

MPhil Thesis: Computational models of Ontology Evolution in Legal Reasoning

Andrew Priddle-Higson



Master of Philosophy

Centre for Intelligent Systems and their Applications

School of Informatics

University of Edinburgh

2010

Abstract

This thesis analyses the problem of creating computational models of ontology evolution in legal reasoning. Ontology evolution is the process of change that happens to a theory as it is used by agents within a domain. In the legal domain these theories are the laws that define acceptable behaviours and the meta-legal theories that govern the application of the laws. We survey the background subjects required to understand the problem and the relevant literature within AI and Law. We argue that context and commonsense are necessary features of a model of ontology evolution in legal reasoning; and propose a model of legal reasoning based upon creating a discourse context. We conclude by arguing that there is a distinction between prescriptive and descriptive models of ontology evolution; with a prescriptive model being a social and philosophical problem, rather than a technical one, and a descriptive model being an AI-complete problem.

Acknowledgements

First, thank you to my supervisors Alan Bundy, Fiona McNeill and Burkhard Schafer for their time and advice. In particular, thank you to Burkhard for his guidance through the field of Law.

I would also like to thank CISA and the Informatics Graduate School for the opportunity to pursue this research. CISA is a wonderful institute filled with intelligent, friendly people, whom I am glad to have worked with.

Finally, thank you to my family for their support.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Andrew Priddle-Higson)

Table of Contents

1	Introduction	1
1.1	Ontology evolution	1
1.2	Legal Reasoning	2
1.3	Thesis	3
1.4	Chapter Overview	4
2	Subject Survey	6
2.1	Ontology	6
2.1.1	Formal Ontology	6
2.1.2	Formalisms	8
2.1.3	Example Ontologies	13
2.1.4	Interoperability	16
2.2	Ontology Evolution	17
2.2.1	Tools	18
2.2.2	ORS	19
2.2.3	Ontology Evolution in Physics	20
2.3	Law	22
2.3.1	Legal Systems	22
2.3.2	Legal Rules	24
2.3.3	Legal Questions	24
2.3.4	Legal Ontologies	25
2.3.5	Legal Reasoning	27
2.3.6	Evolution of the Law	31
2.3.7	Modelling Ontology Evolution in Legal Reasoning	31
2.4	Summary	33

3	Ontology Evolution in Legal Cases	34
3.1	Overview of a Legal Case	34
3.1.1	Ontology Evolution in Legal cases	35
3.2	Some Example Legal Cases	36
3.2.1	Popov v. Hayashi	36
3.2.2	Buchanan v. Babco	41
3.2.3	Corocraft v. Pan American	51
3.2.4	Fothergill v. Monarch Airlines	55
3.3	Features of Legal Reasoning	60
3.4	Summary	62
4	Meaning in Legal Language	63
4.1	Law and Language	63
4.1.1	Ambiguity	64
4.1.2	Vagueness	65
4.1.3	Open texture	66
4.1.4	Context	68
4.1.5	Legal Interpretation	68
4.1.6	Legal indeterminacy	70
4.2	Philosophy of Language	71
4.2.1	Logical Positivism	71
4.2.2	Language Games	72
4.2.3	Meaning	73
4.3	Natural Language Processing	74
4.3.1	Discourse Interpretation	74
4.4	Meaning in Ontologies	75
4.4.1	Authority	76
4.5	Arguments about Meaning	77
4.6	Summary	79
5	Legal Case Modelling in AI and Law	80
5.1	Non-monotonic reasoning	80
5.2	Logic Programming	81
5.3	Argumentation	82
5.4	Argument Schemes	84
5.5	Case-based reasoning	85

5.6	Theory Construction	87
5.7	Coherent reasoning	88
5.8	Meta-level reasoning	91
5.9	Limitations of these techniques	92
5.10	Summary	94
6	Arguments about Meaning in Legal Cases	96
6.1	Arguments about Meaning	96
6.1.1	Ontologies of meaning	97
6.1.2	Authority over meaning	98
6.2	Context	99
6.2.1	Contextual Logics	100
6.2.2	Contexts in Legal reasoning	103
6.2.3	Limitations of existing contextual logics	105
6.2.4	Argumentation and Contextual logics	106
6.3	Commonsense	106
6.3.1	The role of Commonsense in Legal Reasoning	107
6.3.2	Commonsense Knowledge Bases	108
6.3.3	Limitations of current Commonsense Knowledge-Bases	110
6.4	Discourse Context	112
6.4.1	An example implementation of this model	114
6.4.2	Plaintiff	116
6.4.3	Defendant	117
6.4.4	Limitations of this model	119
6.5	Summary	120
7	Requirements on a Theory of Ontology Evolution in Law	122
7.1	Ontology Evolution in Legal Reasoning	122
7.1.1	Arguments about Meaning	123
7.1.2	Contextual knowledge	124
7.1.3	Discourse Context	125
7.2	Types of theories	126
7.2.1	Prescriptive theories	126
7.2.2	Applicability of a Prescriptive theory	127
7.2.3	Descriptive	128
7.2.4	Problems for a Descriptive theory	129

7.3	Developing Applications	130
7.3.1	Kinds of application	131
7.3.2	Developing ontologies	131
7.3.3	Developing reasoning mechanisms	132
7.3.4	Usefulness of Applications	132
7.4	Further Work	134
7.4.1	Ontologies about Meaning	134
7.4.2	Domain-specific examples	135
7.4.3	Contextual logics	135
7.4.4	Interactive tools	135
7.4.5	Commonsense knowledge	135
7.5	Conclusion	136

Bibliography	137
---------------------	------------

Chapter 1

Introduction

The knowledge we have of the world is not static. We continually encounter new situations that don't fit within our established theories. We can reflect upon and spot contradictions in our intuitions. We then need to adapt our theories to resolve these problems, so that we can continue to use them to plan our interactions with the world.

Computer software is based upon a representation of a domain. A programmer will choose a representation that enables them to specify an algorithm on this representation which produces a useful behaviour. When the domain changes the software must be manually changed.

Researchers in Informatics have recently started to investigate ontologies, models of a domain, as a way to define domain knowledge which can be used by generic problem solvers to produce useful behaviour within the domain. However, when the domain changes the ontology must also be manually changed.

We would like to have automated methods to reduce this maintenance burden. But how can we automate this process of adapting a representation to its domain?

1.1 Ontology evolution

This is the problem of ontology evolution:

How can we create automated methods to evolve a model of a domain?

We want the new model to be a *better* representation of the domain than the original one. The original model had a problem: for example, it gave an incorrect prediction or it gave conflicting predictions. We want to replace the original model with one which

doesn't have this problem. But there may be several different possible new models. How do we find these new models and decide amongst them?

The challenge for research in ontology evolution is to create theories about how ontologies evolve. In order to investigate ontology evolution we need case studies: examples of situations in which ontologies have been evolved. We want to use case studies from a domain to investigate how we can build computational models of the ontology evolution in that domain. We survey existing ontology evolution research in section 2.2.

We have been investigating ontology evolution within legal cases. There are many historical cases where the ontology of the Law has changed. We have studied some of these cases to explore how computational models of these changes could be constructed and how useful they would be. Further discussion of the legal domain is in section 2.3.

1.2 Legal Reasoning

The relevant problem in Law¹ is to interpret the ontologies that are used in legislation. These ontologies are frequently incomplete, they can't be clearly applied to some situations. In these problem situations we can't determine whether some object in the domain should be included in the definition of a term in the ontology.

The problem is that there can be conflicting intuitions about how laws should be interpreted. The challenge is to find an interpretation of the law. However, the different parties in a legal case will have different aims when interpreting the law. A lawyer will want to find an interpretation that aids their client. A judge should aim to find an interpretation which uses the methods sanctioned by the legal system to resolve the case.

Consider the following analogous situation. You are walking in the woods with a friend when they suddenly point at a tree stump and exclaim "Great! A chair". You interpret "chair" as referring to the tree stump. Now consider another situation in which you ask your friend to buy a "chair" for your house and they return with a tree stump. In the second situation, was your friend's interpretation of "chair" correct?

In the first situation your friend was using the term "chair" because a tree stump

¹In this thesis, we will generally use the upper-case "Law" to refer to the subject, which studies legal systems, and lower-case "law" to refer to legal rules. So "English Law" refers to the English legal system and "an English law" refers to a rule in English Law.

had the same functional properties that were relevant to your current (shared) context of the forest. When you interpret your friend's remark, you adapt your definition of "chair" to include the tree stump in that particular context. In the second situation you used the term "chair" to refer to an object which would be considered a "chair" in the context of your home. The meaning of "chair" can vary between these situations, and shouldn't refer to a tree stump in the context of your home.

Your friend could argue that since "chair" could refer to a chair in the forest then why shouldn't it refer to a chair in your home. They could further argue that it is fashionable to use a tree stump as a chair, it might give the house an eco-home ambiance. Your friend has adapted their representation of "chair" to include tree stumps in the context of a house. You might counter their argument by arguing that it is an unusual interpretation of what a "chair" is, and so they should have recognised that you probably didn't intend this meaning. These arguments give conflicting interpretations of what the meaning of "chair" is in this context.

In this example the problem of interpreting "chair" is to determine what you meant by "chair" in the second situation. This problem is not too difficult to solve, since you presumably know what you meant. In the legal domain the problem is analogous, but significantly more difficult since there isn't an oracle that can define a term. The background context of a case provides some information about how a law can be interpreted, but this background context doesn't always give a definitive, clear meaning for the law.

We will survey the relevant aspects of Law in section 2.3. We are specifically looking at ontology evolution in legal cases. We give an analysis of legal cases in section 3.1.

1.3 Thesis

We will argue for the following claims in this report:

1. Ontology evolution in legal cases occurs as a side effect of arguments about the *meaning* of legal rules and principles.
2. There is a distinction between prescriptive and descriptive theories of ontology evolution in legal reasoning.
3. A complete descriptive theory of ontology evolution in legal reasoning is an AI-complete problem.

4. A combination of the features of the different theories could provide the basis for a useful system for aiding the construction of arguments about the meaning of legal ontologies. However, any system would require significant human interaction to produce useful output.

Claim one is the main claim of this report. Ontology evolution in legal cases results from a particular process of legal reasoning: reasoning about meaning.

Claim two is about the types of theory of ontology evolution that we can create. We argue that there are two different possible types of theory: prescriptive and descriptive.

Claim three is a claim about the possibility of developing a descriptive theory of ontology evolution in legal reasoning.

Claim four is a claim about what types of system can be developed now, and how useful these systems would be.

1.4 Chapter Overview

Chapter 2 gives a survey of research in ontologies and ontology evolution, and provides the necessary background in Law to understand our research. This chapter can be skipped, or selectively read, by those with the necessary background.

In chapter 3 we give an overview of the Law and some example legal-cases. This provides an analysis of the phenomena we will be trying to model: the argumentation about the meaning of a legal ontology that occurs in a legal case. We also look at the problems of working on ontology evolution in legal cases.

In chapter 4 we discuss the problems of ascribing meaning to legal language, and how these problems affect our attempts to model ontology evolution.

Chapter 5 gives an overview of the research in AI and Law on legal case modelling. There has been some research trying to model the argumentation in legal cases, some of this is relevant to our research. We also discuss the limitations of current techniques in AI and Law for modelling ontology evolution in the Law.

Chapter 6 introduces the techniques we have been trying to use to model the ontology evolution in legal reasoning, and our attempts to apply these techniques to our problem.

Chapter 7 is the culmination of the thesis. In this chapter we discuss the requirements on a theory of ontology evolution in legal reasoning. We distinguish between the different types of possible theory, summarise our arguments for the claims introduced

in section 1.3, and describe further work.

Chapter 2

Subject Survey

The previous chapter gave an informal introduction to the problem of ontology evolution. In this chapter we survey ontology and ontology evolution research, and cover the necessary background on Law to understand our research.

2.1 Ontology

Ontology is the study of conceptualisations, the collections of concepts we use to form representations of our world. It began with the work of philosophers such as Aristotle¹, who tried to form sets of fundamental concepts that could describe all the things that existed. The subject has since been developed by generations of philosophers.

The original concern of the philosophical study of ontology was to discover a single ontology which described everything in the world. Philosophers studied these *upper ontologies* with the aim of discovering the ultimate nature of reality.

We are interested in research into ontologies within Computer Science, where researchers have adapted the subject to aid the development of intelligent computer systems. The concerns of Computer Scientists are more pragmatic than Philosophers: the aim of Computer Scientists is to develop useful knowledge-based systems.

2.1.1 Formal Ontology

Computer Scientists recently became concerned with ontologies whilst trying to build knowledge-based systems, systems which use explicitly declared domain-knowledge to solve problems within a domain. Knowledge-based systems require ontologies to

¹such as his work on *Metaphysics*

describe the application domain. For instance, a financial expert system would require ontologies of credit ratings and consumer habits to define rules about who should be given credit.

The standard definition of an ontology for computer science was given by Tom Gruber as “An explicit specification of a conceptualisation”[52], where a conceptualisation is “an abstract simplified view of the world that we want to represent for some purpose”.

The emphasis is on the idea of making the conceptualisation explicit. As we mentioned in the previous chapter, every program is based upon a representation of the domain it operates in. These representations are frequently implicit within the program, and only discussed in documentation regarding the design of the program. The motivation behind ontologies research is to make these conceptualisations explicit within a representation language.

Gruber’s article goes on to discuss the need for these ontologies to be shared, so that components within a distributed knowledge-based system can share knowledge.

This emphasis upon a shared conceptualisation led to an expanded definition, given in [133]:

“An ontology is a formal, explicit specification of a shared conceptualisation. A ‘conceptualisation’ refers to an abstract model of some phenomenon in the world by having identified the relevant concepts of that phenomenon. ‘Explicit’ means that the type of concepts used, and the constraints on their use are explicitly defined. For example, in medical domains, the concepts are diseases and symptoms, the relations between them are causal and a constraint is that a disease cannot cause itself. ‘Formal’ refers to the fact that the ontology should be machine readable, which excludes natural language. ‘Shared’ reflects the notion that an ontology captures consensual knowledge, that is, it is not private to some individual, but accepted by a group.”

For example, a medical-informatics system might have information about patients stored at different hospitals. The system would need to have protocols for establishing connections between hospital systems. These connections would then need to share information formalised within a common logic, so that the meaning of the expressions passed between the hospitals is the same. The system should now also require that the ontologies used by the hospitals are the same, so that the expressions passed between the hospitals use a common language that is mutually understood.

Without a shared ontology messages might not be mutually understood. For instance, one of the hospitals might send an expression such as “requires(patient12312,drug12321)”

to another. If the other hospital used “needs” rather than “requires” to express a relation between patients and medications, the message would not be understood by the receiving hospital. The receiving hospital wouldn’t understand that “requires” means the same thing as “needs” and so the patient might not get the medication they require.

This was the initial ideal behind ontologies research: we could design distributed systems that could effectively share knowledge by using a common ontology. We shall look at some of the problems with this ideal below (see section 2.1.4).

It should be noted that the term “ontology” was used prior to Gruber’s work to describe various research projects in Artificial Intelligence. For instance, Pat Hayes was attempting to create an ontology of naive physics back in the late 1970s [62]. This notion of ontology is slightly closer to the sense in Philosophy, since it aims to create a fundamental conceptualisation of the world for an AI agent. The more recent sense is about creating ontologies to enable practical knowledge-based systems to function within a specific domain.

The term “formal ontology” is also used to describe modern reincarnations of the Philosophical project, for example the work of Philosophers such as Barry Smith [120] and Nic Guarino at LFO[53]. These projects are mostly concerned with developing philosophical and logical foundations for ontology and building formal upper-ontologies. We are not concerned with these projects, although we shall consider related themes later in this thesis.

2.1.2 Formalisms

Formal ontologies require mathematical logics. A variety of logics have been investigated in the field of Knowledge Representation and Reasoning, a subfield of Artificial Intelligence.

A useful abstract definition, taken from [66], is that an ontology, O , consists of a pair $O = (S, A)$ where S is a signature, the non-logical terms used to express the ontology, and A is a set of axioms within a logic (which use the signature).

This high-level definition gives us a representation of ontology that makes it easy to understand abstract problems relating to ontology. However, if you want to build a knowledge-based system you must use a particular logic. We shall look at description logics, first-order logic, higher-order logic and RDF as knowledge representation languages capable of representing ontologies. Naturally our discussion of the formalisms is quite cursory, readers seeking more rigorous formulations should consult the refer-

ences.

2.1.2.1 Description Logics

Description logics form the core of most research into ontologies, due to the description logic OWL² being promoted as a W3C³ standard. The W3C is trying to create the Semantic Web [13], an extension of the current World Wide Web from being about hypertext documents to being about things and their connections [140].

The Semantic Web requires a formal logic to represent the meta-data in a language that computers can understand. Description Logics were chosen as a suitable formalism for the purposes of Semantic Web.

Description logics are subsets of first-order logic which allow unary relations, called concepts, and binary relations, called roles. The logic is restricted to asserting relations of inclusion and equivalence between concepts, with roles being used to define auxiliary concepts. The reason for these restrictions is that they have nice theoretical properties, this subset of first-order logic is decidable. In practise, it is not clear what the benefit of this restriction is, since the computational complexity of the decision procedures of standard description logics is PSPACE-complete⁴ [36].

Description logics also distinguish between a T-Box, which contains terminological assertions, and an A-Box, which contains assertions about particular individuals. For instance, an assertion that “All Cats are Animals” would belong to the T-Box and an assertion that “Tiddles is a Cat” would belong to the A-Box.

OWL is the merger of the DAML and OIL projects. The language is at the top of the Semantic Web stack (see figure 2.1), the technology stack proposed by the W3C as a basis for the Semantic Web. OWL syntax is based upon XML and RDF, an example of an OWL ontology is shown in figure 2.1.

The formal semantics of OWL is based upon model theory, the same as first-order logic which we discuss in the next section.

2.1.2.2 First-order Logic

First-order logic is the original mathematical-logic, independently created by Gottlob Frege [43] and Charles Peirce [96]. The logic extends propositional logic with objects,

²Web Ontology Language

³The W3C is an organisation founded to guide the development of the World Wide Web. (<http://www.w3.org>)

⁴The PSPACE complexity class contains the class of NP problems

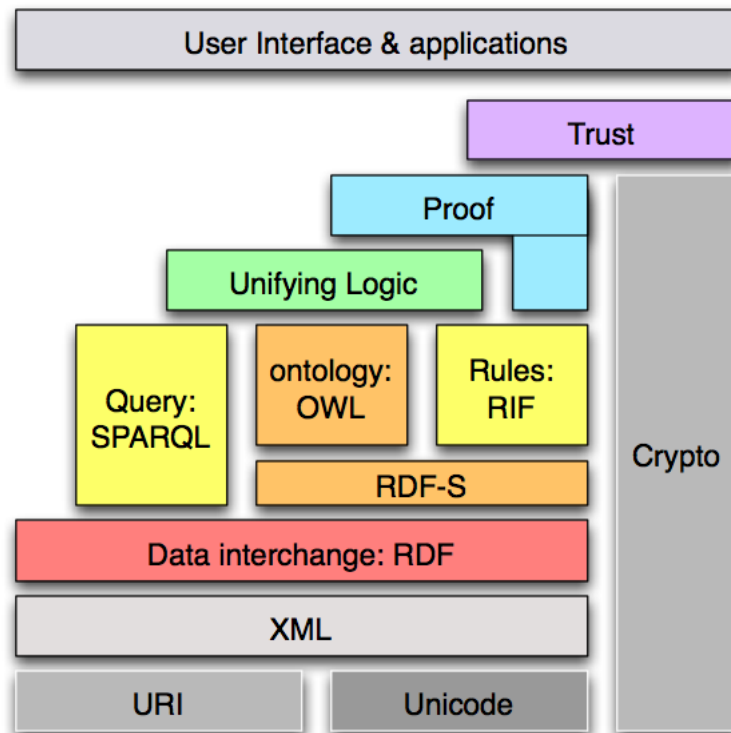


Figure 2.1: The Semantic Web technology stack. OWL is at the top of the ontology section of the stack, above XML and RDF. The stack also includes concepts, such as Trust, which are regarded as being critical to the success of the Semantic Web.

and relations and functions over those objects. The objects form a *domain of discourse* which the propositions are about.

Propositions in first-order logic are formed according to the syntax shown below (quoted from [95]). The main distinction with description logic is the addition of boolean connectives and n-ary relations and functions.

FOL Syntax:

$$\begin{aligned}
\textit{Sentence} &\rightarrow \textit{AtomicSentence} \\
&| (\textit{Sentence} \textit{Connective} \textit{Sentence}) \\
&| \textit{Quantifier} \textit{Variable}, \dots \textit{Sentence} \\
&| \neg \textit{Sentence} \\
\textit{AtomicSentence} &\rightarrow \textit{Predicate}(\textit{Term} \dots) \mid \textit{Term} = \textit{Term} \\
\textit{Term} &\rightarrow \textit{Function}(\textit{Term} \dots) \\
&| \textit{Constant} \\
&| \textit{Variable} \\
\textit{Connective} &\rightarrow \Rightarrow \mid \wedge \mid \vee \mid \Leftrightarrow \\
\textit{Quantifier} &\rightarrow \forall \mid \exists
\end{aligned}$$

The syntax of first-order logic with equality, specified in Backus-Naur form.

Example: Consider the following example theory, with an ontology as follows:

$$\begin{aligned}
\textit{Variable} &\rightarrow X \mid X' \mid X'' \mid \dots \\
\textit{Predicate} &\rightarrow \textit{likes} \mid \textit{person} \\
\textit{Function} &\rightarrow \textit{father} \\
\textit{Constant} &\rightarrow \textit{john} \mid \textit{bananas} \mid \textit{bill}
\end{aligned}$$

where the predicates have their intuitive natural-language meaning, and the following axioms:

$$\begin{aligned}
\forall X. \textit{person}(X) &\Rightarrow \exists X'. \textit{Father}(X) = X' \Rightarrow (\forall X''. \textit{likes}(X, X'') \Rightarrow \textit{likes}(X', X'')) \\
&\textit{likes}(\textit{john}, \textit{bananas}) \\
&\textit{person}(\textit{john}) \\
&\textit{father}(\textit{john}) = \textit{bill}
\end{aligned}$$

These axioms have the consequence that $\textit{likes}(\textit{bill}, \textit{bananas})$. Note the choices that we have made in presenting this theory. Bananas are represented as a constant, rather than as a concept in their own right. This is exemplary of the difficulties of knowledge engineering.

First-order logic provides a logic expressive enough to act as a meta-logic for any other logic [124]. Indeed, several AI researchers have proposed first-order logic as being an adequate foundation for Knowledge Representation in Artificial Intelligence [46] [124] [75]. However, there have been critiques of this position [86] [14].

First-order logic has a Tarskian semantics [134], where the meaning of an expression in first-order logic is based upon the models of that expression. A model of an expression is an interpretation of its syntax in terms of the domain of discourse, where the domain of discourse is a mathematical set of objects. We shall consider this notion of meaning further in section 4.4.

We shall use first-order logic to formalise any examples in this thesis as it is adequate for our purposes.

2.1.2.3 Higher-Order Logic

We mention higher-order logic as there have been some uses of this logic to formalise knowledge, in particular the Ontology Evolution in Physics research we discuss in section 2.2.3.

First-order logic restricts quantifiers to objects in the domain of discourse. However, some expressions are naturally formulated with quantification over relations. For instance, the principle of induction over the natural numbers states that any property which holds for zero and which if it holds for n then it also holds for $n + 1$, therefore holds for all natural numbers. In first-order logic we could only express this property for a particular predicate, whereas the principle should be applicable to all predicates. We need to use second-order logic⁵, which allows quantification over predicates and functions, to formalise this principle.

Higher-order logic extends the quantification over functions and relations to quantification over higher-order functions and relations. This enables expressions that can't be directly stated in first-order logic, but it also creates problems (such as Goedel's incompleteness theorems [93]).

2.1.2.4 RDF

We have just surveyed some quite expressive logics with rigorous meta-logical theories. However, many of the ontologies in widespread use are based upon lightweight representations, such as RDF.

⁵Alternatively, an infinite number of first-order axioms can be used.

RDF is a simple subject-verb-object triple representation formalised in an XML schema. Although RDF is not a mathematical logic - there is no proof theory for RDF - it does have a model-theoretic semantics⁶.

RDF ontologies can be naturally viewed as graphs of RDF resources connected by edges (which represent the triples). An example of an RDF ontology in XML format is shown below ⁷, and its graph representation is in figure 2.2.

```
<?xml version="1.0"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
        xmlns:contact="http://www.w3.org/2000/10/swap/pim/contact#">

  <contact:Person rdf:about="http://www.w3.org/People/EM/contact#me">
    <contact:fullName>Eric Miller</contact:fullName>
    <contact:mailbox rdf:resource="mailto:em@w3.org"/>
    <contact:personalTitle>Dr.</contact:personalTitle>
  </contact:Person>

</rdf:RDF>
```

This RDF ontology uses the namespace at <http://www.w3.org/1999/02/22-rdf-syntax-ns#> to express that Eric Miller is a person who has an email address `em@w3.org` and that he is a Doctor.

2.1.3 Example Ontologies

Ontologies research is motivated by the desire to create practical knowledge-based systems. In this section we consider some example ontologies and their applications.

2.1.3.1 SUMO

The Suggested Upper Merged Ontology (SUMO) project was founded as an attempt to unify various upper ontologies into a single upper-ontology which could enable ontology interoperability [94].

The ideal behind the use of such upper ontologies is that domain specific ontologies can be mapped into the upper ontology and hence the relationships between the

⁶c.f. <http://www.w3.org/TR/rdf-mt/>

⁷Taken from <http://www.w3.org/TR/rdf-primer/>



Figure 2.2: An example RDF graph. The nodes are resources, either things or data. The edges express triples, relating two resources into a Subject Verb Object triple. For example, the edge between “Eric Miller” and “http://www.w3.org/People/EM/contact#me” expresses that the contact resource has the full name “Eric Miller”.

domain ontologies can be formalised using the upper ontology. There is little empirical evidence to support this claim, and in general upper ontologies have not been of much use in solving interoperability problems⁸.

The problem for upper ontologies is that they are too general to be of much use in specific problems. You need to have domain specific knowledge to solve domain specific problems. We discuss the problem of ontology matching further in section 2.1.4.

2.1.3.2 WordNet

The WordNet project was founded by George Millar to create a database of English words and their linguistic categorisations [90]. WordNet has become one of the most

⁸c.f. Barry Smith ([121] p159): ‘The initial project of building one single ontology, even one single top-level ontology, which would be at the same time non-trivial and also readily adopted by a broad population of different information systems communities, has largely been abandoned.’

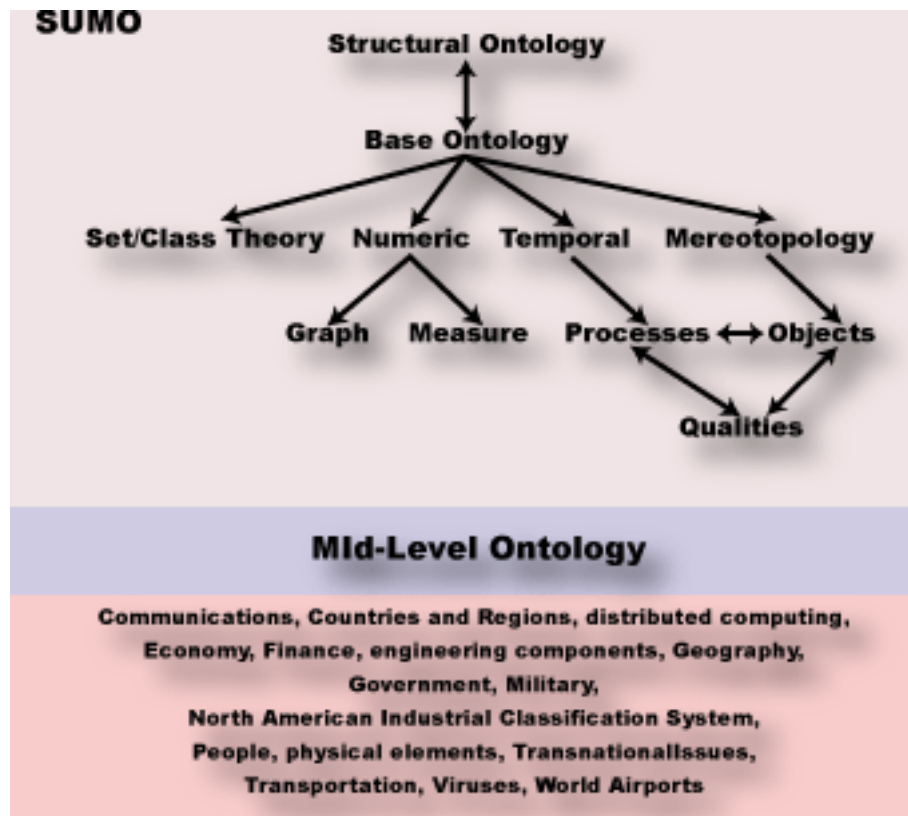


Figure 2.3: A section of SUMO, showing the upper ontology containing the most abstract concepts and the mid-level ontologies for particular domains.

important resources for natural language processing.

WordNet is not a formal ontology in the strictest sense. It doesn't provide logical representations of the concepts that the words denote. It only provides their linguistic categorisations, links to related words and English language descriptions of their meaning. However, it does satisfy a lightweight sense of ontology, since it uses a taxonomy of linguistic terms to discuss a domain of discourse composed of words. WordNet also demonstrates that lightweight ontologies can be very useful, even if they can't be used to perform logical reasoning.

