *aggravated assault* or other intermediate legal concepts are reasoning patterns used in a case. Causes of action are complex elements, relating to legal theories, legal rules, facts, and remedy. We have not, then, explicitly represented legal theories, legal rules, or causes of action, for they merit further consideration and development.

### 3.6 Element

Figure 8 represents the *Element* class, which includes a range of components that play an important role in the case, and as such, *define* the case. In particular, individuals of the *Element* class relate to a case by an object property *case\_role* which associates them to an *argument scheme*, though this does not preclude other object properties



**Figure 8** Element

from holding. The underlying idea is that things of the *Element* class are inputs to overall legal reasoning process results in the case decision. The *Element* class has several subclasses (see the sections below), each of which have further subclasses; parallel to this classification, we have a range of related object subproperties of *case\_role*. The particular *case\_role* defines the class; this follows the ontology design methodology outlined in section 2.2.

While in principle individuals of the *Element* class relate to the case via the decision and argument scheme, in practice we give but two samples of this below – the relationship between testimony of witnesses and the witness argument schemes. Individuals from other subclasses would be tied to additional argument schemes, though we have not spelled these out; to do so would require the argument schemes to be given in detail, which is beyond the scope of this paper. To give an example, if the argument schemes of case based reasoning (Wyner and Bench-Capon (2007)) were represented here, then *Factor* and *Value* individuals would be tied to those schemes. Similarly, *Fact* or *Legislation* would use schemes where the individuals of these classes are found as premises of argument schemes.

In the following, we discuss each of the subclasses in further detail.

### 3.6.1 Evidence

In a legal proceeding, a range evidence is presented to the court in order to prove or disprove the legal issue such as the legal responsibility of a party or possession of some object. In Figure 9, we have provided a sample classification of such evidence. The type of evidence determines the role it plays in the case. We have not attempted to formally define the necessary and sufficient properties of evidence as the properties may vary from one legal domain to another and the complexity of social reality exceeds the expressiveness of OWL (Hoekstra (2009)).

The subclass *Testimony* includes eyewitness testimony that is given by an *eyewitness* and with respect to the *eyewitness argument scheme*. We have a class *Document* which might be those non-legal documents that are used to support an argument in the case such as telephone records, letters, or other items that are entered into evidence. We have as well a range of other sorts of evidence – facts, physical, circumstantial, etc. Over the course of the legal proceedings, the lawyers (as well as perhaps the presiding judge) make arguments and statements about the evidence and their relation to the cause of action; the lawyers' arguments and statements are not, themselves, evidence. What the lawyers agree on may be taken as members of *Fact*; for example, where a knife is introduced as evidence and the lawyers on both sides of the case agree that the knife is the murder weapon, then it can be taken as a fact that the knife is the murder weapon. A judge may make a determination of what is a fact. Facts in the case decision are, then, a