2.1.3.3 SNOMED-CT

The medical community is one of the most significant real-world users of ontologies. There have been several popular medical ontologies, such as GALEN [106], UMLS⁹ and SNOMED-CT [126].

SNOMED-CT (Systematised Nomenclature of Medicine - Clinical Terms) is a

⁹<http://www.nlm.nih.gov/research/umls/>

merger of the original SNOMED-RT ontology with the CTV3 ontology. The ontology was created as part of a joint venture between the College of American Pathologists (CAP) and the National Health Service (NHS) in the United Kingdom.

The ontology provides a comprehensive terminology of medical concepts, it contains approximately 300,000 concepts. The ontology is used for managing knowledge about patients and ensuring consistent usage of terminology across sites¹⁰.

SNOMED-CT is formalised within a description logic [126]. However, the design process of SNOMED-CT, which makes extensive use of feedback from clinicians, creates incorrect relationships and duplicate concepts [27].

There has been some research into the changes that have occurred to SNOMED-CT [26], which we will consider in more detail in section 2.2.1.

2.1.4 Interoperability

As we mentioned earlier, one of the aims of ontologies research was to enable knowledge-based systems to share knowledge. This ideally requires a shared ontology, but real-world systems don't necessarily have shared ontologies. In particular, in a dynamic Semantic web setting it is unrealistic to demand that people ensure that their ontologies are compatible. Such a demand might prevent people from putting ontologies on the Semantic Web.

The World Wide Web was popular as it allowed anyone to put a hypertext document online, without having to consult and agree with others about that document's content. Similarly in a distributed system using ontologies it would be desirable for anyone to put an ontology online and for it be integrated with other ontologies automatically.

This has led to research on various forms of ontology integration [66]. There are many ways in which ontologies can be combined or compared. We shall consider the problem of ontology matching, which demonstrates many of the problems faced by any form of ontology integration.

2.1.4.1 Ontology Matching

The problem of ontology matching is to find matches between terms in two ontologies such that any expression mapped to another ontology has the same meaning as it did in the original one[66].

¹⁰See <http://www.connectingforhealth.nhs.uk/systemsandservices/data/snomed> for a description of the usage of SNOMED-CT within the UK

An example of an ontology matching system is the S-Match system [50]. This system matches terms in two ontologies by comparing the hierarchy of the taxonomy, the term structure and by using external resources. An example of an external resource is WordNet. Predicates in different ontologies might be matched according to their WordNet relations. For instance, if the word used to denote a predicate is a synonym of a word used to denote a predicate in the other ontology then the predicates might be matched as being interchangeable without any loss of meaning.

Ontology matching is based upon ensuring that the meaning of matched concepts is the same in the two ontologies. Since the ontologies are based upon a mathematical logic the meaning is theoretically based upon the model-theoretic semantics of the two ontologies. However, ontologies are created by people using natural language words. The intended meaning of the concept is based upon the natural language words [88]. Matching systems, such as S-Match, can use this feature to match terms using resources such as WordNet. However, it also raises interesting questions about what exactly the meaning of an ontology is. We shall consider such questions further in section 4.4.

Current ontology matching systems have problems with the background knowledge required to match ontologies. A deliverable of the Open Knowledge project [117] noted that

“Recent industrial-strength evaluations of matching systems ... show that lack of background knowledge, most often domain specific knowledge, is one of the key problems of matching systems.”

2.2 Ontology Evolution

This thesis is concerned with the problem of ontology evolution, which is related to the interoperability problems we saw in the previous section.

Ontologies are about an application domain, but our understanding of any domain changes over time, hence our ontologies must adapt to reflect this changing understanding. An ontology that didn't change would soon become useless to us.

This creates a problem since ontologies are naturally static. We need to design mechanisms that manipulate an ontology so that it evolves along with our understanding of the domain.

But what is our understanding of the domain? If the ontology simply is the explicit specification of our understanding of the domain, and so the two are intertwined, how

can our understanding change without the ontology changing?

There is an interesting distinction here between the implicit and explicit knowledge that an agent (whether human or machine) has¹¹. An ontology is part of the explicit, shared knowledge that an agent has. Implicit knowledge is the knowledge that an agent has that is not explicitly shared as part of some social discourse, see section 6.3.3 for further discussion.

The evolution of an ontology will depend upon the implicit knowledge that an agent has. The distinction between implicit (tacit) and explicit knowledge is recognised within the field of knowledge engineering [114] [128].

We are faced with two types of ontology evolution. One in which human beings are the arbiters of domain understanding, and this can change without the ontology being updated. And one in which an autonomous agent has an ontology representing its understanding of the domain, and so the agent must adapt the ontology to better reflect the domain.

In the first scenario the implicit knowledge is stored in people's brains and we need to design tools for ontology evolution which help people use this implicit knowledge to guide the ontology evolution. In the second scenario the implicit knowledge must be provided by the autonomous agent: the agent must determine how to evolve the ontology using the implicit knowledge and reasoning mechanisms they have.

We shall consider examples of these different scenarios in the following sections.

2.2.1 Tools

The first scenario suggests a form of ontology evolution in which humans change their representation of the domain and want to adapt an existing domain ontology to reflect this new representation.

There has been some research into this problem within the ontology research community [132] [131] [41] [130]. Stojanovic et al [130] highlight six phases for ontology evolution, based upon a process of identifying the need for changes, then determining the appropriate changes, then enacting the changes and rechecking the ontology.

As noted previously, there has been some research into the evolution of the SNOMED-CT ontology [26]. This research has focused upon classifying the changes that occur to SNOMED-CT, in particular making a distinction between the changes due to a change in reality and the changes due to a change in our understanding.

¹¹c.f. the traditional distinction in A.I. between declarative and procedural knowledge.

The aim of these research projects is primarily to help people tasked with managing ontologies to evolve their ontologies. The research doesn't address the problems of meaning that arise when considering how an autonomous agent might evolve an ontology independently. Our research is concerned with this problem, since we want to design mechanisms that can function without human intervention, so this research is not particularly relevant to us.

2.2.2 ORS

The second scenario forms the basis of research within the DReaM¹² group. In particular the ORS system [87] and research on Ontology Evolution in Physics [22]. The focus of both research projects is on signature evolution, c.f. the high-level definition of ontology given in section 2.1.2.

The ORS system actually involves both kinds of scenario. The ontologies used by ORS agents were created by people to fulfil some purpose. However, the ORS agents were designed to use these ontologies to autonomously achieve goals within the domain.

The problem that ORS was designed to solve occurred when a planning agent wanted a service from a service-providing agent and there was an ontology mismatch in their communications. The planning agent had to create a plan that involved requesting some services from service-providing agents¹³. The planning agent was tasked with spotting any mismatches that occurred during plan execution and repairing its own ontology to solve the mismatch problem. A mismatch might be a relation with an extra argument or a relation that is a superclass of the expected one in the planning agent's ontology.

When a mismatch occurred the planning agent used a decision tree to decide what repair to make to its ontology. Once the repair was made the planning agent resumed its plan.

There were several types of mismatch that the planning agent could not repair, such as a relation that didn't occur in its own ontology or was not a sub/super-class of another relation in its ontology. The scenarios that ORS could handle were restricted to agents using a different version of the same ontology. The differences between the two ontologies could not be too great as the repair mechanisms would not be able to

¹²Discovery and Reasoning in Mathematics, although the research includes more general AI themes

¹³The actual examples were based upon hand-crafted ontologies for tasks such as booking travel and accommodation for a conference

identify an appropriate repair.

2.2.3 Ontology Evolution in Physics

Ontology evolution in physics [22] is a research project which studies repair plans in the physics domain. The aim is to formalise repair plans which emulate the discovery of new physical concepts by using historical physics experiments as case studies. This research is the most relevant prior work on ontology evolution to our own research.

The basis of the Ontology Evolution in Physics project is the study of repair plans. These are similar to proof plans¹⁴, which formalise meta-level heuristics used by mathematicians in theorem proving. Repair plans are designed to capture heuristics used to spot and enact common repairs to ontologies. Repair plans have been used to make changes to the signature of an ontology, such as adding in new objects or functions, in order to resolve problems with the ontology.

The main components of a repair plan are the trigger formulae (the conditions under which the plan can be used), the schema that is matched to the existing ontology and the change to the schema that is made to the ontology. Repair plans are defined in a higher-order logic, see section 2.1.2.3.

An example of a repair plan is the inconstancy plan [28], which has been used to model the creation of the Modified Newtonian dynamics theory. The inconstancy repair plan is triggered when a function is observed to be dependent upon a parameter that it is defined to be independent from.

The motivating intuition behind the inconstancy repair plan is that there is some underlying parameter that varies between distinct sensory ontologies that is creating the different observed values. The current theory doesn't take this parameter into account, so the repair is to include it into the theory.

The repair is to identify the parameter which causes the unexpected variation in sensory observations and redefine the existing function so that it is dependent upon the parameter.

The inconstancy repair plan is defined below (the definition is taken from [28]):

Suppose we have an ontology O_t representing the current state of a physical theory and some ontologies O_s representing sensory information arising from experiments, such that different sensory ontologies give distinct values for function $stuff(\vec{s}_i)$ in different circumstances. Suppose function

¹⁴A related area of research within the DReaM group. Proof plans are used for meta-level reasoning in mathematical theorem proving.

$V(\vec{s}_i, \vec{b}_i)$ of the i^{th} sensory ontology, where \vec{b}_i contains variables distinguishing among these circumstances, returns distinct values in each of these circumstances, but is *not* one of the parameters in \vec{s}_i , i.e., $stuff(\vec{s}_i)$ does not depend on $V(\vec{s}_i, \vec{b}_i)$. We will call $stuff(\vec{s}_i)$ the *inconstancy* and $V(\vec{s}_i, \vec{b}_i)$ the *variad*. The Inconstancy repair plan establishes a relationship between the variad $V(\vec{s}_i, \vec{b}_i)$ and the inconstancy $stuff(\vec{s}_i)$.

Trigger: If $stuff(\vec{s}_i)$ is measured to take different values in different circumstances, then the following trigger formulae will be matched.

$$\begin{array}{ccc} O_s(V(\vec{s}_1, \vec{b}_1) = v_1 \dots) & \vdash & stuff(\vec{s}_1) = c_1 \\ & \vdots & \\ & \vdots & \end{array} \quad (2.1)$$

$$\begin{array}{ccc} O_s(V(\vec{s}_n, \vec{b}_n) = v_n \dots) & \vdash & stuff(\vec{s}_n) = c_n \\ O_t & \vdash & stuff(\vec{x}) ::= c(\vec{x}) \end{array} \quad (2.2)$$

$$\begin{array}{ccc} \exists i \neq j. O_t & \vdash & stuff(\vec{s}_i) - c_i \neq \\ & & stuff(\vec{s}_j) - c_j \end{array} \quad (2.3)$$

where \vec{x} can be instantiated to \vec{s}_i for $1 \leq i \leq n$, $O_s(V(\vec{s}_i, \vec{b}_i) = v_i)$ is the sensory ontology containing observations made under the condition that $V(\vec{s}_i, \vec{b}_i) = v_i$ and $V(\vec{s}_i, \vec{b}_i)$ is not an existing argument of $stuff(\vec{s}_i)$, i.e., $V(\vec{s}_i, \vec{b}_i) \notin \vec{s}_i$.

Add Variad: The repair is to change the signature of all the ontologies to relate the inconstancy, $stuff(\vec{x})$, to the variad, $V(\vec{x}, \vec{y})$:

$$v(stuff) ::= \lambda \vec{y}, \vec{x}. F(c(\vec{x}), V(\vec{x}, \vec{y})) \quad (2.4)$$

where F is a new function, whose value we will seek to determine by curve fitting against the data from the sensory ontologies.

Create New Axioms: We calculate the axioms of the new ontologies in terms of those of the old as follows:

$$\begin{array}{l} Ax(v(O_s(V(\vec{s}_i, \vec{b}_i) = v_i \dots))) ::= \{ \phi \{ stuff / v(stuff)(\vec{b}_i) \} \mid \\ \phi \in Ax(O_s(V(\vec{s}_i, \vec{b}_i) = v_i)) \} \\ Ax(v(O_t)) ::= \{ \phi \{ stuff / v(stuff)(\vec{y}) \} \mid \\ \phi \in Ax(O_t) \setminus \{ stuff(\vec{x}) ::= c(\vec{x}) \} \} \\ \cup \{ v(stuff) ::= \lambda \vec{y}, \vec{x}. F(c(\vec{x}), V(\vec{x}, \vec{y})) \} \end{array}$$

i.e., the axioms of $v(O_t)$ and the $v(O_s(V(\vec{s}_i, \vec{b}_i) = v_i))$ are the same as for O_t and $O_s(V(\vec{s}_i, \vec{b}_i) = v_i \dots)$ except for the replacement of the old $stuff$ with $v(stuff)$ and the replacement of the definition of $stuff(\vec{x})$ by the definition of $v(stuff(\vec{x}))$ in $v(O_t)$.

Repair plans do not currently deal with the problem of selecting between different possible repairs, or with evaluating the quality of a repair, however this is a current research goal of the project¹⁵.

¹⁵Prof. Alan Bundy, personal communication

We will consider the relevance of repair plans to our research problem in section 3.2.2.2.

2.3 Law

For our purposes, we can consider the Law as fulfilling two main roles within a society. The first is to resolve disputes between members of the society. The second is to ensure that the outcomes of these disputes are “fair”¹⁶.

The members of the society can be individual people or recognised legal-entities, such as a corporation. The members of a society can perform various actions that affect other members of the society, in positive or negative ways. These actions can create conflicts between members of the society. The role of a legal system is to determine what actions are permissible and what sanctions should be brought against a member who performs an illegal action.

A legal system requires mechanisms to determine and specify permissible actions, to recognise and enforce violations, and to judge what the sanctions should be in case of a violation.

Legal systems use legal rules to specify what the law is. Although, how these rules are created, what form they take and how they are applied differs between legal systems.

2.3.1 Legal Systems

The law is enacted through legal systems. They provide mechanisms for people to create and enforce laws within a society. There are two prominent families of legal systems we shall consider: common law and civil law [32]. There are other forms of legal system, but all the cases we consider will involve these two.

In [70] Peter Legrand points out that “In comparative law, the difference in styles of legal reasoning between legal families is often emphasised.” In particular, the perspective upon ontologies differs notably between common law and civil law systems. This thesis studies common law decisions. However, our research indicates that, for our purposes, the distinction between these systems is negligible and that the results should be more broadly applicable.

¹⁶We use scare quotes to indicate that the notion of fairness is contentious, we shall return to this problem later in the thesis.

2.3.1.1 Common Law

Common law has distinctive features both regarding legal sources and legal procedures. The most important sources are precedent, earlier court decisions from higher courts are binding on lower courts. The ideal is that the law should be consistent and so apply similar remedies to similar cases.

Common law uses an adversarial form of case resolution. The lawyers for the two parties generate arguments about the law. The judge acts as an umpire, checking the validity of the arguments and assessing which are superior.

2.3.1.2 Civil Law

Civil law is based upon codifying the sources of law into collections of definitive rules. Applying the law should then become a process of finding the applicable rules to a case and calculating their consequences. However, applying the rules is usually not quite as straightforward as this.

Civil law uses an inquisitorial form of case resolution. The judge should lead the investigation into the case, and the lawyers representing the two parties should assist this investigation.

2.3.1.3 Mixed Systems

There are legal systems which combine elements of different legal systems. For instance, the Scottish legal system, like much of continental Europe, was originally based upon Roman civil law. However, some fields of Scots law follow the English model of precedent based common law.

The E.U. is also potentially an example of a mixed legal system. Each constituent country has its own national legal system, of a common, civil or mixed type, but all recognise the European Court of Justice as the highest court in the E.U. and all are obliged to enact European directives in national legislation. There is some debate about whether the E.U. courts should be considered to be part of each nation's legal system, or whether they are part of international law. Certainly the bindings between E.U. courts and national courts are stronger than the ones to other international courts, such as the International Court of Justice (ICJ).

2.3.2 Legal Rules

From the perspective of Computer Science, legal rules can be viewed as defining constraints upon how a legal system should operate. They relate objects from different domains and put constraints upon how these objects should relate. For example, an Act of Parliament will specify constraints between domain-level objects, such as people, chairs and offices, and legal objects, such as rights, entitlements and prisons. It is by placing constraints between the object-level objects and the legal ones that the law states how social situations should be regulated. Note the normative aspect of the Law, the constraints are not descriptions of how the world is, but instead state how the legal system should respond to various states of the world.

An example of a legal rule might be “if, within an office environment, a person stands on a chair then their employer is not liable for any accidents that result”. A constraint is placed between a situation in which a person, in an office, stands on a chair and the legal liability of their employer for an accident that might result from this situation.

The main problem in law is interpreting legal rules. The legal system is built up upon our natural languages, it can not escape from their limitations to accurately communicate ideas. In the above example the vagueness about what it means for a person to stand on a chair, or what it means for an accident to result from the situation, would create a “hard legal case”, a case that can’t be decided mechanically by a mere “un-thinking” application of the rules [15].

2.3.3 Legal Questions

The ontologies underlying legal rules must be matched with the ontologies used to describe real-world situations. A variety of questions are usually raised when trying to apply a legal rule to a real-world situation. Lawyers recognise two categories of questions: questions of law and questions of fact.

Questions of law are questions about what the law actually is. The legal rule is not clear or there are conflicting rules, in this case the judges must determine what the law is in that case.

Questions of fact are questions about what actually happened in the real-world situation. There might be sources of evidence that lead to conflicting hypotheses. The judges must determine what is most likely to have happened in the real-world situation and judge the case accordingly. The standards of truth can differ between types of legal

cases. In English law a criminal case requires that the real-world situation be established beyond reasonable doubt for the corresponding sanction to be applied, whereas civil cases merely require the balance of probabilities to be in favour of that situation.

2.3.4 Legal Ontologies

We have seen that legal rules are the form in which the law is specified. These legal rules mention various concepts, and in doing so define an ontology.

There are also ontologies that are not explicitly mentioned in any legal rule, but which underlie the legal system, such as the ontologies used to describe the principles of the legal system.

There have been various attempts to formalise legal ontologies which we shall look at in the following section.

2.3.4.1 Formal Ontologies of Law

Thorne McCarty was one of the first researchers working on legal ontologies. He originally worked on models of legal reasoning based upon logic programming. As a result of this research he proposed a Language for Legal Discourse (LLD) [81] as a basis for modelling legal reasoning.

The scope of the projects using Language for Legal Discourse was quite limited. McCarty's models of legal cases were not fully automated, and the work required to hand-craft the representation of a legal case meant that few cases were modelled in this language [82].

There have been other ontologies proposed since then, Pepjin Visser and Trevor Bench-Capon give a survey of the various proposals in [146]. These ontologies were primarily concerned with creating upper-ontologies for modelling legal cases and legislation, and have not been extensively used to model actual legal cases.

The largest scale formalisations of legal cases have been done in quite lightweight formalisms. Zelenkowski's SPLIT-UP project had one of the larger databases of cases, at one hundred and fifty [159]. However, the case representation in this project was quite simple, simply a set of features for each case.

A more recent, and larger scale, project to create useful legal ontologies is the Estrella project¹⁷. The aim here is to create ontologies primarily for modelling legislation, rather than legal cases.

¹⁷<http://www.estrellaproject.org>

The main difficulty in using the foundational legal ontologies is that by themselves they aren't especially useful. Most legal cases involve a real-world domain which must also be formalised as part of a model of a legal case. This problem has been recognised within the AI and Law community [18], and some commonsense ontologies have been produced, but they are not extensive enough for the models of legal reasoning we would like to produce. We will look further into this problem of commonsense knowledge in chapter 6 6.3.

Another interesting example of legal ontologies research is McCarty's modelling of the concept of possession [84]. He analyses the concept of ownership and tries to construct a formal model of the concept. This is particularly relevant for our work, but McCarty doesn't discuss exactly what reasoning is used to evolve the concept, rather he focuses upon what changes occur to the concept and the broad social reasons for these changes.

In [112] Edwina Rissland and Tim Friedman study the changes that occur to a legal rule over time, based upon the features that are added and removed from the rule. Again, this analysis is based upon long-term change to the concept, rather than the mechanisms used to make one particular change. The justifications for the changes are not studied.

The work of Andre Valente and Joost Breuker (with various collaborators) [19] is of particular interest to us. Their work makes a distinction between legal (normative) knowledge and commonsense conceptual knowledge, with legal reasoning involving forming mappings between the legal knowledge used to describe legal rules and the commonsense knowledge used to describe real-world events. They have developed several systems, such as TRACS [20] and ON-LINE [144]. More recent work has focused on the use of OWL ontologies to model legal rules and commonsense concepts in the HARNESS system [145]. This work is quite influential on ours, in particular their distinction between commonsense knowledge used to describe a real-world event and the legal knowledge contained in a source of law and their characterisation of legal reasoning as forming mappings between the legal knowledge and the commonsense knowledge. This influence can be seen in the model we propose in section 6.4. Our model primarily differs from theirs in the use of a contextual logic to provide a basis for the representation.

2.3.4.2 Textual Ontologies

The existing formal ontologies are not extensive enough for our purposes. We are interested in modelling the evolution of legal ontologies in a manner similar to the Ontology Evolution in Physics project. So we must broaden the kinds of ontology we are considering.

The Ontology Evolution in Physics project uses the documentation of historical Physics experiments as a basis for their models. We shall consider the ontologies defined in sources of legal rules as the ontologies whose evolution we will try to model.

These ontologies are textual, in that the ontology is defined by a natural language discourse. The text mentions various concepts and states rules that define those concepts.

Certainly natural language is not a formal logic, and so raises various problems for constructing formal models. In particular, we are hand-crafting the representations for our model, and might ignore relevant complexities present in the actual document. We shall consider these problems later in this thesis.

Note that, we will use the words “word” and “concept” in different parts of this thesis to mean essentially the same thing, i.e. the concept which the word denotes. We use the word “word” sometimes, as this is what appears in the legal texts, although we are usually talking about the concept that the word denotes. We only use this terminology when the word is not ambiguous, i.e. it is agreed that there is one concept that the word refers to, but the definition of this concept is uncertain.

2.3.5 Legal Reasoning

We are focusing upon ontology evolution in legal cases and so we are interested in the reasoning produced in legal cases. In a legal case a dialogue between the lawyers and the judges is conducted. This dialogue is recorded as a discourse about the case in the case proceedings.

The dialogue in a legal case is composed of arguments¹⁸. The legal case raises certain questions. The lawyers in the case must create arguments to support or attack answers to these questions.

Since we are looking at common law legal-cases we are also interested in the reasoning with precedents. We argue that this form of reasoning is a particular form of

¹⁸c.f. [149] for an analysis of argumentation in various forms of dialogue.

argumentation, where the domain of argumentation are the claims about similarity between legal cases.

Legal reasoning is arguably no different from the reasoning we use to solve problems in everyday life. However, even though the mechanisms of reasoning might be no different, the reflective awareness of the reasoning is different. The methods of legal reasoning are frequently made explicit in the discourse. Lawyers are aware of the mechanisms, such as precedential reasoning, they are using and there are rules about how reasoning should be conducted, which are more rigidly enforced within the law than they are within everyday discourse.

The permissible forms of legal argumentation and interpretation are taught at every law school. In particular, the teachings regarding statutory interpretation, such as [39], form a distinct body of knowledge that law students will be taught. The rules of statutory interpretation form an explicit set of meta-rules within a legal system.

2.3.5.1 Argumentation

We have claimed that legal discourse is based upon argumentation, and that legal reasoning involves producing a legal discourse. From a computational perspective, argumentation is connected with logic. For lawyers, the logical aspects of argument are less clear, and common law in particular has often emphasised the rhetorical nature of argument, c.f. [97]. It is a methodological assumption of this thesis that at least some parts of legal argumentation can be reconstructed in a logical formalism. However, this means that not all features of a legal case can necessarily be explained in the frameworks we study.

Recently a field of informal logic has arisen, which studies the patterns of argument in discourse [51]. It is this notion of argumentation that we are using as a basis for legal reasoning.

The argumentation in a legal case is about the claims made by the opposing sides in a legal case. The claimant in a legal case will claim that the defendant has violated some legal rules and that hence certain consequences must be legally enforced. Both sides will then create arguments about these claims and their consequences. The arguments are assessed by the judges in the legal case.

Toulmin's model of argumentation defines a basic structure for argumentation [141]. His model of argumentation has been refined by various authors, e.g. [152], but gives us a basic vocabulary for understanding arguments.

Toulmin's model consists of the following elements:

- Claim
- Data
- Warrant
- Backing
- Qualifier
- Rebuttal

The claim is the proposition that we are trying to argue for, such as “You should carry an umbrella”.

The data is the evidence that we have for our claim, such as “It is raining outside”.

The warrant is the link that allows us to infer the claim from the data, such as “An umbrella will keep you dry in the rain and you would prefer to be dry.”

A backing provides additional support to the warrant, such as “An umbrella creates a barrier between you and the rain”.

The qualifier indicates the strength of the warrant, how well does the evidence justify the claim given the warrant. For instance, we might qualify the warrant above as being quite likely, an umbrella will usually keep you dry in the rain.

A rebuttal is an exceptional condition which prevents the warrant from justifying the claim, such as if it was also very windy, in which case the umbrella would be blown away or inside out and so would not provide protection from the rain.

The most important elements of argumentation are the claim, the data and the warrant. In many arguments warrants are implicit, creating arguments that are called *enthymemes*.

In a legal case the relevant claims are about the questions raised by the case and the legal consequences required in the case. The evidence is primarily provided by the relevant laws, the precedent cases and commonsense knowledge. The warrants are based upon either legal knowledge or commonsense knowledge.

2.3.5.2 Precedent

Reasoning by precedent is reasoning by analogy: a prior case is presented as being similar to the current one. The claims about similarity are attacked by making distinctions between the prior case and the current one.

The aim of the lawyer is to identify the *ratio decidendi* in the prior case, the rule that was used to justify the prior decision. Unfortunately, these rules are not always explicitly stated in the prior case. The lawyer's skill is to argue for a generalisation of the prior case to an abstract rule that covers the current case. This generalisation can be attacked by making distinctions between the current case and the prior one, which are not accounted for by the abstract rule.

For more detail regarding reasoning with precedent see [69].

2.3.5.3 An example

The following example is taken from Twining and Meiers [142], and demonstrates some of the problems faced in interpreting the law. The example is about a boy named Johnny who has recently developed a sweet tooth. One day his mother catches him snacking on a pot of jam in the pantry. She scolds him and tells him that he "Must not enter the pantry". Johnny asks "What does enter mean?", "It means to go into", replies his mother.

In the next few days the following incidents occur: Johnny is caught using a broom to try and hook the jam jar out of the pantry. When caught he argues that he "didn't go into the pantry". Later the family's cat enters the pantry and starts eating the main course for a dinner Johnny's mother was organising. His mother walks into the kitchen to find Johnny standing outside the pantry laughing while the cat devours the meal.

On both of these occasions there is some ambiguity about the applicability of the norm "Johnny must not enter the pantry". In the first situation Johnny is clearly trying to eat the jam, and is using the broom to avoid being physically in the pantry. This violates the intent behind the rule, and arguably also violates the rule: is Johnny going into the pantry by using the broom as an appendage? In the second situation Johnny should be stopping the cat from eating the meal, even if this means entering the pantry.

The example demonstrates some important aspects of legal rules. When Johnny's mother says "You [Johnny] must not enter the pantry" she clearly intended to prevent Johnny from snacking on the Jam, but not to prevent him from saving the meal from the cat. Our interpretation of the rule is controlled by contextual factors which influence how we apply it to a real-world situation.

2.3.6 Evolution of the Law

The simplest form of legal evolution is that rules are added, edited or deleted within a legal system. Particularly interesting examples of this are:

- A new rule affects the interpretation of other rules (e.g. a rule is added to the law which states that new rules take precedence over old ones.).
- A precedent affects the interpretation of a rule (e.g. a precedent case establishes a rule that allows the use of Parliamentary debates to aid statutory interpretation¹⁹)
- A rule which affects the interpretation of other rules is removed.
- A new rule is inducted from precedent cases (e.g. in the common law a *ratio decidendi* is given).

Many of these changes occur outside of legal cases, e.g. when legislation is passed by an official body. The evolution of the Law that occurs in legal cases happens when the available legal rules are inconclusive: when there is a conflict between rules, when the rules are unclear or when the rules are perceived as unjust or absurd. In these cases the competing parties must propose their theories, in favour of competing outcomes, to the deciding audience. The audience will evolve the law by making a decision over which theory is preferable, or by proposing a theory of their own for the case.

2.3.7 Modelling Ontology Evolution in Legal Reasoning

There are two features of the legal domain that are particularly relevant for any attempt to model legal reasoning: the natural-language format of legal discourse and the implicit biases in legal reasoning.

The first feature raises some interesting questions about how we can create a formal model from the natural language discourse. The observations we have of legal reasoning come from the natural-language case proceedings. These case proceedings contain a description of the case and the arguments that were presented in the case.

We must take these textual descriptions and construct logical representations of the reasoning. This requires either someone or some process to translate the text into a logical representation. What we would like to have are models of the reasoning about ontology evolution in the legal case.

¹⁹c.f. *Pepper v. Hart* [1992] UKHL 3

The model might predict the arguments that occurred in a historical case, or suggest plausible arguments. This distinction, between different possible purposes of a model, is the second claim of this thesis, which we shall discuss in more detail later.

The difficulty for us is to justify the mapping from the text to the formal logic. If we are hand-crafting the formal representations and then paraphrasing the formal output into natural language, we are open to the criticism that the output of the model is not related to the actual output in the case, and hence the relationship between the formal model and the actual process is uncertain.

This problem is not specific to our work, but affects AI and Law research generally. We must use our knowledge of the domain to justify the modelling step and to assess the relevance of the model to the actual process.

The second feature of legal reasoning is that the case proceedings might not contain all the information that was relevant to the outcome of the case. The judges in a case might have implicit biases that affect their reasoning. Lawyers will appeal to these biases when they craft their arguments. So the arguments generated in a legal case might only make sense if these biases are recognised.

Examples of these biases are:

- Our sense what is morally acceptable
- Public Opinion
- Political factors in the judiciary.
- Economic concerns.
- Racism, Sexism, and other prejudices.

These biases are implicit constraints upon how we conceptualise the Law, and are distinct from the explicit legal rules. For example, in the ‘Naughty Johnny’ case we would consider Johnny to have violated the rule “You [Johnny] must not enter the pantry” if he had stopped the cat from eating the meal. However, we might not want to punish him for it, as we would be punishing him for being helpful. Our implicit sense of fairness leads us to regard this outcome as undesirable.

Sometimes these biases are an explicit part of the law, and are appealed to as part of legal argument. For instance, when a law uses terms such as “reasonable doubt” the public opinion regarding what is reasonable is part of the law.

Political factors may influence the reasoning of judges. For example, in the case of *Buchanan v. Babco* (see section 3.2.2) it is arguable that Lord Denning's views on importing broad methods of interpretation into UK law were influenced by a political desire to give the judiciary more freedom in interpreting laws set by Parliament.

Economic concerns can be used to explain the evolution of the law, the Law and Economics movement has become quite popular recently. An example²⁰ might be the change to the ownership of land in U.S.A. that occurred when air planes became common. Prior to air travel, ownership of land meant that you had property rights over the sky above the land. This was challenged when a farmer took the government to court for trespassing due to flights that went over his farm²¹. The judges held that it was against common sense for this notion of property to hold, and so now property rights do not extend all the way up into the skies. It is arguable that the judges were influenced by economic concerns; to have upheld the traditional property rights would have made air travel much more difficult if every flight was potentially subject to numerous trespass claims.

The important distinction to make is between the *explicit* legal rules which place constraints upon how the legal system should operate and the *implicit* biases of the agents within the legal system. There is an associated distinction between the context of decision making, in which a judge makes a decision regarding the case, and the context of rationalisation, in which a judge proposes an argument to justify their decision.

2.4 Summary

We have surveyed ontology and ontology evolution research. In particular the Ontology Evolution in Physics project which is most akin to our own.

We have given a high-level overview of what the Law is and how it can evolve. We are focusing upon the evolution of textual ontologies within legal cases, and trying to construct formal models of the reasoning involved and the resulting changes.

In the next chapter we consider some example legal cases to demonstrate some of the problems raised in trying to construct these formal models.

²⁰taken from Lawrence Lessig's book *Free Culture* [72]

²¹Actually there is more to the story. These were low-level flights of military aircraft, and the noise caused by the flights was scaring the farmer's chickens causing them to panic and kill themselves by flying into barn walls.

Chapter 3

Ontology Evolution in Legal Cases

We started this thesis by proposing to study ontology evolution in legal reasoning. However there are many activities that can create and change the law within a legal system, and so there are different events that can lead to the evolution of legal ontologies. We have been studying the evolution of ontologies in legal cases. In this chapter we will look at what happens in a legal case, what reasoning is used to resolve a case and how the ontology of the law evolves in a legal case.

3.1 Overview of a Legal Case

Legal cases are used to resolve disputes about the law. Two parties have conflicting interests which the law must resolve. A legal case provides a forum in which these conflicts can be explored and resolved.

A legal case is raised by one party registering the case with the appropriate court. For instance, within the UK, the Crown Prosecution Service (CPS) will raise a case within the Criminal Courts on the basis of a police investigation. Alternatively, a civil case will be raised when one of the party's lawyers registers the case with the appropriate court.

Once a case has been raised with a court the parties must then prepare their arguments for the case. These arguments will be used during the case to justify that party's preferred outcome. At the end of the case proceedings the judges¹ will decide the outcome of the case.

¹We have been looking at cases in which judges have determined the outcome to the case, so our discussion is based upon judges making all decisions, without mention of juries or other forms of arbitration. Since the evolution of the law will primarily occur in appellate courts, where judges make the decisions, this is a justified use of the terminology and will not affect our conclusions.

The outcome of a case matters not only to the parties involved, but also the wider legal system. In common-law systems the impact of a case can be to establish a precedent on some question of law. However even in legal systems without binding precedent the analysis produced to resolve a dispute in one case may have an impact upon future cases. We have been focusing upon cases in common-law jurisdictions and our discussion, and results, are focused upon the common law. Although, as argued previously, we think that the results should generalise to cases in other jurisdictions.

In a legal case the law must be applied to a real-world event. An event has occurred in which one party, the party pursuing the case, feels that the law can be applied to create a preferable outcome for them. The problem for both parties is to argue for an interpretation of the law, as it applies to the event, that results in a legal win for them.

3.1.1 **Ontology Evolution in Legal cases**

We surveyed some of the ways in which the law can change in section 2.3.6. Historically, much new law within common-law legal-systems has been created by legal cases. However, these days most new law is created via statutes.

What can change are the methods used to interpret and apply the law. Judges must resolve any questions of law that occur in a legal case. It is not possible to leave a question which might affect the outcome of the trial unanswered. In cases where the law is uncertain, judges have to decide what the law is.

The uncertainty in the law occurs because of **inconsistent** law and **incomplete** law. The law is inconsistent if there are multiple laws which each have clear interpretations but which are collectively inconsistent. A law is incomplete if it has multiple interpretations, so we don't know exactly what it means. The problem of *open texture*, which we will study in the next chapter, is created by incomplete law. In particular, incomplete laws in which the terms used are unclear, and for which there are multiple plausible interpretations of how the law relates to the facts of the case.

We are predominantly interested in cases where the law is incomplete, as these are the cases in which the existing law is refined by adapting it to the current case. Although the problem of inconsistent law is also relevant to the cases we study in this chapter, and this thesis.

The evolution of the law occurs in the resolution of this uncertainty. The reasoning in the case is used to argue for interpretations of the law. The outcome of the case refines the law by deciding what it means in that case.

Our analysis of the ontology evolution in a legal case thus focuses on the ontology changes that occur and the argumentation used to make these changes.

3.2 Some Example Legal Cases

This section presents some case studies of legal cases. The cases we have chosen demonstrate some of the possible changes to the law that can occur in a legal case. Our analysis of the following legal cases focuses upon the reasoning relevant to the ontology evolution in the case. We don't give a comprehensive account of the case proceedings.

We start with the case of *Popov v. Hayashi* which has recently been modelled within the AI and Law community [156]. We see that the case raises many problems that make it difficult to construct a formal model of the ontology evolution in the case.

The remaining cases follow a similar theme, all are based upon disputes regarding loss of or damage to an object being transported internationally. These cases involve the interpretation of international law within the English courts.

3.2.1 *Popov v. Hayashi*

The case of *Popov v. Hayashi*² was a dispute over possession of a baseball. On October 7th, 2001, a baseball record was broken for the number of home runs³ hit by a batsman in a game. Barry Bonds hit 73 home runs in that game. The balls from previous record breaking games had been sold for significant sums of money, Mark McGwire's 1998 70th home run ball sold at auction for \$3,000,000.

Some people had anticipated that Barry Bonds would break the record during the October 7th game, and congregated at the places where a ball would be most likely to land, with the intention of catching the ball. The 73rd, record-breaking, ball landed in one of those spots, where Alex Popov and Patrick Hayashi were standing in a crowd. Alex Popov made an attempt to catch the ball, and briefly had the ball in his baseball glove, but was assaulted by the rest of the crowd. The ball left his baseball glove and was lost in the resulting fracas. During the confusion, the ball was spotted by Hayashi, who pocketed it and only revealed it when cameras could record his possession. Popov

²Popov v. Hayashi, 2002 WL 31833731 (Cal.Superior Dec 18, 2002) (NO. 400545)

³A home run is when a batter can run around all bases. The most common type of home run is when the ball is hit outside the pitch, which is what happened in this game. The balls are traditionally gathered by fans and are valued mementos of a game.

subsequently sued Hayashi for possession of the ball.

The issue in the case was whether Popov had possession over the ball, and so a legal claim to it. The legal concept of possession was defined by case law. But prior cases didn't cover all possible situations involving possession, and indeed recognised that requirements for possession differed between domains. So the judges in *Popov v. Hayashi* had to find a satisfactory definition of possession for this case.

The case was resolved by acknowledging that both parties had a claim for possession of the ball, and that they shared possession. The ball was sold and the proceeds divided between them.

3.2.1.1 Argumentation

The ontology evolution in this case is based upon the arguments regarding possession. The existing concept of possession was vague and couldn't be applied to the case. The Judges needed to find a resolution of the situation that led to an acceptable outcome.

Popov had pled causes of action for conversion⁴, trespass to chattel⁵, injunctive relief⁶ and constructive trust⁷. The main basis for these claims was that Popov had possession over the ball, and so Hayashi had subsequently violated that possession, both unintentionally (when he picked it up) and intentionally (when he didn't give it back to Popov). The desired outcome for Popov was that the ball be given to him.

The judge dismissed the trespass to chattel claim, that Hayashi intentionally deprived Popov of the ball, on the grounds that Hayashi had not damaged the ball or interfered with Popov's use of the ball.

With the trespass to chattel claim dismissed, the debate focused on the issue of possession. The ball was initially the possession of Major League Baseball. When it was hit it became intentionally abandoned property. The first person to claim possession of the ball becomes its new owner.

The following definition of possession was proposed, by Professor Brian Gray, for the case:

“A person who catches a baseball that enters the stand is its owner. A ball is caught if the person has achieved complete control of the ball at the

⁴Conversion is the wrongful exercise of dominion over the personal property of another

⁵Trespass to chattel exists where personal property has been damaged or where the defendant has interfered with the plaintiff's use of the property.

⁶An equitable remedy, not necessarily financial, for the injustice. In this case, Popov wanted the ball back.

⁷A particular form of equitable remedy.

point in time that the momentum of the ball and the momentum of the fan while attempting to catch the ball ceases. A baseball which is dislodged by accidental contact with an inanimate object or another person, before momentum has ceased, is not possessed. Incidental contact with another person is contact that is not intended by the other person. The first person to pick up a loose ball and secure it becomes its possessor.”

However, this definition denies Popov possession over the ball since he was still in motion when he had the ball in his glove, and so didn't catch the ball according to the above definition. Popov argued that the requirement for complete control was too strong. He justified this with the opinions of Professors Bernhardt and Finkelman who suggest that possession occurs when “an individual intends to take control of a ball and manifests that intent by stopping the forward momentum of the ball whether or not complete control is achieved.”

This alternative definition of possession applies in cases involving hunting or fishing wild animals or the salvage of sunken vessels. For example, in the case of hunting possession occurs when the animal is mortally wounded, not when it is eventually captured. The alternative definition was needed since a mortally wounded animal can still run for a distance before stopping.

However, the judge rejected this alternative definition on the basis that there was no reason someone couldn't achieve complete control of the baseball. And so Gray's definition applied to the case.

The next problem was that Popov was assaulted by the crowd. Popov could not supply evidence that he would have achieved possession of the ball had he not been assaulted, and so could not claim full possession of the ball. However, the judge couldn't deny Popov possession either, since to do so would be to endorse the violent actions of the crowd.

Since this branch of argumentation was not successful for Popov an alternative approach was attempted. It is possible to pursue an action for conversion where the plaintiff has failed to establish possession or title. Instead, an action for conversion can be brought where the plaintiff has a *right to possession*. So Popov now has to argue that he has a right to possession over the ball.

The court adopted the following rule regarding rights to possession: “Where an actor undertakes significant but incomplete steps to achieve possession of a piece of abandoned personal property and the effort is interrupted by the unlawful acts of others, the actor has a legally cognisable pre-possessory interest in the property. That pre-possessory interest constitutes a qualified right to possession which can support a cause

of action for conversion.”

However a claim of pre-possessory interest for Popov affects Hayashi’s claim to possession. The court claimed that it was required to balance the interests of all parties, and so Hayashi’s claim to possession couldn’t be dismissed. An award of the ball to Popov would be unfair to Hayashi.

The case has been analysed by researchers in AI and Law⁸ [156]. In particular, the conflicting values raised by possible outcomes have been represented within a value-based argumentation framework (see section 5.3). The values attached to the different proposed outcomes of the case are part of the argumentation in the case. For instance, a verdict in favour of Hayashi was attacked on the grounds that it would allow the result of the case to be determined by the crowd’s assault on Popov. However, a verdict in favour of Popov would unfairly penalise Hayashi, who wasn’t part of the crowd who assaulted Popov.

3.2.1.2 **Ontology evolution**

The change to the law is not clear in this case. Prior to the case the concept of possession was uncertain. However, this definition was also known to vary between domains, and so the general uncertainty was accepted, with specialisations for different domains.

A definition for possession was proposed in the case, but we don’t know why the definition was acceptable. What is interesting is that the judges accepted the use of the baseball definition of “caught” to determine the legal possession of the baseball, even though the ball was caught by people not playing within the game.

So what we have is a form of legal evolution through the import of non-legal rules into the legal domain. On a basic level this specifies the meaning of possession in the case. On a meta level we have a precedent for the use of non-legal rules to resolve a domain specific definition. This precedent could be used to justify resolutions of similar disputes in related domains.

The general notion of possession could be formalised as:

$$possesses(Person, Object) \leftrightarrow intends(Person, control(Object)) \wedge control(Object)$$

This states that a person possesses an object if and only if they intend to control the object and they actually do control the object. However to make this definition

⁸The case was the suggested example for the recent workshop on legal case modelling at ICAIL 2009, <http://www.csc.liv.ac.uk/katie/LegalCasesWorkshop09-CFP>

useful for an automated reasoner we would have to define *intends*⁹ and *control* as well. Defining these terms would require defining others, in a process that has no clear ending. This relates to the problem of commonsense knowledge, which we shall consider in section 6.3.

The specific concept of possession proposed for the case by Professor Gray could be formalised as:

$$\begin{aligned}
 \textit{possesses}(\textit{Person}, \textit{Ball}) &\leftrightarrow \textit{catches}(\textit{Person}, \textit{Ball}) \wedge \\
 &\quad \neg \textit{incidental_dislodged}(\textit{Person}, \textit{Ball}) \\
 \textit{catches}(\textit{Person}, \textit{Ball}) &\leftrightarrow \textit{complete_control}(\textit{Person}, \textit{Ball}, t) \\
 \textit{complete_control}(\textit{Person}, \textit{Ball}, t) &\leftrightarrow \textit{momentum}(\textit{Ball}, t) = 0 \wedge \\
 &\quad \textit{momentum}(\textit{Person}, t) = 0 \\
 \textit{incidental_dislodged}(\textit{Person}, \textit{Ball}) &\leftrightarrow \textit{dislodged}(\textit{Person}, \textit{Ball}, e) \wedge \\
 &\quad \textit{cause}(e', e) \wedge \\
 &\quad \textit{incidental_contact}(\textit{Person}, e') \\
 \textit{incidental_contact}(\textit{Person}, e') &\leftrightarrow \exists \textit{Person2}. \textit{contact}(\textit{Person}, \textit{Person2}, e') \wedge \\
 &\quad \neg \textit{intended}(\textit{Person2}, e')
 \end{aligned}$$

We have used event-calculus style notation [91] to formalise the idea of incidental contact. Again, if we wanted to use these axioms in an automated reasoner we would need to define the terms used in the above definition.

The problem for modelling the ontology evolution of the concept of possession within the case is that there is little in-depth argumentation about the meaning of possession. It is recognised that the definition of possession is contentious, but the above specialisation is offered without any detailed explanation. The authority of the Professor seems to be adequate to justify it.

The problem with modelling something like possession is that it is ultimately a commonsense concept that everyone has intuitions about, and these intuitions are hard to model. The following cases feature interpretations of statutory concepts, rather than foundational concepts. In particular we look at cases involving problems with interpreting international treaties. Since such cases don't directly involve commonsense concepts, such as possession and fairness, they may be easier to model.

⁹In particular, the *intends* relation is between a person and a proposition, which requires either meta-level representation (see section 5.8) or a modal logic.

3.2.2 Buchanan v. Babco

The case of *James Buchanan and Co Ltd v. Babco Forwarding and Shipping Ltd*¹⁰ involved the plaintiff (Buchanan) seeking compensation from the defendant (Babco) for the sum of £30000. Buchanan had sold 1000 crates of whisky, to a purchaser in Iran, and had arranged a contract with Babco to transport the whisky from Glasgow to Tehran. The contract was agreed under the terms of the ‘Convention on the Contract for the International Carriage of Goods by Road’, specifically its UK enactment in the ‘Carriage of Goods by Road Act 1965’.

On January 24, 1975, the whisky was taken from Buchanan’s warehouse in Glasgow and loaded into Babco’s lorry. During the following weekend the whisky, having been left unattended in a lorry park outside London, was stolen. Since the theft was in the UK the whisky was liable for excise duty, and Buchanan was subsequently charged £30000 for this.

Buchanan had sold the whisky for £7000 to their buyer in Tehran, and there is no doubt that Babco was liable for this sum, as their negligence had caused the loss.

However, a further problem arose. Since originally the whisky was intended for export, no excise duty on the alcohol was due. But the theft happened in the UK, and Buchanan couldn’t prove that the whisky had left the country, so Buchanan became liable for excise duty on the whisky. Buchanan was therefore subsequently charged £30000 for this by the Inland Revenue. Their loss therefore was the £7000 lost revenue, and the £30000 taxes.

The debate in the case was over whether Babco was also liable for compensation of the excise duty.

The relevant sections of the Act were as follows.

Article 17, paragraph 1, of the Convention provided:

“The carrier shall be liable for the ... loss of the goods ... occurring between the time when he takes over the goods and the time of delivery, ...”

Article 23 provided:

“1. When under the provisions of this Convention, a carrier is liable for compensation in respect of total or partial loss of goods, such compensation shall be calculated by reference to the value of the goods at the place and time at which they were accepted for carriage.

¹⁰James Buchanan & Co. Ltd. v Babco Forwarding & Shipping (U.K.) Ltd.,[1978] A.C. 141

2. The value of the goods shall be fixed according to the ... current market price or, if there is no ... current market price, by reference to the normal value of goods of the same kind and quality ...

....

4. In addition, the carriage charges, customs duties and other charges incurred in respect of the carriage of the goods shall be refunded in full in case of total loss ... but no further damages shall be payable ...”

The case went through the civil court where Master Jacob found the defendants liable for the compensation of the whisky plus excise duty. The case was appealed by the defendants, in which the Court of Appeal (Lords Denning, Roskill and Lawton) agreed with the verdict of Master Jacob. When the case was again appealed to the House of Lords the judges supported this decision by 3 to 2.

The results of the case were that:

- there could be more than one “current market price” at the same time for the same type of goods depending on where they were going.
- the words “other charges incurred in respect of carriage” were loosely drafted and so could be given a broad interpretation in accordance with the intentions of the makers of the Convention; that the French text of the Convention, at which it was permissible to look for assistance, supported a broad interpretation of the Convention.

However the case also raised important questions about how English law could be interpreted. The Act of Parliament was an enactment of an international treaty, which had an English and a French version. The UK has a dualistic approach to international treaties; this means that a treaty is not part of domestic law as soon as it is signed, it has to be enacted by an Act of Parliament. So the law of the UK is purely that stated in the Act of Parliament. However the French version was also a definitive version of the treaty. Should judges be allowed to refer to the French version to help interpret the English version, and if so how?

The French legal system has a different approach to interpreting law from the English legal system. Should English judges be using English rules of interpretation on the French treaty? Or should they use French rules of interpretation? Furthermore was it acceptable for UK courts to use French cases that clarified the French version of the statute?

In the Court of Appeal, Lord Denning advocated using a teleological approach to interpretation¹¹, as he considered this to be the method that would be used by a French court. Under his teleological interpretation the phrase “other charges incurred in respect of carriage” included excise duty. The House of Lords disagreed with his approach, with the majority of judges advocating English methods of interpretation.

The ontology evolution in the law is partly the evolution of the concepts denoted by “current market price” and “other charges incurred in respect of carriage”, but also the evolution of the meta-legal rules regarding interpretation of multi-lingual treaties. These meta-legal rules are far more vague than the law itself. There is no explicit rule that Lord Denning uses to justify his interpretation, rather he introduces the idea of a teleological interpretation and then states that his interpretation is teleological. The rules used are implicit in the discourse between the judges.

3.2.2.1 Argumentation

The plaintiff’s argument was based upon two claims:

- That the “value of the goods at the time and place they were accepted for carriage” was £37000, the price that the whisky would have been sold for in Glasgow.
- That “other charges incurred in respect of carriage” included the excise duty.

At the Court of Appeal¹² case, Lord Denning noted that “The common law takes the value of the goods at the place and time at which they ought to have been delivered by the carrier.”¹³

Regarding the meaning of article 23 paragraph 2 Lord Denning argues:

“Buchanan submit that it was the market price at which the whisky could be sold at the door of the warehouse in Glasgow to a purchaser for the home market. That is, £37,000. ... Suppose that the whisky was not stolen in England but somewhere on the continent of Europe or in Asia before it was delivered to the consignee in Teheran. Buchanan would not have been liable to pay the excise duty of £30,000. ... I cannot think that Buchanan could claim compensation for the £30,000 if they were never liable to pay it. ... That value must be ascertained at that place and time. It cannot vary

¹¹An approach to interpretation that tries to give effect to the *telos*, or purpose, of the statute: what the lawmakers wanted to achieve, even if they haven’t clearly expressed it. See [73]

¹²James Buchanan & Co. Ltd. v Babco Forwarding & Shipping (U.K.) Ltd. [1977] Q.B. 208

¹³*ibid* p212

according to subsequent events, that is, whether they are lost or stolen in England or anywhere else.”¹⁴

Regarding the meaning of article 23 paragraph 4 Lord Denning argues:

“Buchanan submit that the £30,000 excise duty was a charge “incurred in respect of the carriage of the goods.” I must say that, if this article is to be construed according to our traditional rules of interpretation, the £30,000 was not such a charge. Strictly interpreted, those words comprehend only charges for the actual carriage of the goods and other charges incurred in respect of the carriage , such as packing, insurance, certificate of quality, and so forth.”¹⁵

Lord Denning is notable for advocating a European method of interpretation: “We ought, in interpreting this convention, to adopt the European method.”¹⁶ He argues in favour of this method as follows:

“We had a valuable paper on it by the President of the court (Judge H. Kutscher) which is well worth studying: “Methods of interpretation as seen by a judge at the Court of Justice, Luxembourg 1976.” They adopt a method which they call in English by strange words - at any rate they were strange to me - the “schematic and teleological” method of interpretation. ... They go by the design or purpose which lies behind it. When they come upon a situation which is to their minds within the spirit - but not the letter - of the legislation, they solve the problem by looking at the design and purpose of the legislature - at the effect which it was sought to achieve. They then interpret the legislation so as to produce the desired effect.”¹⁷

And regarding a gap in the legislation Lord Denning argues:

“It speaks only of the charges incurred “in respect of the carriage of the goods,” but says nothing of the charges consequent on the loss of the goods. I think we should fill that gap. ... It seems to me that it was intended that the sender should not be limited to the value of the goods as defined in paragraph 1 of article 23 ... But that he should also be compensated for any additional expense that he incurred directly by reason of the loss. ... The carrier negligently left the whisky unattended, and it was stolen ... It would be most unjust that they should have to bear this expense themselves when it has been brought about solely by the negligence of the carrier. The only sensible solution is that the carrier should compensate the sender for the expense. The men who framed the Convention and agreed to it must be presumed to have intended this.”¹⁸

¹⁴*ibid* p212

¹⁵*ibid* p213

¹⁶*ibid* p213

¹⁷*ibid* p213

¹⁸*ibid* p214

Lord Roskill notes that the plaintiffs had initially proposed an argument similar to the one proposed by Lord Denning above.

“Mr. Johnson suggested that unless that excise duty were included in the normal value of that whisky at that place and time, anomalous results would follow. He instanced two lorries coming from Glasgow in convoy on a motorway, one carrying whisky for export and the other carrying identical whisky for distribution in this country, both under identical C.M.R. conditions. The former, he said, at that time would not be liable to duty; the other would have already borne duty. He asked us to imagine that both lorries had been simultaneously hijacked. Unless the duty were included in the normal value of both consignments, one would have identical whisky stolen at the same time with differing values.”¹⁹

So this form of hypothetical argument, involving scenarios in which the whisky was stolen and different charges were liable, was used by both sides to justify alternate conclusions, see Lord Denning’s use of the argument above.

The *eiusdem generis* rule²⁰ was proposed as a method for including excise duty in the list of items in article 24 paragraph 4. On this point Lord Roskill argued:

“I agree that in principle one should not in construing the text of an international convention, even when scheduled to a United Kingdom statute, apply the *eiusdem generis* rule, even if as a matter of construction one could find an appropriate genus. That rule of construction is a peculiarity of English law.”²¹

Regarding the use of the French text to aid interpretation Lord Roskill argues:

“Why then, in those circumstances, when one is in doubt as to the true construction of the English text of the Convention, should one deprive oneself of the assistance that is readily to hand in its French text ... I am entitled to look at the French text of the Convention ... in order to gain what assistance I can from its terms. There I find the words “*les autres frais encourus l’occasion du transport de la marchandise.*” It does not require a profound knowledge of the French language to gain assistance from these French words for they are quite general in their nature and wide in their compass and in my view quite clearly entitle the plaintiffs to recover the excise duty in question.”²²

Regarding methods of interpretation Lord Roskill argues:

¹⁹*ibid* p218

²⁰The rule applies to a list of categories and identifies a similar category as being of the same kind, and hence similar consequences apply to it.

²¹*ibid* p220

²²*ibid* p220

“It would be disastrous if our courts were to adopt constructions of such legislation different from those of other courts whose method of approach is different and far less narrow than ours merely because of over-rigid adherence to traditional - some might call them chauvinist - English methods.”²³

Lord Lawton notes that:

“It must be permissible for an English judge to show a modicum of [sic] knowledge of a European language which for some centuries in its archaic form was the language of our courts and which in more modern times has been the language of diplomacy. The words “*les autres frais encourus l’occasion du transport*” in my understanding cover more than the English translation of “other charges in respect of the carriage of the goods.” The inference which an English lawyer would normally draw from the use of the word “charges” twice in one sentence in a statute, namely, that “other charges” must bear some relation in meaning to “carriage charges” so as to bring them within the same genus if nat [sic] the same species, is dissipated by the use of different words in the French text.”²⁴

At the House of Lords²⁵ the appeal, by the defendants, was again dismissed (with Lord Edmund-Davies and Lord Fraser of Tullybelton dissenting).

The arguments for the defendants at the House of Lords case were:

“The starting point must be to look at the English text, which is English law, and at least attempt to give the words their normal, natural meaning. ... No expert evidence was given here as to the meaning of the French text to Frenchmen.”²⁶

“If there be an ambiguity, even in a case such as this where the English and French texts are equally authoritative, the court can take advantage of the French text. There is no such ambiguity here, and therefore no need to have regard to the French text. We have no expert evidence in this case. Schoolboy French is all very well, but one is here dealing with very subtle shades of meaning. It is no great compliment to the official translators of texts to say that the English is not an accurate reflection of the French.”²⁷

“The export market price is net of export duty. There may be more than one export market. The plaintiffs’ documents show that they differentiate between E.E.C., United States and other markets.”²⁸

²³*ibid* p221

²⁴*ibid* p223

²⁵James Buchanan & Co. Ltd. v Babco Forwarding & Shipping (U.K.) Ltd., [1978] A.C. 141

²⁶*ibid* p144

²⁷*ibid* p145

²⁸*ibid* p145

“The legal concept of “available market” has no narrow or technical meaning. It implies simply that the seller can dispose of the goods”²⁹

The plaintiffs argued:

“It was not the intention of the draftsman in the Convention that in a case of an individual sale the task of the court would be to look at the country to which the goods were going and then at the f.o.b. price and then deduct something for carriage and say: that is the value at the warehouse. Goods at a particular place have a single “normal value.””³⁰

“As an approach to construction, it is clear from *Stag Line Ltd. v. Foscolo, Mango and Co. Ltd.* [1932] A.C. 328 that a court can have in mind that the convention is designed to create uniformity in the law between the contracting states, and that it will construe the Act with that object in mind, which may lead to what is sometimes described as the “broad approach.” ... On the facts of the present case, the English text is enough for the plaintiffs’ purposes. As to expert evidence, if the court is invited to look at the French text, it is not obliged to have expert evidence as to the French language. It is open to the court to apply its knowledge of the French language, although it might decline to do so if the suggested nuance were so slight that it considered it dangerous to do so.”³¹

Lord Wilburforce noted:

“In a case, such as I think the present is, when one is dealing with a nuanced expression, a dictionary will not assist and reference to an expert might also be unhelpful, for the expert would have to direct his evidence to a two-text situation rather than simply to the meaning of words in his own language, so that he would be in the same difficulty as the court.”³²

“But a decision of the Amsterdam Arrondissementsrechtbank (3rd Chamber) of March 30, 1977, *British-American Tobacco (Nederland) B.V. v. van Swieten B.V.* (unreported), in a matter concerning excise duty, on facts very similar to the present, decided against the carriers’ liability ... The court took the French text “encourus l’occasion du transport” into consideration but thought that the case was basically different from one where customs duties - on passing a frontier - were concerned, which case could be covered by the phrase. These cases show that there is no universal wisdom available across the Channel upon which our insular minds can draw.”³³

²⁹*ibid* p145

³⁰*ibid* p147

³¹*ibid* p148

³²*ibid* p152

³³*ibid* p154

“I find that the judgement of Master Jacob carries conviction. ... “In respect of” is wide enough to include the way in which the goods were carried, miscarried or lost.”³⁴

Viscount Dilhorne argued:

“I am not competent to speak on European methods of interpreting legislation but I know of no authority for the proposition that one consequence of this country joining the European Economic Community is that the courts of this country should now abandon principles as to construction long established in our law.”³⁵

“What charges were the words “other charges incurred in respect of the carriage of the goods” intended to cover? They must be charges not covered by the words “carriage charges.” They must be given a wider meaning than charges for carriage. If one gives a strict interpretation to the words “in respect of,” “charges in respect of carriage” has the same meaning as charges for carriage, i.e., carriage charges.”³⁶

“If “in respect of” is given the broad interpretation of “in consequence of,” content can be given to the words in question. They will clearly cover a far wider ambit than carriage charges.”³⁷

Lord Salmon argued:

“For a court to construe a statute is one thing but to graft a provision on to it on the ground that the court thinks it is reasonable to do so would bring the law into chaos and also introduce a like chaos into the business of international carriers and those who contract with them to carry goods by land from one country to another. For the courts to graft a provision on to a statute or a contract is a practice which is entirely foreign to our jurisprudence and, as far as I know, to any other.”³⁸

The following argument of Lord Salmon’s is particularly interesting as it invokes ideas of justice and commonsense to justify the inclusion of excise duty.