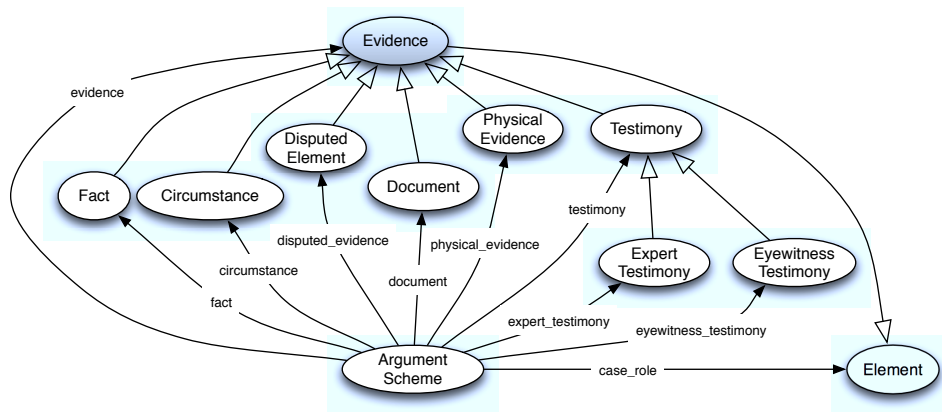


Figure 9 Evidence

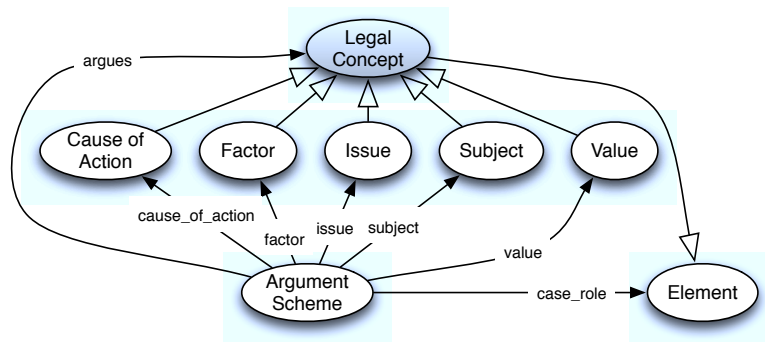


Figure 10 Legal Concept

subspecies of evidence. These are distinct from the evidence which is not agreed on by any party as factual; this evidence need not be consistent, nor need all evidence which is introduced be used to support a judgement. Moreover, given the disputed evidence, not all evidence need be determinate in status (i.e. a member of *Fact* or *Physical Evidence*). Pending further investigation, we presume that the facts are those which either all case participants accept as a fact or which the presiding judge asserts is fact. Similar points can be made concerning *Physical Evidence*, *Document*, and the other classes. It is important to reiterate that our primary objective is to provide an ontology for legal case decisions, taking as given what the legal system has determined is or is not a member of a given class. As implied by our discussion, the determination of what falls into one class or another follows from the argument scheme, though we have not elaborated on this here.

### 3.6.2 Legal\_Concept

The class *Legal\_Concept*, represented in Figure 10, includes a diverse range of elements which are more abstract that describe and qualify a case from a legal perspective among them the cause of action which is cited to bring the case to court, the factors to be taken into consideration in arguing the case, and the legal values that may be relevant to making a determination.

A cause of action is the plaintiff's claim which expresses a legal theory and which justifies bringing the case to court. It must be supported by claims of fact, the legally defined elements of the claim must be met for the claim to succeed, and the legal theory must provide for a remedy. The causes of action are pleaded in the initial complaint. Different areas of law, such as Torts and Contracts, use different causes of action; Torts arise where there are no contractual obligations and cover intentional or accidental actions, where one party holds another party liable for damages. Given our target case, we consider only causes of action related to Torts and our sample case. Because causes of action are particularly complex, we have included them as subclasses of the *Legal\_Concept* class, though this may require some revision.

For the class *Cause\_of\_Action*, we only consider individuals which are raised in *Popov v. Hayashi*: constructive trust (where one person holds property in trust for another), conversion (where one person holds another person's property as her own), injunctive relief (a court-ordered obligation or prohibition on some action, e.g. selling something), and trespass to chattel (where one party has intentionally interfered with possession of property). Each of these bears on property ownership and compensation. In *Popov v. Hayashi*, the judge outlines each of the causes of action, the claims of fact, and other legally defined elements so as to determine whether the cause of action is satisfied and the plaintiff wins the case.

The class *Factor* relates to the factors of legal case based reasoning (Ashley (1990); Alevén and Ashley (1995); Weber et al. (2005); Brüninghaus and Ashley (2005a); Wyner and Bench-Capon (2007)). There are a range of domains where factor analysis has been intensively studied –intellectual property in trade secrets cases (e.g. relating to *Mason v. Jack Daniel Distillery*), property claims (e.g. bearing on *Pierson v. Post* and *Popov v. Hayashi*), or in the distinction between first and second degree murder (Luria (1988)). A notion related to factors are *dimensions* (Rissland and Ashley (2002)), but we leave aside this distinction for the time being. It should be noted that while in some cases, the decision explicitly refers to factors and reasoning about factors, this is not obvious in *Popov v. Hayashi*.