“Reason and justice seem to demand that the burden of paying the £30,000 should rest on the shoulders of the carriers rather than on those of the innocent exporters.”³⁹

³⁴*ibid* p154

³⁵*ibid* p157

³⁶*ibid* p157

³⁷*ibid* p158

³⁸*ibid* p160

³⁹*ibid* p160

Again, the following argument of Lord Salmon's is interesting, this time since it portrays a particular vision of English justice, which may be biasing Lord Salmon's reasoning.

“It is perhaps worth mentioning that foreign traders all over the world having no connection with this country have been constantly entering into contracts which provide that they are to be governed by English law and that any disputes that may arise under them are to be settled by arbitration in London or by our commercial court. It would seem that our system of administering justice enjoys considerable confidence abroad and that we can safely leave our courts to apply their own methods of interpreting the Convention until such time, if ever, as better methods are devised abroad and universally accepted.”⁴⁰

Lord Edmund-Davies argued that:

“But where there is no “gap” or ambiguity in the English text (as I hold the position to be in the present case), the literal approach is preferable to the schematic and teleological approach which Lord Denning favoured”⁴¹

The above argumentation demonstrates that there was some disagreement about both: the meaning of the French text, and the methods that should be used to interpret the French text. There appears to be no legal basis for the decisions taken by the Judges, in the sense that there is no prior law that answers these questions, and so the Judges must recourse to their biases about how the law should be conducted.

We are faced with several choices and combinations:

- Is there a gap in the English statute?
- Should the case be solved on the basis of the English law alone?
- Can we use the French version of the statute as an interpretative aid? If so, should the English court interpret the French statute in the same manner as an English statute? Or should the court interpret the French statute in the manner a French court would.
- Can we use Foreign court decisions that interpret the French Convention as an interpretative aid? If so, should we interpret them the way an English court uses precedent or the the way the Foreign court uses precedent?

⁴⁰*ibid* p162

⁴¹*ibid* p168

Also of note is the language used to discuss the claims about the interpretation of the law. Judges talk of a “gap” in the law, or of terms being given a “narrow”, “technical” or “broad” meaning.

3.2.2.2 Ontology evolution

There were two changes to the ontology of the Act. First, before the case it was thought that an object could only have one “market price”. As a result of the case “current market price” could have at least two distinct values: the market price for a domestic market and the market price for an export market. Second, the phrase “other charges incurred in respect of carriage” was interpreted as including excise duty.

The broader change to the law was the change to interpretation methods for foreign statutes, with the House of Lords judges rejecting a claim that we should be using European methods of interpretation for the Act. Similiar to the case of *Popov v. Hayashi*, the *Buchanan v. Babco* case raises the issue if ontology evolution in the law can happen through the integration of external sets of rules. In both cases, the arguments in favour of the integration involve operations similiar to those used in ontology matching: two different conceptual schemes are integrated by identifying their connections.

We have tried using the inconstancy repair plan (see section 2.2.3) to model the change to the meaning of current market price. In our application the variad is the destination of the whisky and the inconstancy is the current market price function.

Trigger:

$$\begin{aligned} O_{UK}(dest(whisky1) = UK) &\vdash cmp(whisky1) = sp(whisky) + duty(whisky, UK) \\ O_{Iran}(dest(whisky2) = Iran) &\vdash cmp(whisky2) = sp(whisky) \\ O_{Act} &\vdash \lambda Item. cmp(Item) = \lambda Item. sp(good(Item)) \end{aligned}$$

Repair:

$$v(cmp) ::= \lambda Item. sp'(good(Item)) + duty(good(Item), dest(Item))$$

In this application of the inconstancy repair plan (O_{UK}) and (O_{Iran}) represent distinct sensory ontologies in which a consignment of whisky, denoted $whisky1$ and $whisky2$ for the respective scenarios, is being sent to a UK and Iranian market respectively. In

the UK sensory ontology the current market price is observed to be the sale price of the good *whisky*, denoted $sp(\textit{whisky})$, plus the excise duty, denoted $duty(\textit{whisky})$, and in the Iranian sensory ontology the current market price is observed to be just the sale price of the good.

The O_{Act} ontology represents the Act of Parliament, which acts as the theoretical ontology in our application. In this ontology the current market price of an item is defined as being the sale price of the good that the item is an instance of. Note that this is not what the Act actually states, but is a necessary simplification of the situation in order to construct a basic model using ontology repair plans.

The repair is to replace the old *cmp* function, whose observed value varies between the UK and Iranian ontologies, with a new *cmp* function, whose value matches the observed value between these ontologies. The repair specifies the new definition of the *cmp* function, in which the current market price of a good depends upon its destination. Note that, our application omits describing the complete set of changes to the ontologies as the application is intended as a simple demonstration of the use of ontology repair plans to our work.

It is unclear what relevance the repair plan has to the reasoning in the case. The problem for us is that the discourse of the case represents the explicit arguments made about the meaning of the concepts in the legislation, whereas repair plans model the implicit reasoning that suggests a change to the meaning of a concept.

Some process similar to a repair plan may well be occurring in the lawyers' minds when they suggest a modification of the concept of current market price, but this process is only a small part of the ontology evolution in the case. What is more important are the arguments made to state that there is a problem, e.g. the hypothetical scenarios suggested above, and the knowledge used to make the repair, such as the knowledge that there was a distinction between domestic and export market. We have hand-coded that knowledge into the above application of the repair plan, but we wouldn't want to do this for a general model of ontology evolution in legal cases.

3.2.3 Corocraft v. Pan American

The case of *Corocraft Ltd. and another v. Pan-American World Airways, Inc.*⁴² also involved ambiguity in an international treaty. Sub-paragraph (i) of article 8 in the Warsaw Convention stated that an air consignment note shall contain, among other things,

⁴²Corocraft Ltd. v Pan American Airways Inc [1969] 1 Q.B. 616

the following “le poids, la quantity, le volume ou les dimensions de la marchandise”. The United Kingdom Carriage by Air Act 1932, which enacts the Warsaw Convention in UK law, states that the air consignment note must contain “the weight, the quantity and the volume or dimensions of the goods”. The French version has an ambiguity over the “ou” which could be interpreted as demanding either 1 or 3 of the 4 items. The English version clearly demands that 3 of the 4 items must be present.

The facts of the case are as follows: Corocraft had hired Pan-American to transport a carton of jewellery between New York and London; The carton was stolen by an employee of Pan-American whilst being transported; The air consignment note stated that weight and the quantity of the carton, but omitted the volume or dimensions.

Corocraft was suing Pan-American for the full value of the jewellery. However, Pan-American claimed that they had limited liability under the Warsaw convention. Corocraft argued that since the air consignment note was not complete they did not have limited liability in this case.

The issue in the case was how to resolve the ambiguity in the French version, and whether this was necessary. The law in the UK was the Act of Parliament, not the French version of the Warsaw Convention. However, the Judges agreed that the intention of Parliament had been to enact the French version, so the removal of the ambiguity in the French version was a translator’s gloss and should not be considered the law of the UK.

We note the difference between *Buchanan v. Babco* and this case: in *Babco*, an external ontology was suggested to resolve apparent gaps or ambiguities in the target ontology, English law. Here the ambiguity is created by looking at the external ontology. Only when looking at the French law do we realise that our own conceptualisation may be problematic.

The interpretation of the French version was based upon the commercial context of the Convention. The inclusion of details on the consignment note was designed to prevent loss of the goods in transit. However in this case the goods were stolen, and it is unlikely that the theft would have been affected by the inclusion of volume or dimension details on the consignment note. The volume and dimension of the package could also be inferred from the weight, so it would be unnecessary to include them on the note. So the ambiguous French version was interpreted as only requiring one of the 4 details.

The case demonstrates how the intentions of a law can be used to guide its interpretation. Here the intentions of Parliament are inferred from the context of the Act,

namely that it was an enactment of an international treaty with an authoritative French version.

The case also demonstrates the role of context, and commonsense, in supplying the information needed to resolve the ambiguity. The Convention is a treaty created to ensure consistent rules regarding international air transport amongst its signatories. The focus of the treaty is upon creating conditions for trustworthy transactions that benefit the world economy. The absence of a volume or dimension details from the air consignment bill doesn't affect these aims, so interpreting the ambiguity as only demanding 1 out of 4 is in line with these aims.

3.2.3.1 Argumentation

The following facts were agreed by both parties⁴³:

1. "The first plaintiffs and Verit Jewels Ltd. knew the volume and dimensions of the carton at all material times and the defendants knew the same as soon as the carton was received for carriage."
2. "The approximate volume and dimensions of the carton could be inferred from the information contained in the air waybill. The approximate dimensions of the carton were 10-12 inches long; 6 inches wide; 8-10 inches deep."
3. "The omission from the air waybill of the volume and dimensions of the carton did not cause or contribute to the loss."
4. "If the air waybill had stated the volume and dimensions (or either of them) of the goods, the defendants would not thereby have been caused to charge a different rate of freight for the goods, or to alter the mode of carriage or custody."
5. "The carriage was subject to the Warsaw Convention as scheduled in the Carriage by Air Act, 1932 , and the carton was lost during the carriage by air."
6. "The carton arrived at London Airport on about May 12, 1962, and was kept thereafter in a safe at the defendants' premises at the airport pending customs clearance. On June 20, 1962, the carton was taken from the safe for clearance. H.M. Customs officials did not complete clearance on that day and during the night of June 20-21, 1962, the carton was stolen from the defendants' customs clearance office by one Edward Francis Cahillane, a servant of the defendants."

⁴³*ibid* p619

There was an issue regarding a conflict of laws, since English law traditionally did not require the air waybill to explicitly note the volume or dimensions. The plaintiffs (Corocraft) argued there was no conflict of laws, and the defendants argued that even if there was, under both traditional English law and the Act, they weren't required to explicitly note the volume or dimensions.

The conflict of laws is not relevant to the ontology evolution, but it is an example of how lawyers will argue around a potential conflict of laws. The plaintiffs argued against the conflict existing, since it would potentially negatively affect them if there was a conflict. The defendants argued that regardless of what law held it had the same consequences.

Regarding the ambiguity in the French law, the plaintiffs argued that "One might be able to look to the French in order to resolve an ambiguity but not to create one"⁴⁴.

Further, regarding the idea that the Act should be interpreted so as to ensure a uniform interpretation, i.e. by interpreting the French rather than the English text; they argued that "research shows there is no comity"⁴⁵, and cited cases demonstrating that the Convention already had different interpretations in different jurisdictions.

The defendants argued that "on its true construction the Act does give effect to the Convention"⁴⁶. On appeal, the defendants added "in the original French on its true construction required the air consignment note to state only such of the said matters as might be relevant or as might be necessary in the particular case"⁴⁷, which they regarded as not including the volume or dimensions of the package.

The initial judge for the case, Donaldson, consulted two French lawyers regarding the meaning of the French text. Unfortunately they disagreed about the interpretation of the French text.

At the Court of Appeal, Lord Denning argued that "It was plainly the intention of all the parties to the Convention that the French text should be the one official and authorised text; and it was plainly the intention of the English Parliament to give effect to that French text by making an exact translation of it into English ... In order to produce an exact translation, the translator should reproduce the French text faithfully, with all its defects, deficiencies, ambiguities and uncertainties."⁴⁸.

⁴⁴*ibid* p620

⁴⁵*ibid* p620

⁴⁶*ibid* p621

⁴⁷*ibid* p621

⁴⁸*ibid* p652

However, it is never particularly clear what the evidence is to support these arguments. It seems as though the evidence is based upon commonsense knowledge about how Parliament and the Law operate, which led to the belief that the Law should really be what the French text states.

Lord Denning further argues that “Seeing that the French text is ambiguous and uncertain, I should have thought that it should be interpreted so as to make good sense amongst commercial men.”⁴⁹, where good sense here means to only demand the information that is necessary to identify the package, i.e. the volume and dimension are unnecessary.

Again, it is not clear what the evidence for this argument is, the only plausible candidate is commonsense knowledge about the concerns of international trade.

3.2.3.2 **Ontology evolution**

The concepts in the Act have not changed with the outcome of the case. The laws have changed slightly though.

The English law prior to the case stated that the following details were required:

$$weight(X) \wedge quantity(X) \wedge (volume(X) \vee dimensions(X))$$

However, after the case it is not clear what the law states, since the ambiguous French text is now the basis for the law. The required items on the air waybill seem to depend upon the pragmatic concerns of transporting a particular package. If the package’s volume or dimension can be guessed from the weight and quantity of the package, they are unnecessary. So we can’t accurately model the change to the law without formalising the context of international trade.

Lord Denning’s approach to interpretation, namely the use of perceived intentions of Parliament and commercial interests to justify an interpretation of the law, is interesting, but it isn’t clear how we can model this change of meta-legal rules. The interpretation principles in English law have always been somewhat vague (see section 4.1.5), so it isn’t clear how we can model this change.

3.2.4 **Fothergill v. Monarch Airlines**

The case of *Fothergill v Monarch Airlines Ltd*⁵⁰ is another case involving interpretation of an international treaty, in this case over the interpretation of “damage” in the Warsaw

⁴⁹*ibid* p652

⁵⁰*Fothergill v Monarch Airlines Ltd*, [1981] A.C. 251

Convention. More specifically, article 22 (2) of the Carriage by Air Act provides:

“In the case of damage, the person entitled to delivery must complain to the carrier forthwith after the discovery of the damage, and, at the latest, within seven days from the date of receipt in the case of baggage ...”

Fothergill had flown from London to Rome on Monarch Airlines. When Fothergill went to claim his baggage after the flight he noticed that the suitcase had been damaged, and so filled out a form to claim compensation. Under the terms of his ticket he had 7 days to claim for any damage to his luggage. After 7 days had passed, he noticed that there were some items missing from the suitcase. He subsequently tried to add the cost of these items to his previous claim.

Fothergill’s argument depended upon claiming that the loss of the items was not “damage” under the Warsaw Convention. The civil courts and then the Court of Appeal accepted this argument, but the House of Lords overturned these rulings and decided that “damage” included loss of goods.

As in *Buchanan v. Babco* the issue at stake was not so much the actual meaning of “damage” but the acceptable methods of interpretation. The Judges accepted the argument that although within the English legal context damage and loss are different things, since the Warsaw Convention is an international treaty a different interpretation should be applied. The Convention was vague about whether “damage” should include or exclude loss, but the preparatory work for the Convention indicated that “damage” should include loss. So “damage” was interpreted as including loss.

The case is notable for the use of preparatory work for the Convention as an aid to interpretation and the use of a purposive interpretation of the Convention. Since the 7 day limit had been intended to give the airlines a time limit for action against them it was reasonable to interpret “damage” as including loss, otherwise airlines could be liable for a greater time limit than the 7 days.

3.2.4.1 Argumentation

In the lower court proceedings⁵¹ the plaintiffs argued that: “Damage to baggage in article 26 (2) of the Convention means physical and not economic damage and does not include loss of the contents of registered baggage.”⁵²

Regarding the use of the French text the plaintiffs argued that:

⁵¹*Fothergill v Monarch Airlines*, [1978] Q.B. 108

⁵²*ibid* p110

“Although under the terms of section 1 (2) of the Carriage by Air Act 1961 the French text is to prevail in cases of inconsistency, there is no inconsistency or ambiguity here, as the ordinary meaning of damage is clear. If however the court is in doubt as to its meaning, then the meaning of its equivalent “avarie” in the French text is a question of fact and should be proved by expert evidence. The court should call an interpreter or a French lawyer. Dictionaries should not be used and travaux préparatoires are inadmissible as aids to construction”⁵³

The plaintiffs cited the case of *Porter v. Freudenberg* to justify their claim that travaux préparatoires⁵⁴ are inadmissible, and the case of *Rustenburg Platinum Mines Ltd v South African* to justify their claim that the meaning of “avarie” in the French text is a question of fact.

The defendants claimed that “In case of inconsistency between the English and French text, the French text prevails”⁵⁵ citing *Corocraft v. Pan-American Airways* (see section 3.2.3), and “Avarie” in the French text can be interpreted with the aid of French-English judicial and ordinary dictionaries”⁵⁶ citing *James Buchanan & Co. Ltd. v. Babco Forwarding and Shipping (U.K.) Ltd.* (see section 3.2.2).

Regarding the acceptable methods of interpretation the defendants argued “The Convention should be construed as it would be in other countries, applying their principles of statutory construction. The court is therefore entitled to have regard to travaux préparatoires and to foreign decisions, text books and commentaries, unless the meaning of the Convention is obvious.”⁵⁷ They cited the cases of *Post Office v. Estuary Radio Ltd.* and *Ulster-Swift v. Taunton Meat Haulage Ltd.* as justification.

Judge Kerr notes that “I understand the action is in nature of a test case to resolve the difference between the insurance market and the airlines about the effect of the Convention in cases of loss of the contents of passenger’s registered baggage during the carriage.”⁵⁸ This observation indicates that the judiciary are aware of ulterior motives held by the parties in the case, relevant to us as it indicates potential biases in the reasoning.

Regarding the interpretation of “damage” as including loss of contents Judge Kerr noted:

“ I can only think of one line of argument ... though it appears to me to

⁵³*ibid* p110

⁵⁴Preparatory work, documents created whilst drafting the treaty.

⁵⁵*ibid* p110

⁵⁶*ibid* p110

⁵⁷*ibid* p111

⁵⁸*ibid* p111

be highly legalistic. It is common ground that unless the carrier can bring himself within some exempting provision he is liable (subject to limitation) for the loss of contents of registered baggage during the carriage by air. The subject matter of the contract of carriage is the baggage, not its particular contents, about which the carrier knows nothing. Accordingly, before construing article 26, one must have regard to article 18 (1) since this is the provision under which any liability for loss of contents must arise. This refers to “destruction or loss of, or of damage to, any registered baggage...” A case of loss of the contents cannot be loss of the baggage, so that it would follow that the draftsman evidently intended “damage to any registered baggage” to cover a loss of the contents. On this basis it could therefore be argued that the word “damage” must have the same wide (and to my mind unnatural) meaning in article 26.”⁵⁹

Judge Kerr instead argued for an ordinary interpretation of “damage”, arguing that: “Damage” in article 26 (2) is used in the sense of physical injury and must mean damage to the baggage ... The language illustrates the ordinary meaning of article 26 (2) ... This interpretation is reinforced to some extent by the words “in good condition” in paragraph (1) of article 26. These appear to refer to the external condition of the baggage and not to anything which may have been pilfered from it without any externally visible sign.”⁶⁰

At the House of Lords⁶¹, where the ruling in the lower court was overturned, Lord Wilberforce argues “in an English legal context loss is one thing, damage another. But the nature of the text in question does not suggest that it was drafted with strict English meanings in mind.”⁶²

Regarding the use of travaux préparatoires, Lord Diplock notes

“I think the case is one where it is right to have recourse to the minutes of the conference at The Hague ... This said, I do not myself derive any great assistance from this source. With some personal experience of international conferences of this kind, I should not attach any great significance to the fact that two delegates in withdrawing an amendment to article 26 which would have included in the article an express reference to partial loss as well as to damage ... Machiavellism is not extinct at international conferences.”⁶³

Lord Fraser of Tullybelton notes that “Professor Emmanuel du Pontavice, Professor of the Faculty of Law and Economic Science of Nantes, in an article entitled “Air

⁵⁹*ibid* p114

⁶⁰*ibid* p114

⁶¹Fothergill v Monarch Airlines Ltd, [1981] A.C. 251

⁶²*ibid* p273, this quote is of interest to us for its explicit recognition of contexts, with the idea of an English legal context.

⁶³*ibid* p283

Law,” published in the *Revue Trimestrielle de Droit Commercial*, vol. XXI, referred to a decision by the Federal Chamber of Buenos Aires that a partial loss by theft constituted an *avarie*, in the sense of article 26 (2)”⁶⁴.

Lord Scarman notes “We know that in the great majority of the contracting states the legislative history, the “travaux préparatoires,” the international case law (“la jurisprudence”) and the writings of jurists (“la doctrine”) would be admissible as aids to the interpretation of the Convention. We know also that such sources would be used in the practice of public international law.”⁶⁵.

3.2.4.2 Ontology evolution

The change to the Act was based upon the meaning of the word “damage”. The question was whether the definition of “damage” included partial loss of items in the baggage.

So we have a partial formalisation of article 18 (1) as:

$$\begin{aligned} \text{registered_baggage}(X) \wedge \text{damage}(X) &\rightarrow \text{liable}(\text{carrier}, X) \\ \text{damage}(X) &\leftrightarrow \text{destruction_of}(X) \vee \\ &\quad \text{loss_of}(X) \vee \\ &\quad \text{damage_to}(X) \end{aligned}$$

We can also partially formalise article 26 (2) as:

$$\text{damage}(\text{bag}) \rightarrow \text{notify_in_seven_days}(\text{owner}(\text{bag}), \text{carrier})$$

Note that we have had to use quite ad-hoc predicates, such as *notify_in_seven_days*, in order to formalise the article without having to use a more complex temporal representation.

The change to the legislation is that “damage” includes the partial loss of baggage. So we can change the formalisation of article 18 (1) to:

⁶⁴*ibid* p287

⁶⁵*ibid* p294

$$\begin{aligned}
\text{registered_baggage}(X) \wedge \text{damage}(X) &\rightarrow \text{liable}(\text{carrier}, X) \\
\text{damage}(X) &\leftrightarrow \text{destruction_of}(X) \vee \\
&\quad \text{loss_of}(X) \vee \\
&\quad \text{damage_to}(X) \vee \\
&\quad \text{partial_loss_of}(X)
\end{aligned}$$

However, again the major questions in the case were the methods of interpretation, and we again have problems representing these. The arguments regarding the use of travaux préparatoires rely upon quite implicit warrants, such as the use of expert opinion.

3.3 Features of Legal Reasoning

We discussed in section 3.1.1 two distinct kinds of uncertainty that can occur in the law, inconsistent and incomplete law. The cases that we have looked at have involved both problems. The Popov v. Hayashi case was a clear case of incomplete law, in which the concept of possession was incomplete. The other cases involve elements of both incomplete and inconsistent law. In these cases there are both foreign and domestic sources of law that are relevant to the case, so we have multiple sources of relevant law, but these laws are themselves incomplete, the terms they use are unclear. These cases demonstrate that inconsistent and incomplete laws are not mutually exclusive occurrences and that both are relevant causes of the ontology evolution in the law.

We have identified the following features of ontology evolution in law: the use of context to provide resources for interpretation, the incremental nature of ontology evolution, judicial discretion, linguistic ambiguities, and vague concepts.

The use of contextual features as aids to interpretation can be clearly seen in the example cases. In Popov v. Hayashi the context of baseball was used to help define the concept of legal possession over the baseball. And, in Buchanan v. Babco the context of the French legal system featured in the arguments about the acceptable methods of interpretation. However, it is not clear exactly what context is. There are surrounding circumstances to a case, which are the context of a case, but there are also contexts for different domains (such as Baseball, or French Law). In section 6.2 we will study the concept of context and its relevance for ontology evolution.

The incremental nature of ontology evolution can be seen in the cases cited in our example cases. In particular the *Buchanan v. Babco*, *Corocraft v. Pan-American* and *Fothergill v. Monarch Airlines* are interrelated. Along with other cases they have helped to define how international law is interpreted and applied in the UK courts. Collectively these cases define an extensive addition to the ontology of UK law. However, modelling the changes to the interpretation methods seems particularly difficult.

The argumentation in the cases reveals that judges have discretion in the decisions they make in hard cases. When the law is uncertain how can judges determine what it is? They have the ability to decide how a question of law should be resolved. Of course, they can't ignore valid arguments, but they can choose which arguments to accept.

Language plays a central role to the functioning of the law and hence its evolution. All of the cases involved some problem with interpreting language, whether it be a grammatical ambiguity or a vague concept. Much of the argumentation in the case seems to concern the definitions of various words and what methods can be used to determine their meanings.

These problems of interpretation are caused by the problem of incomplete law we discussed in section 3.1.1. The next chapter will study the role of language in law in more detail.

We have also seen the following features of legal reasoning that are relevant to our work: argumentation, commonsense reasoning, rhetoric, and value judgements.

All of the reasoning is based upon arguments regarding questions raised by a case. In the cases we are interested in, the questions are about law, both parties agree upon the facts of the case. The lawyers representing the different parties offer arguments supporting different answers to the questions raised by the case.

The arguments are based upon ontologies of the domain, in particular the arguments are implicitly dependent upon commonsense knowledge that is not formally part of the law. For instance, in *Corocraft v. Pan American* various presumptions about the concerns of international trade seem to be being made.

The argumentation doesn't clearly map onto any particular structure and much of the text of the legal decisions seems more like the judge is thinking out loud than constructing a formal argument. Fitting the ideas in the discussion into a particular argument structure is difficult, as it is not clear what the claims, warrants, and evidence are (c.f. section 5.4). In particular, the warrants seem heavily dependent upon the legal context of the case.

Several arguments appeal to values rather than facts in order to justify their conclusions. The problem in these cases is that there are no facts about what the law is that can resolve the case. The only recourse is to appeal to our sense of what the law should be.

3.4 Summary

The cases involving interpretation of international law were chosen to demonstrate that even within a domain of law that is largely abstracted from commonsense knowledge about real-world situations it is still difficult to construct formal models. The abstract constructions of international law are still dependent upon our implicit commonsense conceptualisations. Any difficult legal case will raise similar issues to the above cases.

The cases demonstrate a variety of different features of ontology evolution in the legal domain which we will analyse in the forthcoming chapters. Chapter 4 looks at the role of language in the law and its relevance for modelling ontology evolution. Chapter 5 looks at existing research in AI and Law, in particular work on modelling legal argumentation, and assesses its relevance for modelling ontology evolution. Chapter 6 surveys contextual logics and commonsense knowledge representation and reasoning to investigate if these techniques could be used to model ontology evolution in legal cases.

Chapter 4

Meaning in Legal Language

In the previous chapter we observed the importance of language to the law. Lawyers depend upon language to communicate their conceptualisation of what the law is, or how they think it should be.

However, language must be interpreted to access its meaning. There are a variety of problems of interpretation, some of which we saw in the previous chapter. In this chapter we survey research in Philosophy, Law, Cognitive Science and Artificial Intelligence on the problems of natural language communication, and assess its relevance to our problem.

Language is used to share conceptualisations, a speaker makes an utterance to create a conceptualisation within the hearer. We must somehow use syntax and words to share these conceptualisations. The problem of understanding meaning is to understand how we associate words with meanings so that we can share conceptualisations.

The aim of this chapter is to survey our understanding of meaning so that we can understand and justify the first claim of this thesis (see section 1.3) and investigate whether the existing models of meaning can help us to model the ontology evolution in the law.

4.1 Law and Language

The Law clearly depends upon natural language to fulfil its purpose. Governments create new laws by drafting and publishing official statutes. Courts generate case proceedings detailing the arguments presented in a trial and its outcome. These documents are referenced in subsequent legal cases. Laws are communicated and applied on the basis of these documents, hence their interpretation determines what the law is.

However, language is not a perfect medium for sharing conceptualisations. There are a variety of problems that arise in our usage of language. We now look at the following common problems: ambiguity, vagueness and open texture¹.

4.1.1 Ambiguity

The term “ambiguity” is frequently used to refer to most problems involving unclear language. We use the term here to describe linguistic objects, such as words and sentences, which can be given more than one meaning. So an ambiguous word² is one which has many different senses and an ambiguous sentence is one which can be parsed in multiple ways.

For example, the sentence “I went to the bank” is ambiguous as it uses the ambiguous word³ “bank”, which could refer to either a financial bank or a river bank. Alternatively, the sentence “every man loves a woman” is ambiguous regarding the quantification over men. Does the sentence mean that there is one woman loved by every man or that for every man there is a woman that he loves?

The problem in the *Corocroft v. Pan-American* case was the ambiguity in the French version of the law. The English drafters eliminated the ambiguity from the English version by choosing one of the possible interpretations.

Whilst ambiguity is a common problem in our use of language it is not of great interest to us. In the case of ambiguity there are different well-defined meanings. The ambiguity is removed by selecting one of the well-defined meanings. The removal of the ambiguity doesn't change the logical structure of any of the underlying meanings, it just enables the lawyers to recognise one of the possible meanings as applicable.

It should be noted that ambiguity is removed from a logical formalisation of a sentence, since the ambiguity is part of the translation between the sentence and a logical representation. Modelling the reasoning about ambiguity requires representing syntax and words as part of the domain of discourse and having relations over these objects.

¹For more details on these problems, c.f. [122] and [40]

²technically called a *polysemous* word

³A funnier example is a newspaper headline reading “March planned for next August” which is ambiguous since “March” could refer to either the month (for an absurd interpretation) or to a gathering of people (presumably the intended interpretation).

4.1.2 Vagueness

Vagueness arises when we know what sense a word has, but there are unclear boundaries to its application. Vagueness is the property of a predicate that it has borderline cases, i.e. there are individuals a such that neither $P(a)$ nor $\neg P(a)$ is definitely true. For instance, the words “rich” and “tall” are vague as we know roughly what meaning they have, but it isn’t always clear whether they should apply to a particular object. Whether someone is rich or tall can depend upon the boundaries a person sets for those terms.

Vague concepts are those for which there exist objects where it is not certain whether or not the object belongs to the category defined by the concept. Typically there is an underlying measure which the object has as a feature. The vagueness arises because we don’t know in what region of the measure we can say the object belongs to the category. For instance, baldness is vague with respect to the number of hairs after which we can no longer say someone is bald, or tall is vague with respect to what height is required to be considered tall.

Vagueness has led to several problems in creating formal accounts of the meaning of terms in natural language. A well-known problem is the *sorites paradox*, also known as the paradox of the heap. Suppose that one has a heap of sand, and that one removes a grain from the heap. After this operation the heap of sand will still be a heap. However, if this operation was applied many times the heap of sand would eventually vanish. There is no single application of the operation after which we would declare that the sand no longer formed a heap. This creates a paradox as we want to declare that the composition of all these operations removes the heap, but any particular operation doesn’t remove the heap. So we have the axioms:

$$\begin{aligned} \text{heap}(n) &\rightarrow \text{heap}(n - 1) \\ \text{heap}(10^{100}) \\ &\neg \text{heap}(0) \end{aligned}$$

These axioms state that if n grains of sand form a heap then $n - 1$ grains of sand are also a heap, 10^{100} grains of sand form a heap and 0 grains of sand don’t form a heap. This theory is inconsistent as repeated application of the first axiom to $\text{heap}(10^{100})$ proves $\text{heap}(0)$.

To resolve this problem we need something other than more axioms defining the

concept. We need extensions to the reasoning mechanisms so that we don't infer that a pile of zero grains of sand is a heap.

4.1.3 Open texture

The notion of *open texture* was introduced by the philosopher Friedrich Waismann in [147]. Waismann was concerned with the problem of verification: how do we verify a logical proposition about the world to be true or false? He argued that our empirical concepts are open textured in that they do not have neat definitions that allow us to verify the statements about them by breaking down their definitions into semantic primitives that can be easily verified using sense data.

His thesis was that the categories in the ontologies which underlie our natural language are not well defined. There are real world situations where we would not be able to say if an object did or did not belong to the category. In particular, Waismann argues that we can't resolve this indeterminacy, and that there are always cases whereby a natural-language concept is vague.

Waismann's thesis conflicts with the classical model of concept definitions, which states that concepts are based upon necessary and sufficient features for belonging to a category. For example, a "cat" might be defined as an animal with whiskers and a tail. So being an animal, having whiskers and having a tail would be necessary and sufficient conditions for being a cat. It should be noted that this model of concepts is still the basis of formal ontologies.

Open texture is similar to vagueness, in that both are used to describe concepts for which there are objects which are not clearly members of the concept. Open texture, as Waismann defines it, is the property of our empirical concepts that they always have the potential for vagueness. So a term such as "cat" may appear perfectly well defined but there might still be objects about which it is uncertain whether we should call them "cat" or not. This is distinct from stereotypically vague terms such as "tall" or "rich" which we know do not match onto a specific range of height or wealth which can be used to classify them.

Waismann gives the example of a friend telling you that there is a cat in his house, but when you go to his house you find a huge cat-like creature which occupies several floors of the house. In this case, the animal has the appearance of a cat, but is vastly larger than anything we would normally call a cat. The expectations you have from your friend's utterance exclude a creature so massive.

Weissman regarded open-textured concepts as having “the possibility of vagueness”[147]. We can understand his claim by noting that open-textured concepts don’t have a specific underlying range upon which the concept is vague, but that in specific uses of the concept we can identify a range which is vague. Thus an open-textured concept is one which always has the possibility of revealing vagueness when we try to apply it in specific cases. For example, the massive cat-like creature reveals a vagueness upon the size of a “cat”.

H. L. A. Hart adopted the term to describe uncertain legal rules [61] using the, now popular, example of “no vehicles in the park”. He envisioned a scenario in which a park had a bylaw stating that people could not bring “vehicles” into the park. In this scenario the term ‘vehicle’ is open textured as there are objects, such as ‘skateboard’, for which it is not clear whether they should be classified as vehicles in this context. He distinguishes between the core of the rule, the situations to which it definitely applies, and the penumbral situations about which there is possibility for debate. He argues that legislators intended to make certain situations illegal but could not foresee all the possible situations to which the law could be applied hence the open texture of the law is an inevitability.

The problem of open texture has long been recognised within the field of AI and Law, for example McCarty discusses it in [80]. There have been various approaches to coping with the problem, which we shall examine in more detail in the next chapter.

The problem of representing and reasoning with open textured concepts has been recognised within AI. John Sowa uses the metaphor of knowledge soup [125, 124] to describe the fluidity of human understanding. Edwina Rissland points out the problem of similarity [109] and the challenges it poses for AI. Both of them point out the need for software systems to be able to reason about and adapt concepts in the same way that humans do. Similarly John McCarthy has argued that most of our concepts are what he terms “approximate” and don’t have necessary and sufficient conditions [79].

The idea of open texture is also related to Wittgenstein’s idea of “family resemblance”. In [154] he argues that our everyday concepts are defined not by the classical model of concepts but instead by the resemblance that the members bear to each other. He cites the example of games, and challenges the reader to come up with a set of necessary and sufficient conditions. For example if you propose that a game must involve more than one player you exclude a whole class of single player games (from the card game solitaire to a whole industry of computer games). If you define a game as requiring enjoyment then you ignore the ‘no pain no gain’ of competitive sports. This

argument against categories defined by necessary and sufficient conditions has been supported by studies in psychology, for instance in work of the Rosch *et al* [113].

4.1.4 Context

The above problems of language are quite regular occurrences, but most of the time we can solve them well enough to communicate our message. What mechanisms do we use?

The role of common ground in language use is widely recognised, for example see [30]. The common ground consists of the shared knowledge that we have which enables us to communicate conceptualisations of the world. This shared knowledge is part of the context of an utterance, along with the situation in which the utterance is made.

We can regard the shared knowledge as consisting of two types: linguistic and social. We have knowledge of how to interpret utterances. This knowledge may not be explicit knowledge, but rather shared mechanisms for resolving problems in interpreting utterances. For example, we know how to resolve the anaphora in the sentence “If the milk is too cold for the baby, you must boil it” so that we boil the milk rather than the baby, even if we can’t articulate the mechanisms we are using.

Social knowledge is the knowledge we have of the social contexts in which we use language. This knowledge goes beyond our knowledge of how to parse language to include knowledge about cultural and institutional contexts. The boundary between social and linguistic knowledge is not precise. In the above example we might regard the knowledge that one shouldn’t boil a baby, but it is ok to boil milk, as social (commonsense) knowledge that is combined with the mechanisms for parsing sentences. We make the distinction as it is useful for our thesis, since we are mainly interested in the social knowledge that is used to interpret the legal language, rather than the natural-language parsing mechanisms.

4.1.5 Legal Interpretation

All of the above problems manifest themselves in legal language, so legal systems have evolved a variety of mechanisms for resolving them. For instance, within English law there are three rules which define the basic methods of interpretation.

1. Literal rule: read and apply the law with its literal meaning.

2. Golden rule: apply the literal rule unless it leads to an absurd situation.
3. Mischief rule: identify what the mischief was that the law was designed to prevent and interpret the law to prevent the mischief.

However, these principles are so vague as to be of limited practical use [158]. Their mention within legal cases is instead indicative of the style of interpretation most commonly used in different periods. Traditionally English law has used the literal rule; however recently there has been a shift to interpretation based upon purpose, arguably due to the integration of European Union law into United Kingdom law.

The principles are arguably used more to justify decisions than to make them. The principles do not provide enough detail to be used to guide decision making, but they can provide a justification for a decision that has been made for other reasons.

In the UK, Acts of Parliament will typically contain interpretation sections which detail the meaning of key terms in the Act. Legal Systems make extensive usage of these meta-legal rules to dictate how the law should be interpreted in different circumstances.

Of course, formal legal rules on statutory interpretation can in themselves be vague and in need of interpretation. Of the cases we discussed, this is most pertinent in *Fothergill v. Monarch Airlines*. The Warsaw Convention is an international treaty. There are rules under international law how to interpret this type of treaty to ensure consistent application in all signatory states, but these rules on interpretation will in turn be interpreted by national courts with their rules of interpretation [44].

One interesting feature of legal interpretation, that is shared by everyday discourse interpretation, is the occurrence of metaphors in describing interpretations. Lawyers will frequently talk of an interpretation of a word being “broad” or “narrow”. It seems as though there are spatial metaphors being used to conceptualise a concept and its meaning.

It is worth noting that the task of interpretation is frequently done by juries. The meaning of ordinary words is a question of fact within English Law⁴, and so juries will interpret these words. This means that many of the problems of natural-language interpretation are only explicitly discussed when they relate to words that have a specific legal meaning that is distinct from any commonsense meaning for the word.

It should also be noted that the task of legal interpretation is different from that of interpreting an utterance in everyday discourse. Our ordinary use of language is

⁴A fact that was established with the case of ‘*Brutus v. Cozens*’ AC 854; [1972] 2 All ER 1297

based upon a speaker-hearer model, in which a speaker makes an utterance which is interpreted by a hearer. The utterance is created by the speaker to convey an intended meaning to the hearer. Legal interpretation is a form of communication between the drafters of the law and the lawyers. However, due to the collaborative nature of legal drafting the single speaker-hearer model doesn't apply to the legal domain. For instance, we are no longer dealing with the communicative intentions of a single speaker, there are several people involved in legal drafting and they may not have the same intentions. This makes the model of the interpretative situation more complex than in the single speaker-hearer case. This is relevant to us as it affects the conceptualisation that lawyers and judges have of the interpretation of the law, and the acceptable forms of arguments that can be used to justify an interpretation.

4.1.6 Legal indeterminacy

The concept of open texture was introduced by Hart into legal philosophy to explain the role of judicial discretion in applying the law. Judges need to resolve the open texture of the law in the borderline cases, such as whether skateboards should be classified as vehicles in a park bylaw.

Judges have some discretion in how they apply the law due to the problems of natural language communication and the lack of conclusive rules governing the interpretation of the law.

The study of the biases that influence judicial reasoning is called critical-legal studies. There are variety of different types of biases⁵, such as the ones we discussed briefly in section 2.3.6.

This research is of interest to us, since it studies the implicit forms of reasoning that are responsible for changes to the law, but it is difficult to incorporate this research into our work. Some of the results⁶ are based upon statistical analysis of court cases, which is difficult to incorporate into a case study of a particular legal case, and some is based upon speculation about the reasoning used in a legal case, which we would also have to formalise and include into our model. We will see later that the significant problems with constructing formal models are understanding the commonsense knowledge used in legal argumentation, which is not discussed in the critical legal studies literature.

⁵For a survey of the different schools of critical-legal studies see [16]

⁶see [158] p354-360 for a discussion of various studies into judicial bias

4.2 Philosophy of Language

The problem of interpretation arises because it is unclear what the meaning of an utterance is. This section briefly surveys the conceptual models of natural language and its meaning that philosophers have studied.

Philosophical discussion on meaning and language has existed since the beginnings of Philosophy⁷. We are interested in the work that began with the development of modern mathematical logic. The development of modern logic provided a basis for formal accounts of meaning, language and algorithms.

Gottlob Frege and Charles Sander Peirce, both independent inventors of first-order logic, each presented their own accounts of meaning and language.

Frege presented a model of meaning that gave a word a sense⁸ and a reference⁹. The sense¹⁰ of a noun phrase was the method by which it determined what the word denoted. The reference of the word was the actual object that it denoted. For instance, the noun phrase “the evening star” had a sense of being the star-like object that could be seen in the evening: we could use the sense of the term to identify the object. The reference of “the evening star” is the planet Venus. The alternative noun phrase “the morning star” has the same reference as “the evening star” but a distinct sense.

Peirce created the field of semiotics as the basis of his theories of meaning. The core idea behind semiotics is that of a meaning triangle (see figure ??) composed of three edges to represent the reference object, the symbol representing the object and the context that connected them. Peirce’s ideas have influenced some modern work on ontologies [123]. The main insight is to emphasize the idea of a context which connects a symbol with what it represents. Peirce’s ideas are similar to Wittgenstein’s language games, which we look at in section 4.2.2.

4.2.1 Logical Positivism

The logical positivists were a group of philosophers in the early twentieth century who wanted to give a formal, rigorous basis for natural language. This work was initiated after the development of modern mathematical logic and the work of the early Wittgenstein on defining the logical basis for language use.

⁷for example, Plato’s dialogue Cratylus, available at <http://classics.mit.edu/Plato/cratylus.html>

⁸what is also termed the *intension* of a concept definition

⁹the set of possible referents of a concept are called the *extension* of the concept

¹⁰Frege’s ideas are quite subtle to explain, but only of general interest to our investigation, the interested reader should consult the references for a more detailed account

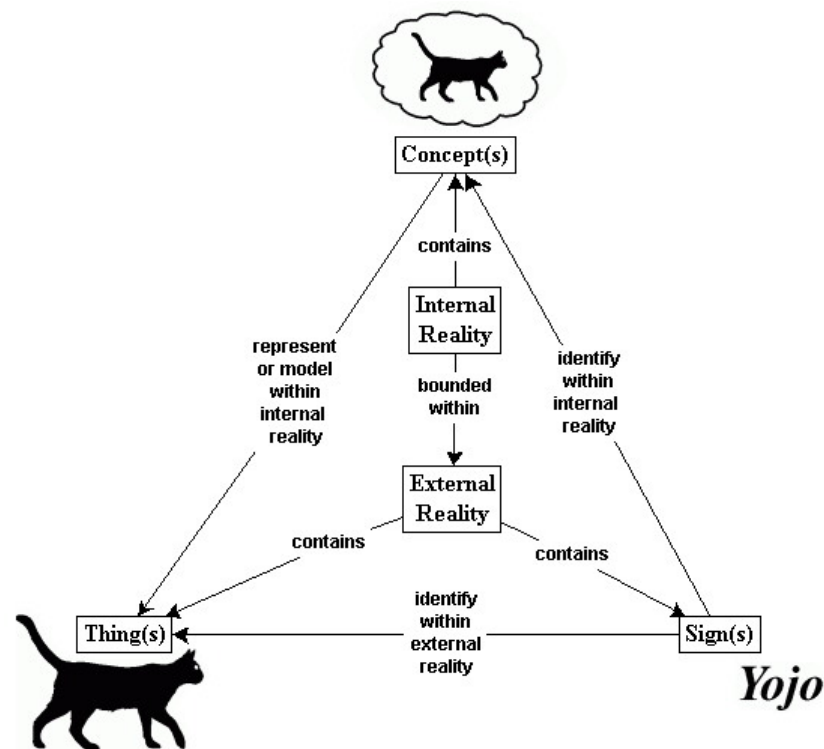


Figure 4.1: An example of Peirce's semiotic triangle. A cat called Yojo is represented by a sign, a particular instance of the word "Yojo". The sign "Yojo" and the cat are part of the external environment. The connection between them is mediated by an agent with a conceptual system which associates its concept of the cat, with the word and the cat itself.

The aim of the logical positivists was to construct logical theories that could be used to give rigorous statements about the state of the world. Rudolf Carnap led these efforts through works such as *The Logical Structure of the World* [25].

However, they discovered that it wasn't a trivial task to give formal definitions of natural language terms. Friedrich Waismann, who was a member of this group, noted the problem of open texture, see section 4.1.3. The logical positivists were the original researchers of formal ontologies and the problems they faced are relevant to contemporary research.

4.2.2 Language Games

Wittgenstein, who was an initial founder of logical positivism, later altered his views on the role of logic in language use. His initial position was expressed in *Tractatus*

Logico-Philosophicus. In this work¹¹, he argued that language has a formal basis in first-order logic statements which represent states of our physical world. Any attempts to talk about things other than the physical world are misuses of the logic of our language.

Wittgenstein introduced the idea of *language games* [154] to describe how meaning can be dependent upon a community of language users. he described communities as playing a language game in their usage of a term. This was a switch from viewing meaning in language as being based upon a logical representation of the terms to a view of meaning based upon the use of terms within a community to achieve various goals.

The *language games* approach to word meaning has influenced research by Luc Steels into language learning within a community of robotic agents [127].

4.2.3 Meaning

The problem of giving a formal foundation to the meaning of natural language terms has been actively researched for quite some time. There are two perspectives on meaning that have been discussed:

- Meaning as denotation
- Meaning as use

Formal logic provides the basis for attempts to formalise the meaning of natural language terms upon their denotations. However the problems of open-texture and vagueness demonstrate that the denotation of a term is frequently uncertain and can't be resolved simply by adding more axioms to the theories defining the term.

We can see the two different notions of meaning within the semiotic triangle. The notion of meaning as denotation is represented as the link between a symbol and its referent. The notion of meaning as use is represented by the links between the symbol and the context and the context and the referent.

Logical positivism was primarily concerned with establishing a foundations for meaning as denotation: by creating a formal language with rigorous semantics which we could use to make clear statements about the world whose meaning would be

¹¹Wittgenstein's ideas are quite subtle and tersely described, the summary here is only intended to give a very high-level description

shared. Wittgenstein's discussion of language games is about meaning as use, where the meaning of a word is determined by what effects its use has within a community.

These problems of meaning have necessarily been discussed within the research field of natural language processing, which we briefly survey next.

4.3 Natural Language Processing

Natural language processing is a subfield of Artificial Intelligence which uses algorithms to process language to produce useful information. Natural language researchers have developed a variety of techniques for resolving the ambiguity in language that arises during parsing attempts.

Of interest to us is the subfield of discourse interpretation. This field deals with some of the above problems of natural language interpretation and application.

4.3.1 Discourse Interpretation

Researchers in discourse interpretation are aiming to create mechanisms that can take an unstructured text as input and produce a representation in a formal logic of the text's meaning. The focus of the research is upon developing formalisms and mechanisms to cope with the above problems of natural language.

An exemplary problem that discourse representation research aims to solve is that of *anaphora resolution*: the problem of determining the referents of pronouns. This is the problem that was present in the "If the milk is too cool for the baby, you must boil it" example. The "it" in the sentence could refer to either the milk or the baby. The usual heuristic to resolve pronoun references, that is to interpret the pronoun as referring to the last mentioned noun, does not work in this example. If we use the usual heuristic we will interpret the "it" as referring to the baby, with disastrous consequences.

The main problems that discourse interpretation aims to solve are related to the interpretation of the syntax of language. We are mainly interested in the reasoning that occurs after the language has been interpreted into a formal representation. The discourse in the law will not make explicit mention of the mechanisms that the lawyers are using to parse the language, since the lawyers are not aware of what mechanisms they are using for this task.

An interesting approach to the problem of discourse interpretation is the use of

abduction, combined with commonsense knowledge to interpret utterances. This approach was pursued by Jerry Hobbs *et al* with the TACITUS system [64]. We will look at abduction further in the next chapter.

4.4 Meaning in Ontologies

We have seen that there are various problems with natural language communication, and that there has been much research within a variety of fields upon these problems. These problems involve the syntax of language, the semantics¹² of language, the categorisations used in language and the pragmatics of language use.

Ontologies are based upon a formal logic, and so avoid the problems of parsing and semantics, but still face the problems of categorisation and pragmatics. Ontologies are used to convey information in a manner not dissimilar to natural language. The formal semantics for most ontologies, based upon Tarskian semantics, does not always capture the intended meaning of an ontology. Ontologies are based upon the classical model of concepts, which has been shown in psychological experiments to be an inaccurate model of how people categorise [92], and Tarskian semantics, which is based upon a denotational model of meaning. The features of open-texture and language games are still present in how people use formal ontologies.

The basis of the problem is that most ontologies use natural language terms to denote the concepts in the ontology. Without these natural language labels the meaning of the terms in an ontology would be opaque, c.f. [88].

This critique of formal representations has been made within logicist AI¹³. Drew McDermott has produced several critiques of research within this field [85] [86], which have been continued by researchers within formal ontologies [151].

The solution people have for these problems is to use the context of an utterance to aid its interpretation. There have been attempts to incorporate the community of agents that uses an ontology into the representation of an ontology, such as Peter Mika's work [89]. However, this research is based upon lightweight formalisms that are not expressive enough for us to represent the evolution in meaning of a legal ontology.

What is interesting to us are the mechanisms used in natural language to negotiate a shared conceptualisation when the meaning of the language is uncertain. In these

¹²by semantics here we mean the meaning that is to be given to logical connectives, such as “and” or “or”, and how they affect the meaning of an utterance.

¹³the school of AI research that uses formal logic as a basis for AI research.

cases we appeal to the context of the use of language to determine a meaning for an utterance.

Formalising these mechanisms goes beyond simply adding more axioms to our domain ontologies and requires us to extend the logic upon which those ontologies are based to include discussion of the meaning of the concepts in the ontology.

There has been some research on dynamically aligning the ontologies of agents when communication mismatches occur, such as the ORS research. This research is of interest to us, and we shall look at some of it further in the next chapter, but we face a different problem to this research. Whereas this research aims to create mechanisms for AI agents to resolve communication problems, we face the problem of modelling the reasoning of people within a particular institutional context.

4.4.1 Authority

One important aspect of meaning in language that we haven't touched upon so far is the role of authority in determining what a word means. We have described the different perspectives regarding what the meaning of language is based upon, but we haven't discussed how specific words in a language acquire meaning. Who determines what the meaning of a word is?

We have seen that judges have some discretion in how they apply the law to a case, due to the uncertainty in the meaning of the laws in some cases and the need to resolve the uncertainty. Judges have authority over the meaning of the law within a legal case, although their judgements may be overruled by higher courts.

More generally, there are members of a community who have authority over what the concepts in an ontology mean. Hilary Putnam discusses what he terms the "linguistic division of labour" whereby some people within a community are responsible for determining the correct usage of a word [104]. For instance, in scientific communities there are usually experts upon a particular domain who know the meaning of technical terms and have authority over their definitions. Yorick Wilks had argued that it is us, rather than scientific authorities, who should determine the meaning of Semantic Web concepts [151]. The challenge for a more democratic form of determining meaning is that we still need to have mechanisms to resolve disputes about meaning.

We can reason about, and change, who has authority within a community. For instance, judges don't have absolute authority over the law, they must defer to governing bodies, such as Parliament within the English Legal system. These governing bodies

have the ability to alter the legal system, and hence alter how the law is interpreted and enforced.

The role of authority in determining meaning leads to distinctions in theories about meaning. There are theories of meaning that aim to define how meaning should be ascribed to language. The logician's approach to meaning is based upon a compositional model of meaning, and a classical model of concepts, that is at odds with how people actually use language. The logician's model is normative: it is not aimed at describing how language acquires meaning but about creating languages which have well-defined meanings.

This hints at the second claim in this thesis: *there are different types of theories of ontology evolution*. The different types of theories arise because there are different approaches to modelling the evolution of the meaning of legal concepts. We can either create a prescriptive theory, that specifies how the ontology of the law should evolve, or we can try to model how the ontology of the law actually evolves.

In the prescriptive case we are playing the role of a government, by trying to specify how the legal system should respond to uncertainty in the law. In the descriptive case we are trying to model the judicial reasoning, to understand what changes to the law a judge would accept.

4.5 Arguments about Meaning

We can now justify the first claim of this thesis:

Ontology evolution in legal cases occurs as a side effect of the arguments about meaning in a legal case.

What exactly do we mean by this claim? We start from the basis that the explicit sources of law are communicated through natural language texts. These texts contain rules which state how the law should affect certain real-world situations. These rules must be interpreted with respect to a particular real-world event, the current case in dispute. An interpretation of a rule is a claim about the meaning of the rule. There can be different interpretations of a rule which are then subject to argumentation. The argumentation process leads to claims about how the legal system should resolve this case. These claims can extend beyond just being about the meaning of the particular rules in dispute and include larger problems, such as what methods can be used to interpret the law. If the judges accept these claims then the ontology of the law evolves.

If we think of the argumentation as a procedure to resolve the current case, the ontology evolution is a side effect of the procedure. The argumentation generates the discourse about the case, which communicates the changes to the meaning of the law. The changes to the meaning are expressed in the arguments about the meaning generated by the lawyers and judges within the case.

The law evolves in a legal case because the existing law is not clear. There is a concept in case law or statutory law whose meaning is uncertain. The case is resolved by determining the meaning of the uncertain concept with respect to the facts of the case.

The arguments about meaning in a case make explicit some features of the context of the case to justify an interpretation of the law. For instance, in *Buchanan v. Babco* the interpretation methods of the French legal system, part of the context of the case, were mentioned explicitly to justify an interpretation of the law.

Our thesis is a definitional claim about ontology evolution in a legal case, based upon a conceptual analysis of the ontologies, ontology evolution and legal cases. We have provided a conceptual model for understanding how ontology evolution occurs in a legal case.

We can now also see the difference between our work and the research on repair plans in section 2.2.3. Repair plans are designed to model the implicit reasoning performed to create new theories. The reasoning about the meaning of the concepts in the theory is performed by the repair plans. The reasoning in a legal case about the meaning of the legal concepts is explicit. Lawyers must form explicit arguments about meaning which they use within the case to achieve an outcome for the case. We need to model both the mechanisms used to create new conceptualisations of the law and the mechanisms used to generate arguments about these new conceptualisations. These arguments are based upon meta-ontologies about the meaning of the ontologies in the Law.

In this chapter we have reviewed various proposals about meaning in language, but they are not immediately of any use to us in creating arguments about meaning. The Philosophical approaches to meaning are foundational: they aim to clarify or create a conceptual basis for understanding meaning in language; they are not concerned with the problem of generating arguments about meaning in any particular context. The approaches to meaning in Natural Language Processing are of interest to us, but they can't solve the problem within the legal domain, since they can't (automatically) supply the domain knowledge and reasoning mechanisms we would need to generate

arguments about meaning.

Our discussion of the different approaches to modelling meaning and the role of authority in meaning implies that there are two types of theory that we might want to construct: a prescriptive or a descriptive theory.

A prescriptive theory might use one of the existing philosophical theories of meaning as a basis for creating a model of legal interpretation which could be used to generate arguments about the interpretation of the law. The arguments produced would not necessarily bear any resemblance to the arguments produced in actual legal cases.

The problem for a descriptive theory is to formalise the models of legal interpretation that lawyers and judges actually use. This is a hard problem, the philosophical models of meaning will not necessarily help us with this task. Although there is plenty of research on the problem of categorisation and speaker-hearer communication this research doesn't address the specific institutional context in which legal interpretation occurs.

4.6 Summary

We have discussed the problems of language and its interpretation, in particular as they apply to legal reasoning. We have argued that ontologies have the same interpretation problems as natural language. And that any legal ontology will depend upon commonsense conceptualisations for its interpretation, unless we also formalise these commonsense conceptualisations.

We have noted that there are methods in natural language discourse for resolving the problems of interpretation, and that there has been some research on similar mechanisms within ontologies research.

These observations have led to our first claim: *ontology evolution in legal cases occurs as a side effect of the arguments about meaning in a legal case.*

We have also noted the different possible theories of meaning, and hence discussed our second claim: *there are different kinds of theories of ontology evolution.* The difficulties of formalising commonsense knowledge have been noted, we will return to this problem in section 6.3.

In the next chapter we will look at research in AI and Law on argumentation, and assess its application to our problem of modelling the arguments about meaning in a legal case.

Chapter 5

Legal Case Modelling in AI and Law

In the previous chapter, we formed a conceptualisation of ontology evolution as being based upon arguments about meaning. In this chapter, we want to understand the existing research on modelling legal reasoning and investigate what is useful for modelling ontology evolution.

There have been a variety of AI techniques used to model legal reasoning. We will look at the following techniques: Non-monotonic reasoning, Logic Programming, Argumentation, Argument Schemes, Case-based reasoning, Theory Construction, Coherent reasoning and Meta-level reasoning.

5.1 Non-monotonic reasoning

Most mathematical logics are monotonic in the sense that adding another proposition to a theory doesn't make anything originally provable in the theory unprovable. The problem with monotonicity for modelling reasoning is that new information can make us reject an inference we previously believed. The classic example of this is the inference from "All birds fly" and "Tweety is a bird" to "Tweety can fly". This is an acceptable everyday inference, but there are exceptions to the "All birds fly" rule. If we learn that Tweety is an exception to this rule, for instance by learning that Tweety is a penguin, then we will reject the inference.