Two aspects of factors should be highlighted. First, the set of factors is said to have a hierarchical organisation (Ashley (1990)), where basic level factors relate to higher level, more abstract factors. To determine whether a given factor holds in a case biases the case decision towards one side or the other of the case. For instance, if parent factor is *Questionable means were used to get the trade secret.*, some of the related subfactors that are organised into a factor hierarchy are *An employee was bribed to get the trade secret?*, *The trade secret was reverse engineered*, *Deception was used to get the trade secret?*. If it is determined that questionable means were used to get the trade secret, then this would bias the decision in favour of the plaintiff. In essence, a more complex, more difficult to determine factor is decomposed into base level factors which are easier to determine. While factors are related in a hierarchical structure, this is not necessarily a structure of strict subsumption so cannot simply be expressed as a class hierarchy. While the factors and the factor hierarchy are central to case based reasoning, we do not articulate them further here for the representation of a *particular* case requires factor individuals (and not the abstract factors, sets of factors, or relationships among factors) that appear in the case. Reasoning with factors as in case based reasoning is outside the scope of the ontology, however, see Wyner (2008), which provides an ontology of factors relative to legal case based reasoning schemes, which is broadly compatible with this ontology.

Though closely related, we treat evidence and factors as distinct for three reasons. First, the introduction of evidence follows its own procedural rules (e.g. admissibility), and evidence appears as particulars of a case. Second, there may as well be conceptual issues as to what is or is not evidence. For example, an obligation to aid a victim of a car accident (relevant in a case concerning a hit-and-run car accident) is not evidence introduced into a case, but a legal “fact” that holds or not relative to the circumstances; it is, though, a factor which is supported by facts that are entered into evidence (e.g. witness testimony). Similarly, whether a driver was driving recklessly is not evidence, but can be supported by evidence such as testimony to support the claim that the car driver failed to heed traffic signs. The factors are crucial in comparing cases and in biasing the judge’s decision to one side or the other. Third, we reason with and organise evidence and factors in distinct ways: the sorts of argument schemes that apply to factors are different from those which apply to evidence (compare the argument schemes in Wyner and Bench-Capon (2007) and Wyner et al. (2007)); factors are organised into factor hierarchies for the purposes of CBR (see Wyner and Bench-Capon (2007) and the literature cited therein); evidence is organised into relationships of “lower level” facts supporting “higher level” facts, e.g. DNA, a finger print, a death by stabbing, and a knife may all be lower level facts that support the higher level fact that one individual stabbed another individual to death. Given these considerations, we have introduced evidence and factors in different parts of the ontology. Further relationships between them ought to be provided. However, as we do not impose class disjunction between them, there is nothing preventing some individuals from being classed as both evidence and factor.

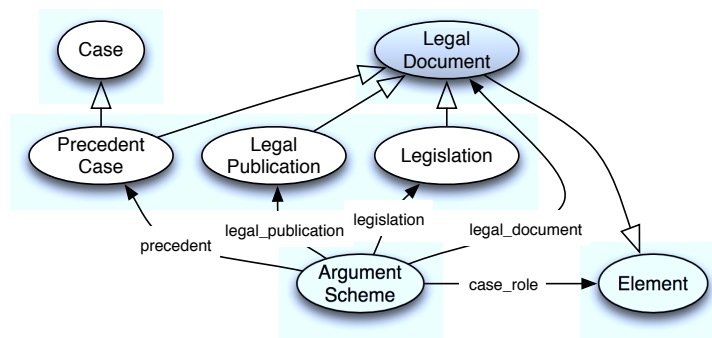
In abstracts of cases, specific legal issues at stake in the case under discussion are usually highlighted. In the case report itself, this might be less explicit. For information extraction and retrieval, it is useful to have a class *Issue* to represent this. A legal issue is some key concept that is central to deciding the case, but is not itself a cause of action or factor. In the instance of *Popov v. Hayashi*, this is *pre-possessory interest in property*, given that Popov had attempted, but failed, to catch the ball. Given that it was determined that Popov did have such an interest, it was determined that he had a legitimate stake in the property that was undermined in the given circumstances. Given this, the causes of action were relevant.

Cases are indexed with respect to the subject area of law they fall under such as employment, family, and so on. We have given a small sample of such members of the *Subject* class.

One line of research in AI and Law gives prominence to legal values such as *fairness* or *the bright line of the law* (Bench-Capon (2002); Bench-Capon and Sartor (2003); Wyner et al. (2007)). In some cases, the legal values are explicitly expressed (e.g. *Popov v. Hayashi*) while in others, the analyst must infer this from the judgement (e.g. cases bearing on hunting wild animals such as *Pierson v. Post* and *Keeble v. Hickeringill*). We suggest that legal values of cases be made explicit. In the ontology, we have a *Value* class and have provided a sample of legal values that a case may have.