There have been several formalisms proposed to tackle this problem. Most are based upon adding in new rules of inference to cope with these exceptions, such as Default Logic [107]. We mention this research as it is related to the argumentation research we will survey later, but it is not directly of any relevance to us. Research into argumentation theory subsumes the relevant aspects of non-monotonic reasoning.

The most interesting aspect of non-monotonic reasoning¹ is the problem that it is trying to solve. This problem, coping with the exceptions to general rules, is precisely the problem that occurs in the law, such as the example of “No Vehicles in the Park” that we discussed in section 4.1.3.

5.2 Logic Programming

Some of the earliest work in AI and Law built upon logic programming. Logic programming is a programming paradigm which identifies programs with logical theories and control flow with deduction. The most popular logic programming language is Prolog, which is based upon horn-clause logic (a subset of first-order logic).

In the 1980s one of the pioneers of logic programming, Robert Kowalski, started using Prolog to model legal rules. His research group modelled the British Nationality Act [115] as a Prolog program.

The idea behind modelling statutes as Prolog programs is that applying the act to a particular case can be done by asking a query at the interpreter. The interpreter can try to deduce the query from the program and so determine whether the law applies in that case.

For instance, a query might be *british(person101)* which asks whether the person identified by the constant *person101* is British. Facts about the person were formalised as Prolog assertions, such as *born(US, person101)*, *father(person101, person102)* and *born(UK, person102)*, and the British Nationality Act was formalised as a set of Prolog rules, such as *British(X) : -born(UK, X)*.

The problem, for us, with this approach to modelling legal reasoning is that all the work is done in formalising the facts of the case. The hard problem of matching the real-world situation to the concepts in the rules is avoided.

There are extensions to the logic programming approach which add in meta-level rules to enable reasoning about these interpretation problems. We shall look at these extensions in section 5.8.

¹Related to non-monotonic reasoning is defeasible reasoning, which discusses the same problem. The differences between these research fields are irrelevant to us, so our discussion applies to both.

5.3 Argumentation

We have already established that Legal reasoning is based upon argumentation. The parties to a legal case will present their arguments regarding the questions raised by the legal case, and will counter each other's arguments. Modelling legal reasoning as a Prolog program doesn't capture this aspect of legal reasoning. Instead, we require a representation that captures the conflicting arguments. This is what argumentation frameworks provide: a representation of conflicting arguments and methods to resolve the conflict.

The work of Phan Ming Dung provides a basis for much contemporary research in argumentation [37]. He formalised an *argumentation framework* as a pair $(AR, attacks)$ where AR is the set of arguments and *attacks* is a binary relation between arguments, which specifies which arguments are conflicting.

An argumentation framework is independent of the underlying logic used to formalise arguments. The logic simply provides well-formed formulae: the set of arguments is a set of well-formed formulae in an underlying logic. The majority of examples of argumentation frameworks, within AI and Law, use a propositional logic.

An example argumentation framework is shown in figure 5.1. This diagram shows the arguments that might be proposed by Johnny and his mother in the case of using a broom to hook a jam jar out of the pantry. The arguments both attack and support the claim that Johnny has broken the rule that he must not enter the pantry.

A semantics for an argumentation framework specifies what set of arguments should be believed given an argumentation framework, with a set of initial arguments AR . A conflict-free set of arguments is one in which no arguments belonging to the set attack another belonging to the set. An *admissible* set of arguments, $S \subseteq AR$, is conflict-free and for every argument $A \in S$ if there is an argument $B \in AR$ which attacks A then B is attacked by S (where a set attacks an argument if there is a member of the set which attacks the argument). An example of a semantics for an argumentation framework is the *preferred extension* which is a maximal subset of AR which is admissible.

A variety of semantics have been proposed for argumentation frameworks. There is some debate about how these semantics should be evaluated with a criticism made in [6] that many are assessed more on intuition (based upon examples) rather than any principled evaluation.

Aside from Dung's model, there are a variety of alternative formalisations of argumentation proposed [103] which explore reasoning with a potentially inconsistent

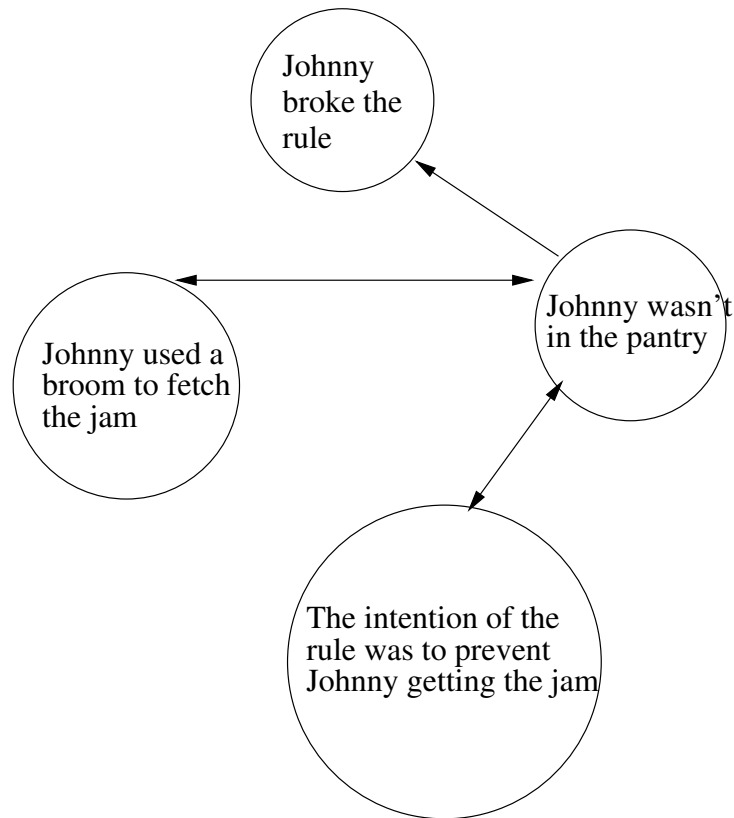


Figure 5.1: A basic argument graph for a few of the arguments proposed in the ‘Naughty Johnny’ case. The nodes are arguments about the case, and the directed edges represent instances of the attacks relation between arguments. Note that the “Johnny wasn’t in the pantry” argument also includes the proposition that Johnny not being in the pantry means that he hasn’t broken the rule. An example of an admissible set would be the set consisting of all the arguments except the “Johnny wasn’t in the pantry” argument.

knowledge base. The basic features of an argumentation system [24] are: an underlying logic, a concept of argument, a concept of conflict amongst arguments, a notion of defeat amongst arguments, and a notion of the acceptability of an argument.

Research on argumentation systems subsumes much research on non-monotonic logics. In [37] Dung shows that Reiter’s default logic can be viewed as a preferred extension over an underlying monotonic logic. In [17] Bondarenko *et al* propose assumption-based frameworks for default reasoning which generalises many prior forms of non-monotonic reasoning.

Recent work in argumentation has included the idea of an audience and their values in the argumentation framework. In [11] Trevor Bench-Capon *et al* introduce value-based argumentation frameworks defined by a triple (AF, V, n) where AF is an

argumentation framework (defined as above), V is a set of values and n is a function $n : AR \rightarrow V$. An audience defines a binary relation over values indicating which values are preferred over others. This framework has been used to analyse legal cases in [10].

Argumentation frameworks are usually based upon a foundational approach to epistemology [129], beliefs which are not attacked act as a foundation for justifying other beliefs. This is distinct from the coherentist approach taken in Paul Thagard's work on coherent reasoning (see section 5.7).

There has been some research within multiagent systems that is of interest to us. In [68] a framework for using argumentation to resolve ontology disputes between agent ontologies is introduced. This work is similar to the ORS system we discussed in section 2.2.2. The framework provides a meta-level ontology to formalise the relations between the agents' ontologies. These frameworks are certainly similar to what we would like to create to model the ontology evolution in legal reasoning. However, it isn't clear whether the meta-level ontologies used in these frameworks can be applied to our problem. We will look further at this problem in the next chapter.

5.4 Argument Schemes

The representations of argumentation we looked at in the previous section are designed for reasoning with inconsistent arguments. However, the content of the arguments is based upon theories in the underlying logic, which these frameworks are independent of. The frameworks don't describe the acceptable patterns of argumentation, as these are considered to be part of the underlying logic.

Prakken points out in [101] that Logic is too abstract for modelling legal arguments. There are recurrent patterns in legal argument that are not expressed when creating a logical representation of the claims the arguments make.

An example of the use of argument schemes is in [156] in which argument schemes for Witness Testimony, Video Tape evidence, Defeasible modus ponens and Purpose are used to analyse the Popov v. Hayashi case. The Witness Testimony argument scheme looks like:

WT: Witness says that P , therefore P .

where P is a proposition. This argument scheme is used in the case analysis to formalise the claims that Hayashi didn't assault Popov made by witnesses, and so Hayashi shouldn't be prevented from possessing the ball.

The challenge in modelling legal cases using argument schemes is to match the argument scheme to the text describing the case. It is sometimes unclear how an argument can be fitted into an argument scheme [102].

The work on argument schemes is relevant to our work as some of the arguments about the meaning of the law might follow a particular argument scheme. For example, a judge in a case with an ambiguous international treaty which had an authoritative French version might ask French lawyers to provide expert testimony regarding the meaning of the French version in French law².

It is unclear how we can use argument schemes. If there are particular argument schemes for arguments about meaning, we should use these to model our legal cases. But it is unclear exactly how arguments about meaning are conducted, and whether the arguments that affect the meaning of a word are distinct from the arguments used to justify other claims. For instance expert testimony might be used to justify a claim about meaning in the same way that it would be used to justify any other claim in a legal case, and so might not need a distinct argument scheme.

There is some discussion of argumentation about the meaning of natural language expressions in [148] [152]. However, these are philosophical analyses that rely upon implicit knowledge to analyse the arguments, and so it isn't clear how these analyses can be used to automate argument generation.

Another problem is that Argument schemes hide the underlying common-sense ontologies upon which they are based. For instance, an argument scheme about expert opinion is a black box, the argument scheme can't be deconstructed into underlying presumptions about epistemology. However, in actual legal argumentation the argument can be decomposed into common-sense propositions about the reliability of experts. The ontologies behind the argument scheme are hidden, whereas we want them to be explicitly available for reasoning.

5.5 Case-based reasoning

In common law one of the primary sources of law are the verdicts of prior legal cases. There has been much research on using case-based reasoning to model the citing of precedent legal cases as a form of legal argument.

The basis of case-based reasoning in law is to find cases which are similar to the current case and to hence argue that a similar verdict should apply to the current case.

²c.f. *Corocroft v. Pan-American*

This involves determining the *ratio decidendi*, the rule which lead to the verdict in the prior case and which is binding upon future cases in lower courts. If this rule can be argued to apply to the current case, a similar verdict must follow.

The challenge in this form of reasoning is to identify the reason in the prior case and argue that the relevant features in the prior case hold, or not, in the current case.

An exemplary Case-based Reasoner is the HYPO system developed by Kevin Ashley and Edwina Rissland [3]. HYPO was created for the domain of trade-secret law. HYPO searched for similarities between cases based upon their common collections of facts.

HYPO had a database of 30 cases. Each case was represented as a legal-case frame, which contained the factual predicates for the case, which formalised the facts of the case. Reasoning was done using *dimensions*. The facts of the case determined whether a dimension was applicable to a case. If a dimension was applicable, a subset of the case facts would further determine in what way the dimension effected the case. The dimension offered a range of possible values which reflected how strongly it effected the case and which side it benefited.

When the system is applied to a new case, as input, it checks what dimensions apply to the case and finds cases which share factors with the new one. The system then identifies the most relevant ones to the new case and generates “3-ply” arguments, i.e. an exchange of arguments about precedents for 3 steps (Plaintiff,Defendant,Plaintiff).

An example of a HYPO dimension is:

Secrets-Disclosed-Outsiders: Plaintiff strengthened the fewer disclosures to outsiders it has made of confidential information.

This dimension is formalised by specifying the facts that have to hold in the case for it to apply, facts such as “There is a corporate plaintiff”. The dimension would weaken the plaintiff’s case the greater the number of disclosures were made by the plaintiff. See [3] for more details.

The CABERET system [108] was later created to combine rule-based reasoning with case-based reasoning, specifically to tackle the problem of legal interpretation for open-textured rules. The system used case-based reasoning to determine the extension of open-textured concepts, and rule-based reasoning to determine the consequences of the open-textured rules. CABERET was applied to the domain of income-tax law, and had a database of 23 cases.

Other systems have been based upon a different underlying representation of a legal case [110] and its argumentation. The CATO system [4] (developed by Kevin Ashley

and Vincent Aleven) was based upon a representation of cases using *factors*. A factor represented a collection of facts which held in the current case and was labelled as being either pro-plaintiff or pro-defendant. The important distinction between a factor and a dimension is that a factor is unary, it's presence is a benefit to one side and there is no variation in how strongly it benefits that side, whereas a dimension can benefit either side and to a varying degree.

The challenge for case-based reasoning is to define metrics of similarity between cases [57]. The algorithms that calculate these metrics need not be based upon inference with explicit ontologies; and they can be a form of implicit knowledge. In legal reasoning, much as in everyday reasoning, we often can't explain exactly why two situations seem similar, they simply do.

A problem with using case-based reasoning is the difficulty of forming representations of legal cases, although this is a problem for all the techniques we survey here. All the above case-based reasoning systems are based upon small sets of cases whose representations have been hand-crafted.

For our purposes, we can regard case-based reasoning as a source of arguments, claims about the current case and its relationship to others. The connection between case-based reasoning and argumentation is discussed in more detail in [119].

5.6 Theory Construction

One of the pioneers of AI and Law research, Thorne McCarty, advocated a conception of legal reasoning as theory construction [83]³. McCarty argued that there is no right answer to the legal problems in hard cases. In these cases lawyers are engaged in theory construction: they are trying to take the facts of the case, and the relevant laws, and combine them into a theory that justifies their desired outcome.

This model of legal reasoning is similar to contemporary accounts in legal philosophy, such as Dworkin's model of Law in [38].

For example, Alison Chorley [29] describes the AGATHA system that constructs case-law theories. The system searches through a case database to identify cases that are similar to the current one, and also identifies counter-arguments to the similarity, using a representation of cases similar to Ashley's HYPO.

A relevant form of reasoning for us is *abduction*, the form of reasoning that gener-

³“... Thus the important process in legal reasoning is not theory application but theory construction.”
p276

ates hypotheses to explain observations. Abduction differs from the other fundamental forms of reasoning: induction, the process of inferring a generalisation from examples, and deduction, the process of inferring propositions using logical rules.

Arguments about meaning could be constructed abductively, by generating claims that explain how the rules are connected with the facts of the case.

The problem with a model of legal reasoning as theory construction is that we then need to account for why certain theories are accepted and others are rejected. This problem has been discussed in research in AI on abduction as the problem of *hypothesis evaluation*.

The problem of hypothesis evaluation can be solved using a *coherentist* approach to epistemology. We mentioned earlier that argumentation frameworks are typically based upon a foundationalist approach to epistemology. In a foundationalist epistemology a proposition is only known to be true if it can be logically deduced from a foundational set of propositions. A coherentist epistemology doesn't recognise any privileged set of propositions and instead argues that a proposition is believed if it belongs to a set of propositions that "fit" together; we shall look at what that means in the next section.

5.7 Coherent reasoning

In order to model legal reasoning as a form of theory construction there must be a method to evaluate the resulting theories. There has been discussion in legal philosophy of the use of measures of coherence to evaluate theories produced as legal solutions to a case. However, this philosophical work typically doesn't result in computational implementations. But, there has been some research into implementing computational models of measuring the coherence of a theory.

Paul Thagard's work on coherence [136] is an attempt to provide a formal basis for the various philosophical intuitions [31, 105] in favour of coherence as a basis for knowledge. Thagard's work constructs the coherence problem as one of constraint satisfaction.

Thagard's model of coherence regards the problem as one of finding a maximal coherent subset of a set of propositions with coherence links between them [138]. The propositions have weighted coherence relations between them which indicate the degree of coherence or incoherence that propositions have with each other. The constraint satisfaction problem is to find the subset of propositions which maximises the sum of

these edges.

For example, if the domain is mathematics then the links between propositions would simply be logical consistency and the maximal subset would be the largest consistent subset.

This framework is very similar to that of argumentation, in that both are based upon representations of propositions as nodes in a graph with edges indicating an epistemological connection. The distinction between them is that the coherence view takes justification to be a two way link. The basis for this is that there is a problem of using our sensory observations to justify our knowledge of the world, since our sensory observations can be false or incomplete; so we must use our background knowledge to determine what is likely to be true. Neither the background knowledge nor the sensory observations can be used as a basis to justify the other.

The primary motivation of Thagard's work has been in Cognitive Science where coherence has been proposed as a model of how people reason. This raises the question of providing a justification for why the theory with the most explanatory coherence should be regarded as the 'truth'. The view in Thagard's work is that the correspondence theory of truth⁴ is correct but that we only have limited contact with reality and hence must go on the balance of evidence to choose between conflicting hypotheses.

Thagard's work has been applied to modelling conceptual change within the sciences [135]. This has led to the formation of principles of explanatory coherence which are used to determine how propositions can 'cohere' with one another. Taken from [136] they are:

Principle E1: Symmetry. Explanatory coherence is a symmetric relation, unlike, say, conditional probability. That is, two propositions p and q cohere with each other equally.

Principle E2: Explanation. (a) A hypothesis coheres with what it explains, which can either be evidence or another hypothesis. (b) Hypotheses that together explain some other proposition cohere with each other. (c) The more hypotheses it takes to explain something, the lower the degree of coherence.

Principle E3: Analogy. Similar hypotheses that explain similar pieces of evidence cohere.

Principle E4: Data Priority. Propositions that describe the results of observations have a degree of acceptability on their own.

⁴The theory that true statements are those that correspond to states of the world.

Principle E5: Contradiction. Contradictory propositions are incoherent with each other.

Principle E6: Competition. If p and q both explain a proposition, and if p and q are not explanatorily connected, then p and q are incoherent with each other (p and q are explanatorily connected if one explains the other or if together they explain something).

Principle E7: Acceptance. The acceptability of a proposition in a system of propositions depends on its coherence with them.

Explanatory coherence has been applied to the legal domain [137] where it has been used to analyse the process of legal decision making. The models are, like argumentation, based upon formalisations of the legal case in propositional logic.

For our purposes we must adapt the reasoning towards the agent's goals. In developing a theory, lawyers are trying to justify the best outcome for their client. This motivating element in developing theories has been discussed in [135] and would correspond to biasing the coherence of a theory to ones which are coherent with our goals.

Coherence theories of law have been proposed, e.g. in [67]. These aim to describe the process of legal decision making as one of seeking coherence within the law. J.C. Hage presents a formalisation of legal coherence in [56].

In [12] Trevor Bench-Capon and Giovanni Sartor experiment with Thagard's early model of explanatory coherence [135]. This model is based around the use of neural networks to calculate the explanatory coherence of a theory. They raise a few technical questions, regarding the neural network implementation, which are not applicable to the general model of coherence as constraint satisfaction presented in [136].

The problem with applying Thagard's model is that the abstract coherence principles must be connected with the argumentation in the legal case. We still have to hand-craft the model of the case and the coherence edges between propositions. Thagard's model does not provide any mechanisms for generating theories, it only provides a method to evaluate them.

We also have the same problem as the argument schemes: the coherence principles are used to create the coherence edges, but are not part of the discourse themselves. This is not the case in the law, where the principles involved in the case will be part of the discourse.

5.8 Meta-level reasoning

Legal reasoning and discourse is recognised as being highly reflective [111]: legal rules reference other legal rules and lawyers will reflect upon their own reasoning and values in a case.

For instance, Acts of the U.K. Parliament contain interpretation sections, which state what the meaning of the concepts used in the other sections of the Act are. These are rules which dictate how other rules should be interpreted.

It is also quite common within a legal system to have rules about the priorities of various sources of law. For instance, in the U.K. new statutory-law has a higher priority than old case-law.

We mentioned in section 2.1.2.2 the concept of the domain of discourse of a logical theory. These are the objects that the logic makes assertions about. The objects usually denote things in the real-world⁵, such as a particular skateboard or a particular person. A legal rule might make an assertion about a physical state of the person, such as that the person is riding a skateboard in a particular park, and a material consequence of this state, such as a £10 fine.

The rule doesn't describe how it should be applied. That information is contained in the meta-logic of the logic that the rule belongs to. For instance if the logic is first-order logic, the meta-logic will usually be the natural language discussion of first-order inference.

The idea of meta-level reasoning is to formalise this meta-logic within another formal logic. A further extension of this idea is to formalise the meta-level reasoning in the same logic and include rules to enable inference at the meta-level to affect inference at the object level, and vice-versa.

The FOL system was one of the first AI systems to explore this idea [150]. It included "reflection principles"⁶ allowing the reification⁷ of object-level propositions into propositions in a meta-level representation⁸.

For example, we might have a rule which states that "Park by-laws don't have affect on a Sunday". This rule describes how park by-laws should be applied, and would prevent the person above from being fined for riding a skateboard through a

⁵Although abstract objects, and objects which have a social meaning, such as money, are also possible.

⁶Their usage of the term originates from the work of the logician Sol Feferman

⁷Reifying a concept or proposition means taking an assertion using the concept or proposition and making it an object in the domain of discourse, at the meta-level.

⁸c.f. [46] for an axiomatisation of meta-level reasoning.

park on the Sunday, since it would make the park by-law inapplicable.

There has been prior work on using meta-level reasoning to control legal interpretation [100] [157]. This work is based upon hierarchies of meta-levels, with rules at higher levels controlling the application of rules at lower levels. These higher-level rules are used to resolve problems with applying the lower-level rules, such as the existence of multiple inconsistent rules.

The most relevant work for us is the research by Jonas Barklund and Andreas Hamfelt [5]. This research uses meta-logic programming, an extension of the logic programming with meta-level reasoning. They represent meta-legal rules as meta-level schemas which are applied to lower-level rules in the case of conflicts or ambiguities in the law, such as the Swedish principle of “*lex posterior legi priori derogat*” which means that more recent laws override older laws. Their system checks whether a legal rule is applicable in the current case and also whether there are any relevant rules in the meta-levels which override the application of the rule.

An extension of these ideas by Andreas Hamfelt in [60] is particularly interesting to us as it includes a “Meaning” predicate, which relates a natural language expression of a law with a proposition in formal logic. However, that system requires user input to create the formalisations of the law.

Another interesting use of meta-logic is in [155] in which the authors note that argumentation frameworks are a form of meta-logic, since the argumentation framework controls what can be inferred from the underlying logical theory.

5.9 Limitations of these techniques

The above techniques are primarily aimed at two tasks: rational reconstruction of legal cases and automation of legal reasoning. The work on argumentation and coherence is usually about creating a rational reconstruction of the reasoning in a legal case. The work on logic programming, case-based reasoning and meta-level reasoning is usually about trying to automate various aspects of legal reasoning. However, the tasks are not mutually exclusive.

The aim of creating a rational reconstruction of a legal case is to improve our understanding of the reasoning in the case by creating a formal model. The modeller must take the discourse of the case and create a formal representation from it.

The representations used in a rational reconstruction are typically based upon a propositional logic and will not formalise the commonsense warrants in an argument

at a level of detail suitable for automated reasoning. The representations are hand-crafted and the systems developed do not generate arguments autonomously.

The systems that try to automate some aspect of legal reasoning are typically based upon quite small databases of legal cases, whose representations have been hand-crafted. The system expects input in a given format, which must also be hand-crafted.

The feature of legal reasoning that we are interested in modelling is the argumentation about meaning in the legal case. As we have argued in section 4.5, this argumentation depends upon a conceptualisation of what meaning is. We have also argued that the context of a case will provide the information needed to create arguments about meaning. Are the above techniques useful to us?

The problem with many of them, with the exception of meta-level reasoning, is that they are based upon a conceptual model of legal reasoning that the system can't introspect upon. This means that they can only automate the reasoning in routine cases, in which only standard reasoning is required, but not to the kinds of hard legal-cases we have been studying, in which the ontology of the law changes⁹.

For instance, the coherence model proposed by Thagard relies upon various coherence principles. These principles are used to create coherence links between propositions. It is the modeller's task to apply these principles to a particular case. The resulting coherence graph doesn't include statements about the presumptions the modeller was making when applying the principles to a case.

Also, the value-based argumentation model proposed by Bench-Capon *et al* requires the modeller to determine what values were present in the case. The argumentation framework doesn't provide a mechanism to engage in argumentation about the values, and whether they actually are present in the case.

However, our discussion of meta-level reasoning also raised the question of what knowledge and reasoning should be introspectable. We need to have a conceptualisation of meaning and this conceptualisation should interact with the reasoning, so that arguments about meaning affect what inferences can be made in the case.

The techniques we have surveyed are typically based upon *prescriptive* models of reasoning, which specify exact standards of reasoning rather than modelling how people actually reason. This feature, again, raises the issue of what kind of model of ontology evolution we are trying to construct. We will return to this question in section 7.2.

⁹c.f. Ann Gardner's thesis [45] which described a system which could identify situations in which rules and/or precedents came into conflict, and so identify that there was a hard legal case, but the system could not solve the problem.

The techniques don't seem to allow different perspectives on the case. There are models, such as AGATHA, which represent the argumentation as a game, so there is an association between the agents and the arguments they make. But these models don't provide a detailed representation for an agent's perspective on the case, so couldn't be used for a task such as crafting arguments for a particular judge.

The meta-level reasoning techniques proposed so far don't combine meta-level reasoning with object-level reasoning. They are based upon a layered approach in which the higher-level rules are used to resolve conflicts at lower levels. This doesn't seem to match the use in a legal case. For instance, in *Popov v. Hayashi*, Popov argued against the definition of possession proposed by Prof. Gray, since it would exclude his possession claim, and instead advocated another definition which would allow him possession. This reasoning combines reasoning about the meaning of a concept, a form of meta-level reasoning, with reasoning about the material outcome of the case, a form of object-level reasoning.

These limitations are not fatal to the usefulness of the techniques. However, they do require careful consideration over the scope of any application they are used in. So far we haven't narrowed the task of reasoning about ontology evolution in legal reasoning down to any particular domain of law. We have also not specified whether we are aiming to rationally reconstruct the reasoning in a legal case, or trying to automate the reasoning in a legal case.

The problem of automating the reasoning is harder, and so offers us a more general problem to consider for the next chapter. However, we shall return to the problem of the kind of theory of ontology evolution we are aiming for in chapter 7.

It is worth noting that we haven't surveyed some of the techniques used within formal logic and philosophy for modelling legal reasoning. In particular, there has been much research on deontic logics, which use a modal logic to formalise legal rules [65]. However, this research is not usually aimed at automating reasoning, and we subscribe to the AI philosophy that these features can be expressed within a more general logic [78].

5.10 Summary

We have surveyed the techniques proposed within AI and Law for modelling legal reasoning, and evaluated their relevance to ontology evolution.

We have argued that the current techniques have various limitations, such as lack of

introspection upon the reasoning mechanisms or lack of argument generation, which make them inapplicable to generating the arguments about meaning in the hard legal cases we have been studying.

We have again seen the problem of prescriptive vs. descriptive theories, and noted that we haven't clarified exactly what kind of theory or application we are trying to develop. We will consider the more general problem, of argument generation, in the next chapter and survey some AI techniques that might be applicable to our problem.

Chapter 6

Arguments about Meaning in Legal Cases

In the previous chapter we have surveyed existing research on modelling legal argumentation. In this chapter we look at the arguments about meaning that occur in a legal case.

Our claim is that *ontology evolution in a legal case occurs as a result of the arguments about meaning proposed in the case*. We further claim that the context of a legal case supplies these arguments, and that commonsense knowledge is essential to modelling these arguments.

We survey some of the existing research on contextual logics and commonsense knowledge bases, and discuss its relevance to creating computational models of ontology evolution in Law.

6.1 Arguments about Meaning

The argumentation frameworks discussed in the previous chapter are meta-logical (c.f. [155]), in the sense that they are used to control the reasoning within a logic but are independent of any particular logic. This implies that they are independent of the ontologies that those arguments are based upon, since the ontologies are theories within the logic. These argumentation frameworks have been used to model certain legal cases, but the majority of these models rely upon simple propositional representations of the domain. There is no (explicit) formal ontology that underlies these domain representations.

The arguments that particularly interest us are the arguments about meaning. These

are the arguments that make, and justify, claims about interpretations of the meaning of legal ontologies. These arguments must relate the legal ontologies to facts about the case. In order to relate these ontologies there must be ontologies about meaning. These ontologies contain rules that can link an object-level statement about the facts of the case with a meta-level statement about the meaning of a law with respect to these facts.

6.1.1 Ontologies of meaning

If we are to have arguments about meaning, those arguments must be based upon ontologies about meaning. These ontologies must describe a conceptualisation of what meaning in natural language is. But what concepts will this ontology contain, and what definitions will the concepts have?

Firstly, it is worth noting that discussion of the meaning of “meaning” is not new. As we discussed in chapter 4, problems of ascribing meaning to natural language utterances have been extensively discussed in various different branches of Philosophy and Computational Linguistics. We must clarify what form these arguments take in particular legal cases.

In the case of *Buchanan v. Babco* (see section 3.2.2) we saw that lawyers argued over the meaning of “other charges ...”. These arguments made reference to the “scheme” of the text, and whether other charges “included” excise duty.

Similarly, in the case of *Popov v. Hayashi* (see section 3.2.1) there were arguments over the meaning of “possession”, which used words such as “definition”.

Words such as “scheme” seem to denote meta-linguistic concepts which represent features of the structure of a text. The word “included” seems to refer to a subsumption relation between two categories. These concepts belong to an ontology of meaning, since they are used to make claims about the meaning of a law.

But what role does an ontology of meaning play in legal reasoning? The meta-level terms like “definition” and “includes” must be related to the object-level theories in order to be useful. For example, consider the ‘Naughty Johnny’ example case discussed in section 2.3.5.3. If Johnny can convince his mum that the meaning of “enter” does not include his use of the broom then his mum can’t conclude that he entered the pantry. The claims about the meaning of a concept will affect the inferences about appropriate legal consequences.

However, any claim must be justified. In the case of ‘Naughty Johnny’, Johnny can

argue that the definition of “enter” requires that his whole body is in the space, since this is how it is used in other contexts. So the claim about the meaning of “enter” in the case with the broom is justified by appealing to evidence of a similar usage in other contexts. There is an underlying warrant here that the usage of the concept in other contexts implies it should have a similar meaning in the current context. But why do we presume this warrant should hold?

There seems to be an underlying commonsense conceptualisation of what meaning is that supplies these warrants. So certain arguments make sense to us, but without us understanding why exactly they make sense.

6.1.2 Authority over meaning

Once we have concepts in our ontologies of meaning we need to define those concepts. For example, when exactly should a concept X be defined by some proposition P ? In answering these kinds of questions lawyers are determining the definition of the object-level concept X . We can assume that lawyers will want concepts to have a definition that serves their (client’s) interests.

For example, Johnny will interpret “enter” as meaning a state in which a body is entirely located within another, and Johnny’s mother will interpret “enter” as meaning that a part of a body is located within another. Although there are some constraints upon how these words can be interpreted since the open texture of natural language allows both of these conflicting interpretations of the borderline case. Since the borderline case was not foreseen, lawyers are free to alter the definition of concepts to suit their purposes.

However, although this story makes sense for lawyers it doesn’t resolve the problem of determining what concepts should mean; what are the axioms in an ontology of meaning? Here we need to consider the role of judges, since they have the authority to accept or reject arguments about the meaning of legal ontologies.

Even for lawyers there are constraints upon what claims about meaning they will make. Since their goal is to convince the judge that a concept has a given meaning they will not make claims about the meaning that they don’t believe the judge would accept. For example, Johnny would not try to alter the definition of “enter” to simply being any state that he wasn’t in, as this definition is unlikely to be accepted by his mother.

There is an issue of authority over meaning not just at the level of the object-

level definitions, of concepts such as enter, but over the meaning of meta-ontological concepts such as meaning. We discussed the role of authority in determining meaning in section 4.4.1.

It seems as though there are difficulties determining what the meaning of the concepts in the ontologies about meaning are, and that there different kinds of ontologies of meaning. We can either prescribe axioms and concepts for an ontology of meaning, or we can try to describe what axioms and concepts are in the ontologies of meaning that people use. We discuss this problem further in the next chapter. In section 6.3 we discuss the possible use of commonsense knowledge bases for a descriptive ontology of meaning.

6.2 Context

We have discussed ontologies of meaning as being essential to forming arguments about meaning. We have noted that finding axioms for these ontologies is hard, but necessary if we want to create arguments about the meaning of a legal term. One source of these axioms is to use the context of the case to justify a particular interpretation.

For instance, Johnny might argue for an interpretation of “enter” that is similar to that used in related contexts; similarly, in the *Popov v. Hayashi* case the definition of “possession” in the context of baseball is used to aid the interpretation of “possession” in the Law; or in *Buchanan v. Babco* the French version of the text is used to aid the interpretation of the English version of the text. In all these cases the context of the case is being used to justify an interpretation. In particular, other contexts are being used to provide evidence to justify an interpretation.

But what exactly is the context of a legal case? In this section we investigate contextual logics and their use to model the arguments in a legal case.

There has been some discussion of the role of context within AI and Law. For instance, in [55] Donald Berman and Carole Hafner present a case-based reasoning system that represents the effect later verdicts have upon the strength of a precedent case. In [9] the authors highlight the situational context of a legal case as one of the avenues for future research in AI and Law.

6.2.1 Contextual Logics

The idea of context has been used for some time in linguistics as a catch-all term for the factors that affect the meaning of a term in a particular utterance [1]. Contextual logics are an attempt to formalise the affect of context upon meaning and reasoning.

AI research on contextual logics began relatively recently: John McCarthy began the research field of context in logical AI in the 1980s [76]. Many researchers have studied contextual logics since then and there is now a sizable community, with a biennial conference¹.

The main approaches to contextual logics are those investigated at Stanford by John McCarthy and others, and at Trento by Fausto Guinchiglia and others. There are some alternative approaches to including context within a logic which we will also briefly look at.

6.2.1.1 Stanford

The problem of context is related to the problem of non-monotonicity in reasoning, which we briefly discussed in section 5.1. John McCarthy advocated the use of context to cope with exceptional cases [77] to general rules. We can create rules that apply in general contexts, and overrule them in specific contexts.

The basis of McCarthy's contextual logic is a $ist(c, p)$ predicate, which asserts that proposition p is true in context c , and linking axioms, which assert a relationship between ist predicates for different contexts².

McCarthy further advocated the idea of combining reflection with contexts and the idea of an outer context. All propositions are asserted with respect to a context, with the outer context as the containing context for all contextual assertions.

McCarthy's student Ramanathan Guha used these intuitions to form a contextual logic which became the basis for the CYC knowledge base [54]. The CYC knowledge base is the main application of this model of context and we shall investigate CYC further in section 6.3.2.

There are some unresolved theoretical issues with the McCarthy model of contexts. The ist predicate introduces a theory of truth into the logic, which potentially brings problems such as Tarski's paradox. These problems have not been fully investigated [1] and we shall look at them again later in section 6.2.3.

¹c.f. <http://mainesail.umcs.maine.edu/Context/context-conferences/>

²e.g. $ist(c_1, p_1) \rightarrow ist(c_2, p_2)$

6.2.1.2 Trento

Fausto Guinchiglia had some similiar intuitions to McCarthy regarding the importance of context to enabling robust reasoning in AI, these intuitions being developed with his involvement in the FOL project [150]. This led Giunchiglia to propose his own model of context [48], which has since been developed by himself and other researchers, mainly at Trento University.

The Giunchiglia model of context is based upon the idea of a multi-context (MC) system. These systems are based upon similiar intuitions to the McCarthy model of context, but rather than introducing a *ist* predicate MC systems have context as part of the meta-logic and introduce “bridge rules” to enable inter-context reasoning. This difference also means that contexts and bridge rules are not, by default, first-class objects in a MC system

A context consists of a logical language (e.g. first order logic), a set of axioms in this language and a set of inference rules (e.g. the rules of natural deduction). So in an MC system it is possible to use different logics in different contexts, whereas McCarthy’s model is based upon first-order logic.

The notation $c : \Phi$ is used to express that formula Φ is true in context c . Reasoning between contexts is done using bridge rules, rules whose premises and conclusion belong to different contexts. The general form of a bridge rule is:

$$\frac{c_1:\Phi_1, \dots, c_n:\Phi_n}{c_{n+1}:\Phi_{n+1}}$$

where c_1, \dots, c_n are contexts, Φ_1, \dots, Φ_n are propositions, $c_1 : \Phi_1, \dots, c_n : \Phi_n$ are the premises of the rule and $c_{n+1} : \Phi_{n+1}$ is the conclusion.

An example of a MC system is given in [49] where an agent, John, has a context representing the beliefs of another agent, Mary. John has the beliefs $B_m(P)$ and $B_m(P \rightarrow Q)$ where $B_m(P)$ denotes that Mary believes proposition P . John uses reflective bridge rules to assert P and $P \rightarrow Q$ in a different context, which represents a simulation of Mary’s reasoning, and performs local inference there to derive Q . This can then be reflected up to the context with the belief predicate to give us $B_m(Q)$.

Local Model semantics has been developed to provide a semantics for MC systems, a necessary development since MC systems have a distinct proof theory. The basic principles behind local-model semantics are: *locality*, reasoning only uses a part of the resources available; and *compatibility*, reasoning performed in different contexts should be compatible. A model for a MC system is a compatibility relation which describes the relationship between the models of the different contexts.

6.2.1.3 Alternatives

There have been several logics which attempt to incorporate some notion of context [1]. We shall look at Situation theory and Buvac and Mason's modal contextual logic as distinct approaches.

Situation theory [98, 35] arose from related intuitions to research on contextual logics within AI. The view was that traditional semantics for logic ignored the situated nature of belief and reasoning. In particular the partiality of representation: that we do not know the truth or falsity of all propositions. This led to the development of situation theory and situation semantics to handle this problem. The theory was based around the idea of *infons* which formalise what is true in a given situation and the idea of constraints holding between infons. The work on constraints was extended to an account of the logic of information flow in distributed systems [8], using local logics with channels between them to formalise the interactions between systems. This approach has been compared with the contextual logics above in [33] where local logics are identified with contexts and channels with bridge rules.

What is most interesting about the Channel Theory approach is the use of tokens of information in justifying the assertion of a proposition within a local logic. What we are interested in is the connection between the objects in different contexts and how we can justify identifying the signatures of different contexts. For instance in the 'Buchanan v. Babco' case the excise duty, as an object within the context of the real-world event, might be identified with the category defined in the text by the phrase "charge incurred in respect of carriage".

The difficulties with using the Channel Theory approach is the relatively complicated, and specific, meta-theory for the logic. Unlike AI approaches, where this kind of meta-theory would be handled by the use of meta-level reasoning within a conventional logic, Channel Theory introduces a complex meta-theory specifically for studying information flow. It is unclear how this work could be adapted to modelling ontology evolution in the Law, however the intuitions behind Channel Theory are useful.

Buvac and Mason were students of John McCarthy who formulated a contextual logic which treated the *ist* predicate as a modality and gave it a Kripke semantics [23]. This provides a more sound theoretical basis to the contextual logic, but removed some of the more flexible features (contexts as first-class objects). The use of modal-logics also went against McCarthy's advocacy of a contextual logic as an alternative to modal logics.

6.2.2 Contexts in Legal reasoning

If we want to use contextual logic to model the reasoning in a legal case, we must investigate what kinds of context exist in legal reasoning.

In the examples of contextual propositions we give we use the Trento notation $c : p$ to indicate that proposition p is true in context c . However we use the McCarthy approach to contexts, so that contexts can be treated as first-class objects within the logic. The notation $c : p$ should be considered shorthand for the assertion $ist(c, p)$.

We can think of the following kinds of context that might be used to model legal reasoning:

- Legal systems
- The current case
- The legal theories presented in a case
- The agents' belief states
- Precedent cases
- Sources of evidence
- Legislation
- The real-world event

A legal-system context would define the rules that govern the reasoning and procedures within a given legal system. These rules place constraints upon how lawyers can behave in a legal case, within that legal system. For example there might be the following rule in the English Legal System:

$$english_law : \forall X, Y. similar(X, Y) \wedge \exists V, V'. verdict(X, V) \wedge verdict(Y, V') \rightarrow .V = V'$$

which states that if two cases are similar then their verdicts should be the same. This rule (partially) expresses the principle of binding precedent in English law.

The current case is the legal case that is currently in disputed. The legal case involves a real-world event, that caused the legal case, and the legal system. What we want to represent is the discourse in the current case, the theories that the lawyers produce in the case.

The theories that the lawyers produce in a case are about what should be done in the current case, and why these actions should be taken. More specifically, for an advocate the theory states why the audience should behave in a certain way; for example the theory might state:

$$\text{discourse} : \text{theory} : ((\text{rwe} : \text{stole}(\text{person}, \text{money})). \rightarrow \text{.convict}(\text{person}))$$

This formula states that in the discourse context there is a theory, proposed by one of the lawyers, which states that if a particular person (denoted *person*) stole some money in the real-world event (denoted *rwe*) then they should be convicted. The theory makes a claim about what should happen in the current case.

Legal theories are essentially plans for the current case with justifications for the actions. A theory states what should be done in the current case and why it should be done. To ‘win’ a legal case is to have your theory accepted and thus change the behaviour of the judge(s).

A lawyer’s, or judge’s, belief state also constitutes a context. A belief context could be used to represent the potential biases that exist on a judge’s thinking. A lawyer could create a discourse context which uses those biases to create a desired belief about what the appropriate outcome to the case should be.

Precedent cases could be represented as contexts which can be used to justify a verdict in the current case. We could represent the discourse context in a precedent case, and use a bridge rule to relate it to the current case.

Legislation could also form a context. There are different possibilities for representing legislation as a context. We could represent the legislation directly as a set of propositions, which state the rules contained in the legislation. However, this representation would require us to formalise the legislation, and hence remove any ambiguities in the text. Alternatively, we could represent the legislation as assertions about the linguistic content of the text, and use bridge rules to represent the interpretation of the linguistic content.

An example of the former approach might be an Act with the following rule:

$$\text{act} : \text{stole}(X, Y). \rightarrow \text{.jail}(X, 12\text{months})$$

which states that if you steal something then you should go to jail for 12 months.

Lastly the real-world event is a context, a particular event that occurred which provoked the legal case. This should be fairly straight-forward to represent as a context, since it is just a description of the event. For example:

$$\text{rwe} : \text{stole}(\text{man}, \text{purse}) \wedge \text{own}(\text{purse}, \text{woman})$$

which states that in the real-world event a man stole a purse which was owned by a particular woman.

6.2.3 Limitations of existing contextual logics

There are some problems with the existing contextual logics. We want the following features from a contextual logic:

- Contexts as first class objects
- Ontology of contexts
- Reification of bridge rules

We would like contexts to be first-class objects and to be part of an ontology of contexts. This is useful for representing the different kinds of context and their relationship. For instance, legislation contexts should be part of a category of contexts representing legislation, with a default bridge rule for importing legal rules into the current case.

We would like to be able to reify bridge rules. Lawyers are free to argue about the rules that are invoked to justify a claim, and so any formalisation using contextual logics must allow the contextual bridge rules to be argued about.

These features are present in the McCarthy model of context, when combined with a meta-level first-order logic, but the meta-logical consequences of their inclusion are not known.

MC systems have a well-researched proof theory and semantics, but they do not include these features by default. In the MC systems approach contexts are part of the meta-logic and are not directly accessible without introducing a theory of contexts, as a theory within a particular context. It is unclear what the logical properties of such a theory would be, e.g. would there be paradoxes, such as Tarski's paradox for embedding a theory of truth in first-order logic (c.f. [139] [34])?

The main limitation of current contextual logics is the simplicity of their example applications. For instance, the most elaborate example of MC systems is the magic box example [47] which is fairly trivial and does not suggest any useful real-world application. The only exception to this is the CYC project, which has constructed a large-scale implementation of a contextual logic. We shall consider the CYC project further in section 6.3.2.

6.2.4 Argumentation and Contextual logics

With a contextual logic we can represent the contextual knowledge in a legal case, but we need to incorporate the reasoning into this representation. We have argued in the previous chapters that legal reasoning is fundamentally a process of argumentation, so we would like to incorporate this argumentation into the contextual logic representation.

There has been some research into combining contextual logics with default reasoning, a form of non-monotonic reasoning. Brewka et al introduce a form of Reiter's default logic for MC systems [21].

Since argumentation formalisms are independent of the underlying logic it should be straightforward to combine an argumentation formalism with a contextual logic to represent legal reasoning. However, the obvious way to combine the two, that is just use the contextual logic as the logic underlying the argumentation formalism, does not allow the argumentation itself to be easily reified. What we want is to be able to formalise arguments about the arguments in the current case, as these form an important part of legal discourse. We shall return to this problem in section 6.4.

One possible advantage of using argumentation as a reasoning framework for contextual logics is that it might get around the problems of possible inconsistency caused by the use of an *ist* predicate. Argumentation frameworks can provide different resolutions of Tarski's paradox, depending upon the semantics of argumentation used [103]. We have not investigated this property ourselves, but it seems like a reasonable approach.

6.3 Commonsense

We have noted above that the ontologies of meaning seem to depend upon underlying commonsense conceptualisations of meaning. In this section we shall consider the role of commonsense knowledge in legal reasoning, and survey some existing commonsense knowledge bases and their limitations.

Commonsense knowledge is the knowledge of basic facts about the world that we all share, facts such as "grass is green" or "people have two arms and two legs". However, these examples also demonstrate that commonsense knowledge is defeasible (grass can be painted blue, people can lose limbs).

In the following sections we will look at the role of commonsense in legal reason-

ing, some example commonsense knowledge bases and the limitations of commonsense knowledge.

6.3.1 The role of Commonsense in Legal Reasoning

There are extensive rules governing how legal reasoning should be conducted. However, legal reasoning relies heavily upon our innate reasoning abilities and our commonsense knowledge to reach conclusions. It isn't possible to define the meanings of all the terms that appear in legislation, and would create a bootstrapping problem if it were required. Instead, lawyers rely upon their commonsense knowledge to understand the meaning of the law.

There are different kinds of commonsense knowledge that are relevant to the legal domain. There is basic commonsense knowledge, of the kind McCarthy envisioned, that enables lawyers to understand legislation. In particular, the use of the word "reasonable" in legislation relies upon lawyers having the commonsense knowledge to determine what reasonable behaviour would be in that context. This knowledge has to cover knowing what events could occur in the domain and how available actions can affect events.

Aside from this commonsense knowledge, there is also common knowledge about the legal domain. This knowledge covers the legal system, the courts, legal cases, etc. This is knowledge that is specific to the law, but not to any particular piece of legislation. This kind of knowledge is considered to be of high value when conducting a legal case, leading to the advice to advocates of "know your court". A lawyer needs this knowledge to generate arguments that will be persuasive to the judges in the case.

Beyond knowledge of basic facts about the world there is also "knowledge" about human morality. Irregardless of whether people share a common sense of natural justice, most people are capable of assessing whether others would regard an action as good or bad. This common knowledge can be used as an implicit or explicit support for arguments. As an explicit support it might be stated as a support for an argument. An implicit support it might be used by a lawyer to generate an argument that exploits biases in judicial reasoning.

For example, in the *Popov v. Hayashi* case there was common knowledge that the assault on Popov was wrong, a sign of excessive greed. This sense of injustice was used to argue against awarding full possession to Hayashi, as it would be rewarding the violent behaviour of the crowd.

In the *Buchanan v. Babco* case, it might have been common knowledge that it was unfair for the distillers to be charged excise duty due to a mistake on the part of the carrier. This sense of injustice may have affected the reasoning, even though it was not part of any explicit justification for the decision.

In general, there are limitations to what can be stated in legislation and case reports. The main observation we make is that the law requires commonsense knowledge to function, and that this commonsense knowledge can be controversial. We will discuss the limitations of commonsense knowledge in section 6.3.3 and in the next chapter.

6.3.2 Commonsense Knowledge Bases

The CYC³ project is by far the most comprehensive attempt to create a knowledge base of commonsense knowledge. CYC began as an attempt by Doug Lenat, and others, to formalise all the knowledge in an encyclopedia⁴. They soon discovered that there was a great deal of knowledge that was being presumed by encyclopedia editors, and which was necessary to understand the articles. They then switched their focus to trying to formalise this commonsense knowledge.

Knowledge is entered into CYC by knowledge engineers, people who specialise in formalising knowledge into logical notation. The aim of the CYC project is to get a critical mass of commonsense knowledge so that CYC can automatically acquire commonsense knowledge from natural-language texts [71].

The CYC project has been ongoing for the past 20 years. There have been no major applications resulting from CYC, but there is currently some interest in applications within the intelligence community [118].

Research copies of the CYC knowledge base are available, and we have used one to experiment with CYC. Our experience with CYC has not been particularly promising. The documentation for the project is quite old, and not particularly clear.

As a test we looked up the concept of skateboard, to see whether we would be able to find useful commonsense knowledge to argue about whether a skateboard is a vehicle.

The skateboard collection⁵ is an instance of the type-of-vehicle collection⁶, which is a second-order collection. So it seems as though CYC would interpret “vehicle”

³www.cyc.com

⁴Hence the name, enCYClopedia

⁵In CYC there is a distinction between collections and relations, see [74] for more details, this distinction is not relevant to our discussion.

⁶see <http://sw.opencyc.org/concept/Mx4rvViyvZwpEbGdrcN5Y29ycA>

as including “skateboard”, presuming that it associated the word “vehicle” with the type-of-vehicle collection.

However, it isn’t clear why CYC has this definition of skateboard. There is no justification of this definition within CYC, and it isn’t clear how we could use it in a legal argument. CYC does not have any knowledge specific to the law. This means that whilst it could be useful to supply some basic commonsense knowledge, it would still be necessary to formalise all the common knowledge about the legal domain if we wanted to automatically generate arguments about the meaning of the law.

Another problem with using CYC in any application is that it isn’t clear what authority CYC definitions have. For instance, CYC has a concept of skateboard which appears to include it as a vehicle, so a system based upon CYC might simply propose that the rule “No vehicles in the park” applies in the case of skateboards. The problem with using this argument in a legal case is that the CYC definitions do not have legal authority. An application based upon CYC would need to find further evidence to support the definition of the term.

The knowledge in the CYC knowledge base is only useful to the extent that it is common and uncontroversial. However, in the hard legal cases we have been looking at the concepts that are evolving are controversial, and so the knowledge CYC has about them may not be of any use.

As we noted in section 2.3.4, there has been discussion with the AI and Law community about the importance of commonsense ontologies to legal reasoning. In particular, in [63] John Henderson and Trevor Bench-Capon discuss the creation of a general-purpose ontology for modelling legal cases. They develop a system which uses a general ontology of occupations to describe the duty of care someone owes within their occupation. This ontology is general purpose in that the conceptualisation of occupations is not specific to this system. In their paper they describe the difficulties of using WordNet for this purpose, as it is difficult to relate the natural language descriptions of WordNet concepts to a case. They also discuss the need for the ontology to contain concepts specific to the law.

Their proposal is interesting, but faces the difficult problems raised in this thesis. In general, the attempts at creating a commonsense ontology within AI and Law have been focused upon creating commonsense terminologies of limited scope, and don’t formalise any commonsense rules, which restricts their usefulness for argument generation.

6.3.3 Limitations of current Commonsense Knowledge-Bases

We noted in the previous chapter that some of the reasoning mechanisms appeared to be ‘black boxes’ in that they presented an opaque representation of an underlying phenomena that any system using them couldn’t introspect upon. For example, argument schemes are based upon an underlying commonsense model of the domain of argumentation. The commonsense model that lead to the argument scheme can’t be introspected upon by an automated reasoning system using the scheme. Similiarly Thagard’s model of coherence uses various coherence principles to create weighted links between propositions, but the resulting system can’t reflect upon the coherence principles used to set the weights.

We mentioned in section 2.2 the role of implicit knowledge in ontology evolution. However, we didn’t distinguish between implicit knowledge in the form of facts and implicit knowledge in the form of procedures⁷.

An implicit procedure is a mechanism that we have to perform some reasoning task, but that we can’t introspect upon. For example, most people can learn how to navigate their way around a city, but they can’t introspect upon the algorithms they are using to learn routes, or the algorithms they use to plan a route. These tasks are typically performed automatically and subconsciously.

We have argued above that modelling ontology evolution in legal cases requires ontologies about meaning, and that these ontologies about meaning are based upon our commonsense representations of meaning. Some aspects of these ontologies might be based upon implicit procedures.

The problem for modelling ontology evolution in legal reasoning is to understand what aspects of the implicit knowledge should be in the form of commonsense facts, such as “grass is green”, and what should be should be in the form of implicit procedures.

The CYC knowledge base contains a large body of formalised commonsense facts, but it doesn’t have any implicit procedures for reasoning. CYC can tell you that if you drop a ball then it will fall to the ground, but it might not be able to estimate how long that ball will take to fall as it doesn’t perform any physical simulations. We can, as an example, contrast this with the AI research on qualitative physics [42], which studies models for making these kind of estimates based upon features of the ball and the distance it is going to fall.

⁷c.f. the distinction in AI between declarative and procedural knowledge

Some aspects of human moral reasoning may well be based upon such implicit procedures. Some situations can seem wrong without us being able to describe exactly why we think that they are wrong.

For instance, Jonathon Haidt has constructed the following example scenario as a test of people's moral judgements [58]:

Julie and Mark are brother and sister. They are travelling together in France on summer vacation from college. One night they are staying alone in a cabin near the beach. They decide that it would be interesting and fun if they tried making love. At very least it would be a new experience for each of them. Julie was already taking birth control pills, but Mark uses a condom too, just to be safe. They both enjoy making love, but they decide not to do it again. They keep that night as a special secret, which makes them feel even closer to each other. What do you think about that, was it OK for them to make love?

People who hear the story have an immediate moral reaction that what Julie and Mark did is wrong. However, when asked about their reasons for their judgement they frequently give answers that are inconsistent with the story, such as complaining about the risk of children with birth defects, which is unlikely due to the use of birth control. Most people believe that the scenario is wrong, but struggle to come up with a justification for their belief.

More generally, Timothy Wilson [153] points out many circumstances in which we rationalise subconscious instincts. There is certainly some evidence that people form beliefs on the basis of intuitions that they can't rationally justify. This affects any attempt to construct a commonsense knowledge base, as we need to account for these common intuitions.

As we noted above, knowledge about the legal system, and even about the particular court that you are arguing within, is very important for generating legally valid arguments. This knowledge goes beyond having an upper ontology about the law, to knowing the dark art of advocacy. There are no knowledge bases that represent this knowledge.

It is important to note that the CYC project has been a massive investment of time and money. The project has produced some interesting work, but does not appear to have produced a general purpose commonsense-reasoning system⁸. This suggests that creating a commonsense-reasoning system is hard, and certainly well beyond the scope of this project.

⁸The CYC project has not created a system capable of understanding everyday stories in natural language, which might be regarded as the gold standard for a system purporting to have commonsense.

The problem is that it is difficult for us to determine the structure of our common-sense knowledge. A member of the Vienna Circle (see section 4.2.1), Otto Neurath, used an analogy with repairing a ship at sea to describe the problem. There is no place for the sailors to dock the ship to begin repairs, everything that is replaced is currently still being used to keep the ship afloat. Similarly, we rely upon our commonsense knowledge to help us determine what that commonsense knowledge is.

Even if we can create commonsense knowledge bases, we are still limited by the acceptability of this knowledge in legal arguments. A judge would currently have no reason to accept an argument which relied upon an axiom in CYC, unless they also regarded this axiom as self-evident.

So we are stuck on two points. Firstly, the content of a general commonsense knowledge base, which would be adequate for modelling the argumentation in any legal case, is very difficult to create. And secondly, even if we can create this content, it isn't necessarily an accurate reflection of what would be acceptable in a legal case.

The consequences for our work is that creating a general ontology of meaning, which we have argued is essential for generating arguments about meaning, is currently infeasible, and it is unknown how to construct one. We will look at this problem further in the next chapter.

6.4 Discourse Context

We can now present a model of argumentation about the meaning of the Law based upon the idea of creating a discourse context. The discourse context represents the social discourse that results from a legal case.

The purpose of the model is to consolidate our discussion of contextual logics, and demonstrate their utility in legal case modelling.

In this model there are propositions representing claims about the meaning of the law in the current case. Arguments are proposed by the lawyers to support or attack these propositions. The arguments create theories for the competing parties in a case. The theories refer to the context of the case, in order to justify their arguments about the meaning of the law.

The discourse context contains claims made within various theory contexts about the outcome to the current case. The discourse is about the current case, and so refers to the contexts of the real-world event, the relevant legislation and the other contexts relevant to the case.