### 3.6.3 *Legal\_Document*

In the class of *Legal\_Document* in Figure 11, we have several subclasses – *Precedent\_Case*, *Legal\_Publication*, and *Legislation*. Individuals of these classes are related to an argument scheme, but also *cited\_by* the case, which allows us to identify those documents which are used in the case. This class does not comprise all legal documents,



**Figure 11** Legal Document

nor does it define legal documents independent of cases and argument schemes. In this perspective, some individual can only be a legal document if it is cited by a case (though that is not a sufficient condition).

### 3.7 Summary

We have outlined the structure of the ontology. In the next section, we provide an instantiation of the ontology with respect to *Popov v. Hayashi*.

## 4 Ontological Representation of *Popov v. Hayashi*

In this section, we give elements of the instantiation of *Popov v. Hayashi* with respect to the ontology which we have just presented, giving the members or data associated whether found overtly in the text or presented in the literature on the case. For clarity, we present a selection of the results at the output from ACEView, allowing us to read information which is found in the case and relative to the ontology.<sup>6</sup>

- expert\_testimony argument\_scheme popov\_v\_hayashi.
- eyewitness\_testimony argument\_scheme popov\_v\_hayashi.
- popov\_v\_hayashi\_hearing hearing\_of popov\_v\_hayashi.
- popov\_v\_hayashi address "City and County of San Francisco, Department 306, California, USA".
- popov\_v\_hayashi citation\_index "Popov v. Hayashi, 2002 WL 31833731".
- popov\_v\_hayashi cites pierson\_v\_post.
- popov\_v\_hayashi cites state\_v\_shaw.
- popov\_v\_hayashi\_decision favours alex\_popov.
- popov\_v\_hayashi\_decision favours patrick\_hayashi.
- popov\_v\_hayashi\_defendant\_solicitor jane\_doe.
- popov\_v\_hayashi\_defendant patrick\_hayashi.
- popov\_v\_hayashi\_expert\_witness professor\_bernhardt.
- popov\_v\_hayashi\_fact home\_run\_barry\_bonds.
- popov\_v\_hayashi\_fact home\_run\_catch\_possibility.
- popov\_v\_hayashi\_judge kevin\_mccarthy.
- popov\_v\_hayashi\_legally\_argues constructive\_trust.
- popov\_v\_hayashi\_legally\_argues injunctive\_relief.

<sup>6</sup>We have slightly "cleaned" the output of ACEView, which was not the key objective of our development.

- `popov_v_hayashi` name "Popov versus Hayashi".
- `popov_v_hayashi` plaintiff `alex_popov`.
- `popov_v_hayashi` plaintiff\_solicitor `john_doe`.
- purpose\_argument\_scheme `popov_v_hayashi`.
- video\_tape\_argument\_scheme `popov_v_hayashi`.

A range of queries that can be put to the knowledge base such as:

- What is `popov_v_hayashi`'s `citation_index`?
- What is `popov_v_hayashi`'s `decision`?
- What is `popov_v_hayashi`'s `name`?
- What is `popov_v_hayashi`'s `cause_of_action`?
- What is `popov_v_hayashi`'s `fact`?
- What is `popov_v_hayashi`'s `legal_issue`?
- What is `popov_v_hayashi`'s `plaintiff`?
- Who is `popov_v_hayashi`'s `expert_witness`?

Some inferences from the knowledge base:

- `popov_v_hayashi` benefits `patrick_hayashi`.
- `popov_v_hayashi` benefits `alex_popov`.
- `popov_v_hayashi` eyewitness `tom_doe`.
- `state_v_shaw` is a decided case.

## 5 Discussion and Future Research

Early research identified some key elements for case base representation and reasoning (Hafner (1981, 1987); Dick (1991); Rissland et al. (1996); Costa et al. (1998)). The line of work stemming from Ashley (1990) (see also Rissland et al. (2006)) focuses largely on case based reasoning with respect to factors. This earlier work on knowledge representation did not, by and large, make the ontological structure explicit. In more recent research, ontological analysis has been central. The LOIS project derived a rich legal lexicon from a range of legal documents, however, case representation is not central (Peters et al. (2007); Schweighofer and Liebwald (2007)). In a similar vein, Lame (2004); Saias and Quaresma (2004); Peters (2009) apply a range of natural language processing techniques to corpora of legal texts (usually legislation) in order to identify or enrich legal ontologies; they report varying degrees of success, with much work remaining. The Legal Knowledge Interchange Format (LKIF) takes a broad scope, encompassing abstract elements of legal knowledge and upper level ontologies, but does not include cases (Breuker et al. (2004); Hoekstra et al. (2007, 2009)).

Several recent proposals focus more on legal case ontologies and case based reasoning. Zeng et al. (2005) discusses factors and information retrieval from a case base, however, the ontological structure is not made explicit. Sierra (2008) primarily introduces bibliographic elements. The ontology of Shen et al. (2008) must be inferred from the paper, does not relate to case based reasoning, and does not illustrate the richer range of elements we provide. Wyner (2008) focuses on the factor analysis and comparison which inputs directly into case based reasoning argumentation schemes (Wyner and Bench-Capon (2007)). Finally, we also incorporate research on legal argumentation as in Wyner et al. (2007).

An ontology represents a conceptualisation of a domain. There will be alternative conceptualisations with varying degrees of compatibility. One avenue of future research

is a further detailed comparison and contrast among legal ontologies to determine where they merge or diverge. For example, we have not attempted to work with upper-level ontologies or more abstract legal ontologies. As we outlined in section 2.4, our approach is *modular* in that distinct subcomponents of the ontology can be independently developed. We have noted in the presentation of the abstract ontology various ways in which elements can be added to the ontology; in particular, the structure of argument schemes, the factor hierarchy, legal theories and rules, and the integration with a case based reasoning ontology ought to be further examined. Further reasoning within the case ontology must be examined. No doubt other components that are relevant to case law can be found.

Another avenue of investigation is to instantiate the ontology with further cases to test the design decisions, populate the ontology, and to consider further additions. In general, working with a corpus of instantiated cases, preferably a well-known set of related cases such as that for CATO, would test the current ontology and suggest modifications. A related approach is to tie ontology development and population to natural language processing techniques. While an ontology represents systematic, general, abstract conceptualisations of the domain, it is infeasible to require an ontological representation for all aspects that might be found in a case. For this, information retrieval and extraction techniques would be required to discover individuals, classes, and properties within and among the cases that are not already explicitly coded in the ontology (Peters (2009)). Along these lines, the ontology can be used along with a text annotation system such as the *General Architecture for Text Engineering*.<sup>7</sup> With the annotation system, we would develop a *gold standard* of cases which can either be used to examine reasoning relative to the corpus or to *seed* further text annotation using *machine learning* techniques. Related to our presentation, we will further develop compatibilities between ACEView and other ontology development tools. Finally, it would be worth drawing the attention of the community of legal academics and professionals to check that the ontology represents key aspects of legal knowledge as drawn from practice. In all, much remains to be done in the area of ontological representation of and reasoning with cases.

## References

- Aleven, V., Ashley, K. D., 1995. Doing things with factors. In: ICAIL '95: Proceedings of the 5th International Conference on Artificial Intelligence and Law. ACM, New York, NY, USA, pp. 31–41.
- Antoniou, G., Assmann, U., Baroglio, C., Decker, S., Henze, N., Patranjan, P.-L., Tolksdorf, R. (Eds.), 2007. Reasoning Web, Third International Summer School 2007, Dresden, Germany, September 3-7, 2007, Tutorial Lectures. Vol. 4636 of Lecture Notes in Computer Science. Springer.
- Antoniou, G., van Harmelen, F., 2004. A Semantic Web Primer. The MIT Press.
- Ashley, K., 1990. Modelling Legal Argument: Reasoning with Cases and Hypotheticals. Bradford Books/MIT Press, Cambridge, MA.
- Ashley, K. D., 2009. Ontological requirements for analogical, teleological, and hypothetical legal reasoning. In: ICAIL '09: Proceedings of the 12th International Conference on Artificial Intelligence and Law. ACM, New York, NY, USA, pp. 1–10.

<sup>7</sup>See [www.gate.ac.uk](http://www.gate.ac.uk)

- Bench-Capon, T., Sartor, G., 2003. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence* 150 (1-2), 97–143.
- Bench-Capon, T. J., Visser, P. R., 1996. Deep models, ontologies and legal knowledge based systems. In: *Legal Knowledge Based Systems. JURIX 1996: The Ninth Annual Conference*. Tilburg University Press, pp. 3–14.
- Bench-Capon, T. J. M., 2002. The missing link revisited: The role of teleology in representing legal argument. *Artificial Intelligence and Law* 10 (1-3), 79–94.
- Berman, D., Hafner, C., 1993. Representing teleological structure in case-based legal reasoning: the missing link. In: *ICAAIL '93: Proceedings of the 4th International Conference on Artificial Intelligence and Law*. ACM, New York, NY, USA, pp. 50–59.
- Breuker, J., Valente, A., Winkels, R., 2004. Legal ontologies in knowledge engineering and information management. *Artificial Intelligence and Law* 12 (4), 241–277.
- Brüninghaus, S., Ashley, K., 2005a. Reasoning with textual cases. In: Munzo-Avila, H., Ricci, F. (Eds.), *Proceedings of the International Conference on Case-based Reasoning 2005*. No. 3620 in LNAI. Springer-Verlag, pp. 137–151.
- Brüninghaus, S., Ashley, K. D., 1997. Finding factors: learning to classify case opinions under abstract fact categories. In: *ICAAIL '97: Proceedings of the 6th International Conference on Artificial Intelligence and Law*. ACM, New York, NY, USA, pp. 123–131.
- Brüninghaus, S., Ashley, K. D., 2005b. Generating legal arguments and predictions from case texts. In: *ICAAIL 2005*. ACM Press, New York, NY, USA, pp. 65–74.
- Chorley, A., 2007. Reasoning with legal cases seen as theory construction. Ph.D. thesis, University of Liverpool, Department of Computer Science, Liverpool, UK.
- Coi, J. L. D., Fuchs, N. E., Kaljurand, K., Kuhn, T., 2009. Controlled english for reasoning on the semantic web. In: *REWERSE*. Springer, pp. 276–308.
- Costa, M., Sousa, O., Neves, J., 1998. An architecture to legal distributed case representation. In: Hage, J., Bench-Capon, T., Koers, A., de Vey Mestdagh, C., Grtters., C. (Eds.), *Legal Knowledge Based Systems: JURIX: The Eleventh Conference*.
- Daniels, J. J., Rissland, E. L., 1997. Finding legally relevant passages in case opinions. In: *ICAAIL '97: Proceedings of the 6th International Conference on Artificial intelligence and Law*. ACM, New York, NY, USA, pp. 39–46.
- Dick, J., June 1991. Representation of legal text for conceptual retrieval. In: *ICAAIL'91: Proceedings of the 3rd International Conference on Artificial Intelligence and Law*. Oxford, pp. 244–252.
- Gangemi, A., 2007. *Trends in Legal Knowledge: The Semantic Web and the Regulation of Electronic Social Systems*. European Press Academic Publishing, Ch. Design Patterns for Legal Ontology Construction.
- Gangemi, A., Sagri, M., Tiscornia, D., 2005. A constructive framework for legal ontologies. In: Benjamins, V., Casanovas, P., Breuker, J., Gangemi, A. (Eds.), *Law and the Semantic Web*. Springer Verlag, pp. 97–124.

- Gordon, T., Prakken, H., Walton, D., 2007a. The carneades model of argument and burden of proof. *Artificial Intelligence* 171, 875–896.
- Gordon, T. F., 1993. The pleadings game: formalizing procedural justice. In: *ICAAIL '93: Proceedings of the 4th International Conference on Artificial Intelligence and Law*. ACM Press, Amsterdam, The Netherlands, pp. 10–19.
- Gordon, T. F., Prakken, H., Walton, D., 2007b. The carneades model of argument and burden of proof. *Artificial Intelligence* 171 (10-15), 875–896.
- Gruber, T. R., 1993. A translation approach to portable ontology specifications. *Knowledge Acquisition* 5 (2), 199–220.
- Guarino, N., Welty, C., 2002. Evaluating ontological decisions with OntoClean. *Communications of the ACM* 45(2), 61–65.
- Guarino, N., Welty, C. A., 2004. An Overview of OntoClean. In: Staab, S., Studer, R. (Eds.), *Handbook on Ontologies*. Springer Verlag, Ch. 8.
- Hafner, C., 1987. Conceptual organization of case law knowledge bases. In: *ICAAIL '87: Proceedings of the 1st International Conference on Artificial Intelligence and Law*. ACM, New York, NY, USA, pp. 35–42.
- Hafner, C. D., 1981. Representation of knowledge in a legal information retrieval system. In: *SIGIR '80: Proceedings of the 3rd annual ACM conference on Research and development in information retrieval*. Butterworth & Co., Kent, UK, UK, pp. 139–153.
- Hoekstra, R., June 2009. *Ontology Representation – Design Patterns and Ontologies that Make Sense*. Vol. 197 of *Frontiers of Artificial Intelligence and Applications*. IOS Press, Amsterdam.
- Hoekstra, R., Breuker, J., Bello, M. D., Boer, A., June 2007. The LKIF core ontology of basic legal concepts. In: *Legal Ontologies and Artificial Intelligence Techniques*. Stanford University, Palo Alto, CA, USA.
- Hoekstra, R., Breuker, J., Bello, M. D., Boer, A., 2009. Lkif core: Principled ontology development for the legal domain. In: Breuker, J., Casanovas, P., Klein, M. C. A., Francesconi, E. (Eds.), *Law, Ontologies and the Semantic Web*. Vol. 188 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, pp. 21–52.
- Jackson, P., Al-Kofahi, K., Tyrell, A., Vachher, A., November 2003. Information extraction from case law and retrieval of prior cases. *Artificial Intelligence* 150 (1-2), 239–290.
- Kralingen, R. W. V., Visser, P. R. S., Bench-Capon, T. J. M., Herik, H. J. V. D., 1999. A principled approach to developing legal knowledge systems. *International Journal of Human-Computer Studies* 51, 1127–1154.
- Lame, G., 2004. Using nlp techniques to identify legal ontology components: Concepts and relations. *Artificial Intelligence and Law* 12 (4), 379–396.
- Luria, D., 1988. Death on the highway: Reckless driving as murder. *Oregon Law Review* 799, 821–22.

- Maynard, D., Li, Y., Peters, W., 2008. NLP techniques for term extraction and ontology population. In: *Proceeding of the 2008 conference on Ontology Learning and Population: Bridging the Gap between Text and Knowledge*. IOS Press, Amsterdam, The Netherlands, The Netherlands, pp. 107–127.
- Moens, M.-F., Boiy, E., Mochales-Palau, R., Reed, C., 2007. Automatic detection of arguments in legal texts. In: *ICAAIL '07: Proceedings of the 11th International Conference on Artificial Intelligence and Law*. ACM Press, New York, NY, USA, pp. 225–230.
- Motik, B., Patel-Schneider, P., Parsia, B., Bock, C., Fokoue, A., Haase, P., Hoekstra, R., Horrocks, I., Ruttenberg, A., Sattler, U., Smith, M., 11th June 2009. OWL 2 web ontology language structural specification and functional-style syntax. Tech. rep., World Wide Web Consortium, candidate Recommendation.
- Peters, W., 2009. Text-based legal ontology enrichment. In: *Proceedings of the workshop on Legal Ontologies and AI Techniques*. Barcelona, Spain.
- Peters, W., Sagri, M.-T., Tiscornia, D., 2007. The structuring of legal knowledge in LOIS. *Artificial Intelligence and Law* 15 (2), 117–135.
- Rahwan, Iyad, C. R., Zablith, F., 2007. On building argumentation schemes using the argument interchange format. In: *Working Notes of the 7th Workshop on Computational Models of Natural Argument (CMNA 2007)*, Hyderabad.
- Rissland, E. L., Ashley, K. D., 2002. A note on dimensions and factors. *Artificial Intelligence and Law* 10 (1-3), 65–77.
- Rissland, E. L., Ashley, K. D., Branting, L. K., 2006. Case-based reasoning and law. *The Knowledge Engineering Review* 20, 293–298.
- Rissland, E. L., Skalak, D. B., Friedman, M. T., 1996. BankXX: Supporting legal arguments through heuristic retrieval. *Artificial Intelligence and Law* 4 (1), 1–71.
- Saias, J., Quaresma, P., 2004. A methodology to create legal ontologies in a logic programming based web information retrieval system. *Artificial Intelligence and Law* 12 (4), 397–417.
- Sartor, G., 2006. Fundamental legal concepts: a formal and teleological characterisation. *Artificial Intelligence and Law* 14 (1), 101–142.
- Schweighofer, E., 1999. The revolution in legal information retrieval or: The empire strikes bac. *The Journal of Information, Law and Technology* 1, online.
- Schweighofer, E., Liebwald, D., 2007. Advanced lexical ontologies and hybrid knowledge based systems: First steps to a dynamic legal electronic commentary. *Artificial Intelligent and Law* 15 (2), 103–115.
- Schwiter, R., Kaljurand, K., Cregan, A., Dolbear, C., Hart, G., 1–2 April 2008. A Comparison of three Controlled Natural Languages for OWL 1.1. In: *4th OWL Experiences and Directions Workshop (OWLED 2008 DC)*. Washington.

- Shen, Y., Steele, R., Murphy, J., 2008. Building a semantically rich legal case repository in owl. In: Richardson, J., Ellis, A. (Eds.), Proceedings of AusWeb08, The Fourteenth Australian World Wide Web Conference. Southern Cross University, Lismore, New South Wales, Australia, pp. 97–108.
- Sierra, S., 2008. Owl case features, <http://github.com/lawcommons/altlaw-vocab/tree/master>.
- Solan, L. M., 2003. Cognitive foundations of the impulse to blame. *Brooklyn Law Review* 68, 1003–1029.
- Solan, L. M., 2005. Language and law: Definitions in law. In: *Encyclopedia of Language and Linguistics* (2nd ed). Elsevier.
- Solan, L. M., Darley, J., 2001. Causation, contribution and legal liability: An empirical study. *Law & Contemporary Problems* 64 (4), 265–298.
- Sporleder, C., Lascarides, A., 2006. Using automatically labelled examples to classify rhetorical relations: An assessment. *Natural Language Engineering* 14 (3), 369–416.
- Uschold, M., Gruninger, M., 1996. Ontologies: Principles, methods and applications. *Knowledge Engineering Review* 11 (2), 93–155.
- Valente, A., 2005. Types and roles of legal ontologies. In: Benjamins, V., Casanovas, P., Breuker, J., Gangemi, A. (Eds.), *Law and the Semantic Web*. Vol. 3369 of *Lecture Notes in Artificial Intelligence*. Springer, pp. 65–76.
- Walton, D., 2002. *Legal Argumentation and Evidence*. The Pennsylvania State University Press.
- Weber, R. O., Ashley, K. D., Brüninghaus, S., 2005. Textual case-based reasoning. *Knowledge Engineering Review* 20 (3), 255–260.
- Wyner, A., December 2008. An ontology in OWL for legal case-based reasoning. *Artificial Intelligence and Law* 16 (4), 361–387.
- Wyner, A., 2009. An OWL ontology for legal cases with an instantiation of *Popov v. Hayashi*. In: Atkinson, K. (Ed.), *Proceedings of the ICAIL 2009 Workshop on Modeling Legal Cases*. IDT Series. Huygens Editorial, Barcelona, Spain, to Appear.
- Wyner, A., Bench-Capon, T., 2007. Argument schemes for legal case-based reasoning. In: Lodder, A. R., Mommers, L. (Eds.), *Legal Knowledge and Information Systems. JURIX 2007*. IOS Press, Amsterdam, pp. 139–149.
- Wyner, A., Bench-Capon, T., Atkinson, K., 2007. Arguments, values and baseballs: Representation of *popov v. hayashi*. In: Lodder, A. R., Mommers, L. (Eds.), *Legal Knowledge and Information Systems. JURIX 2007*. IOS Press, Amsterdam, pp. 151–160.
- Zeng, Y., Wang, R., Zeleznikow, J., Kemp, E. A., 2005. Knowledge representation for the intelligent legal case retrieval. In: *Knowledge-Based Intelligent Information and Engineering Systems*. pp. 339–345.