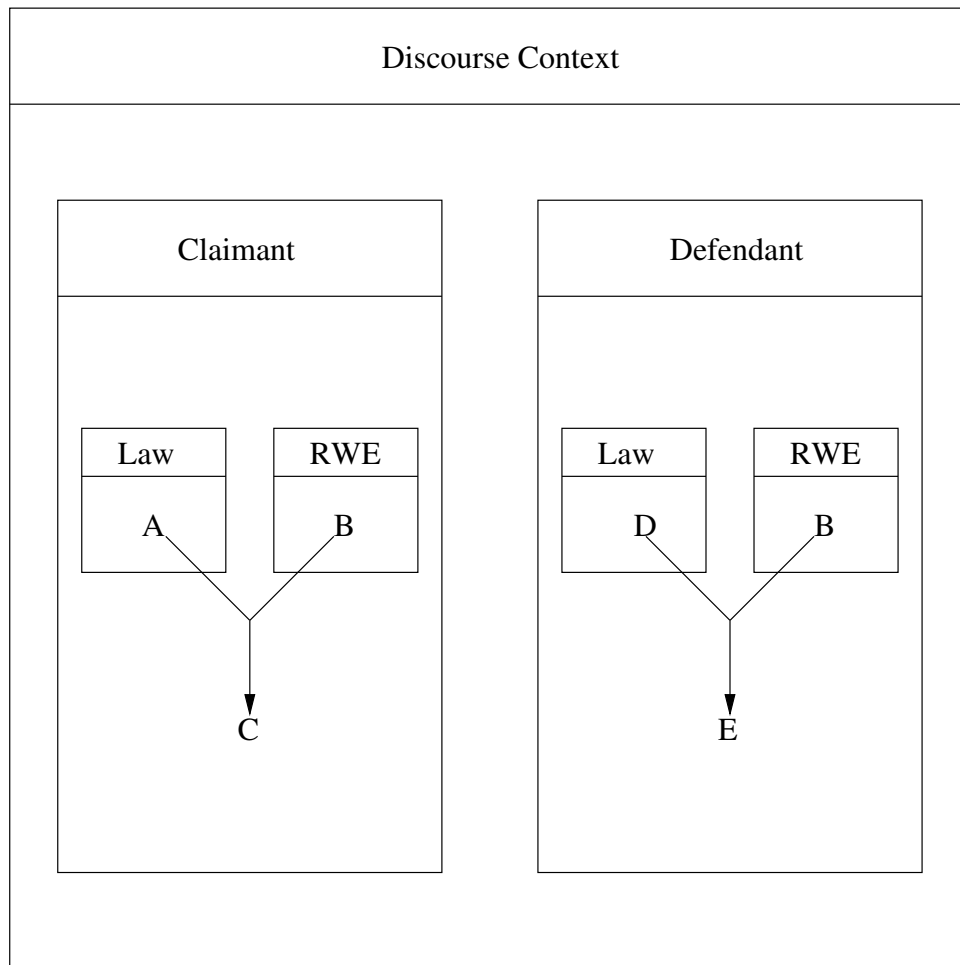


Figure 6.1: A diagram showing an abstract instance of our discourse context model for legal reasoning. The boxes represent contexts, so propositions within boxes should be interpreted as *ist* assertions. The arrows represent linking axioms between propositions in different contexts. In the abstract implementation, the Claimant makes a claim about what the law states and what happened in the Real-World Event and uses these claims to justify another claim within their theory. Similarly, the Defendant makes a claim about what the law states and uses this distinct claim to justify a claim within their theory. The distinct claims about the law and the real-world event represent distinct interpretations of the law and characterisations of the real-world event. These claims have not been justified in this abstract implementation, but there might be further contextual sources of evidence used to justify these claims. We have only shown a claimant and a defendant theory, but the model could also include judicial theories.

The main intuition behind our model is that lawyers can create theories regarding the current case as theories within the discourse context. They can use the representation of the discourse context to determine the consequences of a theory within the current case. For instance, the discourse context allows reasoning about the consequences of a claim about the meaning of the law, such as whether that claim would enable the law to apply, or not, to the real-world event.

Our model is designed to represent both of the kinds of uncertainty in the law that we discussed in section 3.1.1. Inconsistent law can be represented as multiple distinct sources of law, represented as contexts, which are referenced in a legal-case context as being relevant to the real-world event. The inconsistencies between the multiple sources of law create a problem that must be resolved. Incomplete law is represented by the existence of multiple conflicting mappings between a source of law and the real-world event.

We presume a general contextual linking axiom that connects any assertion made in a theory context with an assertion made in the discourse context, i.e. the theories are flattened the discourse context. So since the claims made by the different lawyers will conflict we have an inconsistent discourse context.

A lawyer could then apply an argumentation framework to the contextual logic theory representing the discourse context. A semantics for the argumentation framework would resolve the conflicts between the theories and justify an outcome for the case. However, we don't specify how this should be done within our model, since as we have noted in previous chapters there is no consensus about the correct way to resolve these problems. We discuss this problem further in section 7.2.

6.4.1 An example implementation of this model

Consider again the Buchanan v. Babco case 3.2.2. How could we implement a model of this case as a discourse context? Here are two example theories, which might be produced by the plaintiff⁹ and defendant in the case. We formalise a few of the arguments from section 3.2.2.1 into the discourse context model we described above.

We use the following contextual signature to formalise the theories:

Contexts:

- *english* denotes the English legal system.

⁹The discourse context model uses "Claimant" since this is the new term for plaintiff within English law.

- *french* denotes the French legal system.
- *bvs* denotes the case, in the Dutch courts, of *British-American Tobacco Co. (Nederland) B.V. v. van Swieten B.V.*
- *act* denotes the Act of Parliament enacting the Convention.
- *convention* denotes the French original of the Convention.
- *rwe* denotes the real-world event that led to the case.

Predicates:

- *openTexture(X)* means the phrase *X* is open textured.
- *treaty(X, Y, C)* means that the phrase *X* has a corresponding phrasing as *Y* in an authoritative version of the international treaty *C* that the law is an enactment of.
- *includes(X, Y)* means that noun phrase *X* includes the concept *Y*.
- *purposeful(C, X)* means that the interpretation of legislation context *C* as including the proposition *X* is a purposeful interpretation of the legislation.
- *expert(P)* means that an expert has claimed proposition *P*.

Functions:

- “*othercharges*” denotes the phrase “other charges incurred in respect of carriage” in the Act.
- *excise* denotes the concept of excise duty.
- “*encourus*” denotes the French phrasing of “other charges ...” in the Convention.

Note that we don’t give formalisations of these concepts, since, as we noted above, the commonsense definitions are too difficult to formalise currently. This is in part why we don’t specify reasoning mechanisms for our model, since any automated reasoning system would require more extensive formalisations of the signature terms that we use in order to generate interesting arguments.

The contexts represent the social, interpretative and situational contexts in the case. The *english* and *french* contexts represent legal systems which are social contexts representing a conceptualisation of the social institution of law. The *act* and *convention* contexts are interpretative contexts representing the interpretation of legislation. And

the *bvs* and *rwe* contexts represent situations which have occurred and whose details are relevant to the case.

There are intuitive relationships between these contexts. The legal system contexts contain the legislation contexts, for instance the *act* legislation context is within the *english* legal system context. We have not attempted to formalise these relationships due the difficulties of representing them; in particular, if we wanted to represent these relationships we would need an ontology of contexts as discussed in section 6.2.3. However, these relationships would be used by the modeller to create formal arguments within this representation.

We don't present full axiomatisations of the arguments we cover. Our aim is simply to discuss the main axioms, so that the difficulties involved in formalising the arguments are clear.

6.4.2 Plaintiff

The plaintiff's theory includes some of the arguments that were proposed by Buchanan to justify their interpretation of the law. We use the same notation as in section 6.2.2.

We first try to formalise the argument that it is acceptable to consult the French version of the convention. We use the following axioms, all of which are asserted within the Plaintiff-theory context:

1. $act : openTexture("othercharges")$
2. $convention : includes("encourus", excise)$
3. $treaty("othercharges", "encourus", convention)$
4. $act : openTextured(X) \wedge treaty(X, Y, C) . \rightarrow .$
 $(C : includes(Y, A) \rightarrow act : includes(X, A))$

These axioms could be used to justify the following claim:

$$act : include("othercharges", excise)$$

Note that we have omitted some of the necessary axioms to formally derive this.

The axioms state the following:

1. That the phrase "other charges ..." is open-textured.

2. That in English law if an Act is open-textured and is part of an international treaty with an authoritative version then if the interpretation of that phrase in the treaty includes a concept A then the Act of Parliament should also include the concept. This axiom is designed to express the idea of uniform interpretation between local enactments of a law and an international treaty.
3. That the phrasing of the law in the Convention includes excise duty.

We next try to formalise the argument that we can use the French methods of interpretation when consulting the English law.

1. $french : purposeful(convention, includes("encourus", excise))$
2. $french : purposeful(C, P) \rightarrow C : P$
3. $french : purposeful(convention, includes(X, A)) \wedge openTextured(X) \rightarrow includes(X, A)$

These axioms could be used to justify the following claim:

$$act : include("othercharges", excise)$$

The axioms state the following:

1. That there is a purposeful interpretation of the convention in the French legal system which would include excise duty within the French phrasing of "other charges ...".
2. That if there is a purposeful interpretation of a legislation context in French law, that interpretation is valid.
3. We should use this purposeful interpretation within the English legal system if there is open texture in the English phrasing.

6.4.3 Defendant

For the defendant, Babco, we first try to formalise the argument that the Dutch precedent could be used to justify the exclusion of excise duty from the phrase "other charges ...".

1. $openTextured(X) \wedge treaty(X, Y, C) \wedge ForeignPrecedent : includes(Y, A). \rightarrow . includes(X, A)$
2. $bvs : \neg includes("encourus", excise)$

These axioms could be used to justify the following claim:

$$act : \neg includes("othercharges", excise)$$

The axioms state the following:

1. That if a phrase is open textured, is part of an international treaty and there is a foreign precedent in which the phrase was interpreted as including a concept then we should interpret the phrase as including this concept. *ForeignPrecedent* is a contextual variable that could be replaced by any foreign precedent case which discusses the interpretation of the phrase.
2. The Dutch case did not include excise duty in "encourus ..."

We next try to formalise the argument that there was no expert evidence given to the meaning of the French law, and so we can't conclude what it means.

1. $\neg expert(french : convention : includes("encourus", excise))$
2. $(french : convention : P) \rightarrow expert(french : convention : P)$

These axioms could be used to attack the following claim, since no expert evidence for the interpretation of the French Convention exists.

$$act : \neg includes("othercharges", excise)$$

The axioms state the following:

1. There is no expert evidence that the meaning of the French Convention would include excise duty within the French phrasing of "other charges ...".
2. If a claim is made about the interpretation of the convention in French Law then there must also be an expert who supports this claim.

6.4.4 Limitations of this model

Our model can be used as a framework to represent the discourse regarding the meaning of the law in a legal case. However, there are some features of legal reasoning that are not reflected in the model, at least as we have currently presented it. In particular, we can see the following issues arising:

- Choice of contexts
- Perspective on the case
- Lack of commonsense knowledge
- No semantics or reasoning mechanisms

The contexts might be viewed as a bit ad-hoc, simply created for modelling purposes and bearing no relation to the reasoning that actually occurred. In some respects this is a general presumption behind contextual logics: that people simply do chunk the world up into different contexts and limited their reasoning to local ones. What is more of a concern is whether the contexts in our particular contextual model are an accurate reflection of the distinctions made in the case.

We have presented the discourse context as a representation of the case, but whose perspective on the case are we modelling? We can maintain a certain neutrality by not specifying the reasoning mechanisms used in the case, so we don't prescribe how the arguments should be resolved, but it isn't clear what perspective we are representing.

We argue that we are representing an agent's perspective on the discourse context, the agent who is doing the modelling of the case. If the agent were a lawyer tasked with representing one of the sides in the legal case, the discourse context would be their representation of one possible discourse based upon that lawyer's presumptions about the reasoning in the case.

We have not included judicial theories, however this is a straightforward addition to the model. A modeller would just need to include a context representing the judge's theory in the case.

A more interesting omission from our model is the lack of belief states. It would be interesting if we included contexts representing the belief states of the lawyers and judges in the case, and could use them to resolve the reasoning in the case. This would provide a mechanism to represent the biases that can affect the reasoning in the case.

We have not investigated this possibility, since the belief states are not part of the explicit discourse in the case, but it would be an interesting addition to the model.

As we have noted above, commonsense knowledge is a problem for constructing models of legal reasoning. We have avoided the problem in this model, and left it to the modeller to create commonsense warrants for their arguments. This is a major problem for any attempt to automate the reasoning, since we require a person with commonsense to input the information we would need to generate arguments.

Another problem for automation is the lack of reasoning mechanisms. The idea of the discourse context is that it will inevitably form an inconsistent theory, since the claimant and defendant theories will conflict and we assume contextual rules which allow the inference of any proposition in one of the theory contexts into the discourse context. Our intention is that the reasoning is done by mechanisms outside the contextual logic. The logic just provides a mechanism to check whether a theory is consistent.

For instance, the conflicts in the discourse context could be resolved using either the argumentation frameworks or Thagard's coherence model from section 5.7. However, for Thagard's coherence model it would be the task of the modeller to determine the weights and the links between propositions.

So, it seems that our model has several of the limitations that affect the techniques we surveyed in the previous chapter. The novel feature of our model is that it can represent the contexts in the reasoning of the case, and so can represent some aspects of legal interpretation. In particular, our model can represent some aspects of the different legal systems and the relationships between them; and so is well suited to representing the interpretation of laws which have been enacted in multiple legal systems, such as in *Buchanan v. Babco*.

6.5 Summary

In this chapter we have extended the core claim of this thesis: *ontology evolution in the law occurs as a result of the arguments about meaning proposed in a legal case*. However, we have also discussed the limitations of this representation of the problem. Ontology Evolution is now more understandable, since it is simply another domain of argumentation in a legal case, but also more complex, since the ontologies these arguments depend upon are so heavily dependent upon commonsense knowledge.

We have looked at the usefulness of existing A.I. technologies, in contextual logics and commonsense reasoning, to modelling the content of these arguments about

meaning. In particular, we have proposed a model of legal reasoning based upon a *discourse context*. However, we have also noted the limitations of this model: primarily that the contexts and the commonsense knowledge required to understand a legal case are hard to make explicit. We have also not specified how the discourse context should be created for a legal case. In the next chapter we will look at whether these limitations could be overcome to construct computational models of the ontology evolution in a legal case.

Chapter 7

Requirements on a Theory of Ontology Evolution in Law

As a reminder we proposed the following theses in the introduction:

1. Ontology evolution in legal cases occurs as a side effect of arguments about the *meaning* of legal rules and principles.
2. There is a distinction between prescriptive and descriptive theories of ontology evolution in legal reasoning.
3. A complete descriptive theory of ontology evolution in legal reasoning is an AI-complete problem.
4. A combination of the features of the different theories could provide the basis for a useful system for aiding the construction of arguments about the meaning of legal ontologies. However, any system would require significant human interaction to produce useful output.

In this chapter we summarise our arguments for these theses and discuss their consequences for developing applications and future research.

7.1 Ontology Evolution in Legal Reasoning

Our first claim, the main claim of this thesis, is about what ontology evolution in legal reasoning actually is. At the beginning of this thesis we gave an intuitive idea of the process of ontology evolution and how it occurs in the law. This intuitive idea was

developed into a conceptual model of how ontology evolution occurs in legal cases and a discussion of how we can construct formal models of this process.

The next few sections describe the main components of our model and summarise the arguments we have presented to justify our model.

7.1.1 Arguments about Meaning

We have claimed that ontology evolution in legal cases occurs as a side effect of the arguments about the meaning of a law. But what exactly does this claim mean?

We have argued that the reasoning in a legal case is driven by argumentation. Lawyers will present arguments to support their claims about the case. These claims are created to support the outcomes they desire for the case.

Lawyers can make claims about how the ontology of the law relates to the commonsense ontology that describes the real-world event in the case. These claims are about the meaning of the legal rules in the relevant legislation. The claims provide an interpretation of the law with respect to the real-world event.

Claims about meaning require meta-ontologies to formalise their content. For instance, lawyers will argue about whether a term such as “vehicle” denotes a concept which includes skateboards. Terms such as “includes” form a meta-ontology, an ontology used to describe the meaning of another ontology.

There are meta-ontologies used in everyday dialogue and there are meta-ontologies which are specific to the legal domain, such as interpretation sections in an Act of Parliament. Lawyers may also use meta-ontologies related to literary interpretation, for instance they might talk about the “scheme” behind a statute¹.

The claims are specific to the context of the current case. The claims do not comprehensively define the meaning of the term, but only clarify it enough to resolve the uncertainty in the current case.

The claims are justified by using evidence in the context of the case. The context of the case is used to generate the arguments about the meaning of the law in the current case. For instance, a lawyer might argue that the intentions of the Council were to ensure the safety of people using the park, and that this safety is not affected by skateboarders. The argument uses a claim about the surrounding context to the case, namely the intentions of the Council which created the park bylaw, to justify a claim about the meaning of the law.

¹c.f. *Buchanan v. Babco*

If an argument about the meaning of the law is accepted by those who have authority over the law, the meaning of the law can change. In common law jurisdictions this change is made through the principle of binding precedent, future cases are bound by the decision. In other jurisdictions the change is more subtle, the analysis in the case can influence the analysis in future cases, even if it is not binding.

Our claim is that this process is how ontology evolution occurs in legal cases. Our claim is based upon the definitions we have given of ontology evolution and legal reasoning. In the beginning we presented an intuitive account of the process. We then analysed what ontology, ontology evolution and legal reasoning were, and based upon this analysis we presented an account of how legal ontologies could evolve in legal cases. However, this claim just describes a mechanism for ontology evolution. It does not define how ontology evolution will actually occur in specific cases. In section 7.2 we look at the kinds of theories that this claim allows.

7.1.2 Contextual knowledge

Our conceptual model of arguments about meaning depends upon the idea of context. We have argued that claims about meaning are specific to a particular case, and that the claims are justified using contextual sources of evidence.

What makes the sources of evidence contextual? Our argument here is based upon the motivating intuitions for contextual logics, in particular the idea of knowledge representation being contextual. The intuition is that we organise our background knowledge into sets of propositions specific to a context. We then use inter-context rules to infer information between contexts.

There is some psychological evidence to support a contextual model of representation. Barsalou has demonstrated that categorisations are situated, in that people can change their categorisation of an object depending upon the situation they are in [7]. This gives some justification to the use of a contextual logic to represent the arguments in a legal case, where we want to allow a concept to have different meanings depending upon the context.

Another advantage of a contextual logic for our model is that we can represent different interpretations of the law as being different theories within the same discourse context. The model also extends to modelling different conceptualisations of the discourse context, and hence different strategies to affect the judges' conceptualisation of the discourse context.

7.1.3 Discourse Context

Our model of legal reasoning in hard legal cases is based upon the idea of a discourse context. The idea of legal reasoning as forming a discourse in which lawyers present theories is not new. Our addition to the idea is to demonstrate how such a model can be formalised using contextual logics.

We presented our model in the section 6.4. We use a simple contextual logic representation with an *ist* predicate and rules for inferring between contexts.

The basis of our model is a discourse context representing the arguments in the current case. These arguments make reference to relevant background contexts in order to justify a claim about how the meaning of the law, in a context representing the law, relates to the commonsense description of what occurred in the real-world situation.

There are, of course, limitations to this model. The contexts we have used are debatable, for instance separating theories into only claimant and defendant ignores the judicial theories in a legal case. We have also not investigated the role of reasoning about different representations of the discourse context. For instance, a lawyer might represent the beliefs a judge has about the discourse context and plan their arguments using this representation.

Another feature of legal reasoning we have not fully investigated is the role of reflection. The theories within the discourse context might have to refer to the discourse context itself, and make claims about what should be true in the discourse context. For instance, a lawyer might want to claim that a precedent should not apply to the current case due to differences between the precedent case and the current case. This claim could be formalised by asserting a proposition within a theory that states that the discourse context (for the current case) is distinct from the discourse context in the previous case due to the differences between the real-world events in the cases. This would be a claim within the discourse context, which makes a claim about what is true in the discourse context. This raises various difficulties for formalising the contextual logic, as there is a risk of paradox. We have not investigated these issues, but they are worth considering in future research.

We argue that we can make additions to our model to incorporate these features, as the contextual logic framework is quite flexible, but we have not sufficiently investigated these claims, instead we have left them for future investigation.

Our model does not specify how the discourse context should be created. Competing claims are made, but our model does not describe which claims should be accepted.

In the next section we will argue that there are two different possible specifications for modelling legal reasoning using our model.

7.2 Types of theories

We have seen that theories of reasoning and creating meaning can either be based upon modelling how people perform these tasks, or constructing a normative model of how the task should be performed, see chapters 4 and 5. These theories are termed “descriptive” and “prescriptive” theories respectively.

Note that, this distinction is far from novel. Early logicians recognised that their formal logics were not accurate models of human reasoning, but rather models of how reasoning should be done. Both Paul Thagard in [136] and John Pollock in [99] give discussions of the distinction and its relevance to creating computational models of intelligent agency. We are interested in how the distinction specifically affects our attempts to create computational models of ontology evolution in legal reasoning.

The model we presented in the preceding section gives a framework for understanding ontology evolution in legal cases, but does not specify how the process of ontology evolution should be done. We have not specified what semantics should be used to accept or reject arguments about meaning in the discourse context. Since ontology evolution depends upon reasoning and meaning we argue that it too has prescriptive and descriptive theories. We describe these different kinds of theory in the following sections.

7.2.1 Prescriptive theories

A prescriptive theory of ontology evolution is one which presents a model of how the meaning of the law *should* evolve in cases where the meaning is uncertain. This theory can ignore the process of ontology evolution as it is currently carried out by lawyers in legal cases.

Most of the techniques we surveyed in chapter 5 are prescriptive theories of reasoning. For instance, the models of argumentation are not designed to perfectly mirror real-world arguments, but are instead intended to describe how argumentation should be performed. Exceptions to this are Walton’s research on argumentation schemes, which capture existing patterns of argumentation, and Thagard’s coherent reasoning, which attempts to model how people assess theories.

A prescriptive theory of ontology evolution in legal reasoning would have to prescribe both the acceptable meta-ontologies and the semantics of the argumentation about meaning.

For instance, a prescriptive theory might specify a description logic formalism for representing legal and commonsense concepts. This formalism would then have a meta-ontology based upon a formalisation of the description logic's meta-logic. The theory might also specify a semantics for resolving conflicting contextual arguments, for example based upon weighting the for and against arguments and selecting the side with the highest collective weight.

Such a theory would determine what arguments are acceptable or not in a given case, but it might not select the arguments that most people would agree with.

Although our model does not specify the semantics of the argumentation, it does specify a contextual framework for ontology evolution in legal reasoning. This framework makes some arguments easy to express, but might make some arguments awkward. Since the framework is an attempt to formalise the natural language argumentation that already occurs, it is in part prescriptive as it places constraints upon how the argumentation should occur.

7.2.2 Applicability of a Prescriptive theory

It is hard to create a prescriptive theory as one must have a theory and a philosophical foundation to justify that theory.

In the case of first-order logic the justification was an intuitive idea that this formalism could represent mathematical arguments correctly. This intuition was later formalised by Tarski using the idea of a model-theoretic semantics.

In the case of a theory of ontology evolution one would need a philosophical foundations for the ontology and for the epistemology, to represent the concepts and determine what one can know about their relationships. We have seen that foundations for ontology and epistemology are not universally agreed upon.

As has been noted in AI and Law research, any theory regarding legal reasoning must take account of legal theory. There is no universally agreed theory of legal reasoning. There are several distinct theories about what the law is, and so we can not base an uncontroversial theory of ontology evolution upon them.

A prescriptive theory could be enforced as a requirement for participating in a legal system. For instance, if legal arguments were required to be represented in our

framework, and there was a semantics for resolving conflicting arguments, then our framework would be a prescriptive theory for ontology evolution in legal cases within that legal system.

There is an analogy here with standardisation efforts for data interchange, such as XML². The standard becomes a prescriptive format into which data must be converted for it to be accepted by others.

However, the applicability of such a theory depends upon it being enforced by those who have authority, such as Parliament for the UK legal system. Without this authority a prescriptive theory might not be much use, as it wouldn't predict how the actual judges in the case would decide the arguments and hence wouldn't be useful to a lawyer preparing arguments for a case.

A framework, such as the one we have presented, may be useful as an aid to a lawyer to help them organise their arguments. An example of this is the use of Wigmore argument schemes to help construct legal arguments as shown in [143].

7.2.3 Descriptive

A descriptive theory would describe how people actually argue for changes to legal ontologies, and would make predictions about what changes to the law would occur to resolve the uncertainty in the law with respect to a real-world event.

In order to construct such a theory we would need to know what claims about the meaning of the law, and their supporting arguments, would be accepted by an adjudicator. We focus upon adjudicators as they have authority over what changes to the law will be accepted, and hence affect the outcome of the case.

However, there is an immediate problem with this approach: people differ in how they want to resolve the uncertainty in the law. Some judges will accept arguments that others reject.

So a descriptive theory would need to accommodate the different possible responses that an adjudicator might have to an argument. Some will reject it, and others accept it. A descriptive theory might have to explain why a particular judge would accept or reject an argument, what biases or values are invoked by the argument.

These explanations go beyond the problem that we have discussed so far in this thesis, since we have to account not only for the arguments about meaning in a legal case but also the affect of those arguments upon the mental state of the judges in a

²<http://www.w3.org/TR/REC-xml/>

case. This task raises various AI-complete problems that we shall consider in the next section.

7.2.4 Problems for a Descriptive theory

The third claim of this report is that a descriptive theory is an AI-complete problem. A problem is AI-complete if it requires human-level AI in order to provide a solution. We argue that a descriptive theory is an AI-complete problem as it requires a theory of the meaning of natural language, i.e. that we would have to solve the problem of natural language understanding.

Why would a theory of ontology evolution require solving natural language understanding? As we have noted, the Law uses natural language to communicate the laws that should hold in a legal case. Interpreting those laws is one of the main problems faced by lawyers in applying the law. In particular, the uncertainty in the law arises due to the open-texture of the natural-language legal-rules.

7.2.4.1 Meaning in Language

Arguing for the changes to make to a law requires understanding the meaning of the existing natural-language legal rules, and how they could be interpreted by other people. Applying the law involves understanding the natural-language texts used to describe the real-world event and connecting the facts in these texts with legal rules.

An automated system would have to read and understand legal texts and texts describing real-world situations, and it would have to understand how a person might read and apply those texts. This is far beyond the current state of the art.

7.2.4.2 Human Context

Understanding how a person might read and apply a legal rule to a real-world event requires being able to model their mental state and reasoning processes. This requires some very hard problems, which are far beyond the current state of the art. In particular, understanding human moral intuitions and how people assess similarity are hard problems which would have to be solved to understand how another person might approach a problem of legal interpretation, and how an automated system might be able to craft arguments to affect their interpretation.

People seem to share some values, but there are still significant differences between groups. For instance, Jonathon Haig argues that there are differences between conser-

vatives and liberals in what behaviours they consider to be morally acceptable [59]. These differences make it harder to create a descriptive theory, since an automated system would have to account for the values of the individual judges when creating arguments, rather than being able to construct a single account of human morality which could be used in all cases.

There is much research discussing the importance of analogy to understanding human reasoning. The role of analogy is particularly important in common law legal systems, where assessing similarities between cases is the primary form of reasoning. The difficulty in modelling this reasoning is that we don't know how people assess similarity, since there is still much psychological research to be done in this area. This prevents us from constructing a predictive system with the current state of the art, as we don't know how to form convincing arguments about similarities between cases.

There is also the institutional knowledge that is common among the participants in a legal system. An automated system, which had to construct arguments that would be accepted within a legal system, would have to understand that legal system.

7.2.4.3 Reflection

A problem with making claims about how people ascribe meaning to language is to account for our ability to reflect upon how we are interpreting an utterance and change our interpretation.

Reflection raises the problem of epistemological foundations. We have noted the foundationalist and coherentist approach to this problem, but we have not tried to incorporate either into our model. The problem is not so much what foundations should hold over our reasoning, but what foundations do hold in most people's reasoning.

Ultimately this is a hard problem that is well beyond the scope of this thesis. We briefly discuss possible approaches to this problem, within the scope of ontology evolution in legal reasoning, in the further work section.

7.3 Developing Applications

Our final claim was that we could use a combination of theories to develop applications to aid lawyers in constructing arguments about the meaning of open-textured legal rules.

7.3.1 Kinds of application

What kinds of task would a theory of ontology evolution be useful for? We can think of the following examples:

- Helping lawyers generate arguments about meaning
- Checking argument consistency
- Finding counter-arguments
- Suggesting interpretations of the law

Lawyers need to generate arguments to support their desired outcomes in a case. A system which could help the lawyer construct these arguments might be useful. Such a system could offer varying degrees of assistance, from providing a lightweight framework for constructing arguments to full automation.

Once arguments have been created it is useful to check them for consistency. This is the kind of task that could be relatively easily automated if the arguments have been formalised, although finding inconsistencies might be difficult on large theories.

A related task to consistency checking is finding counter-arguments. For instance, a counter-example found when checking consistency could be used to create a counter-argument.

The hardest task to automate is finding interpretations of the law. Automating this task requires formalised ontologies about meaning and abductive reasoning mechanisms to find interpretations of the law that would achieve a goal within the case.

For any of these tasks there are some common implementation problems, which we shall discuss next.

7.3.2 Developing ontologies

We have noted that modelling argumentation requires ontologies to describe the domain. Any application would require ontologies about the legal case and the relevant laws.

There is an active research field in AI and Law which is creating legal ontologies. However to be useful a system would have to have commonsense ontologies too, so that it could formalise the relationship between the real-world event and the law. These commonsense ontologies are much harder to form. There has been some discussion of

commonsense ontologies within AI and Law, but there are no significant attempts to solve the problem.

A practical approach would be to have the lawyer input the commonsense knowledge and the system provide a partial formalisation of the legal knowledge. The lawyer could enter their arguments in a subset of English that could be ambiguously parsed. However, it might be difficult to extract the underlying commonsense knowledge that the lawyer uses in their arguments, so the resulting formalisations might be incomplete.

There would be interface issues with this solution, and a learning curve for new users to create unambiguous descriptions of their arguments. However, the alternative is to try to work with pure natural language, which is currently unfeasible.

7.3.3 Developing reasoning mechanisms

The complements to any ontology are reasoning mechanisms to make inferences from it. We have surveyed many different methods of reasoning in chapter 5.

The main problems are incorporating reflection in the argumentation and the reasoning about coherence of a theory. We can create argumentation frameworks, and methods for theory evaluation, but we also need to provide a way for automated systems to reflect upon and argue about the arguments and methods of theory evaluation. These too are part of the domain of legal discourse.

This presents a theoretical challenge, of formalising methods of reflection upon argumentation and theory evaluation, and a technical challenge of creating systems which can perform this task.

7.3.4 Usefulness of Applications

It is unknown how useful any applications would be, since we can't fully automate any aspect of the reasoning within the domain. The most useful applications appear to be quite lightweight; they aren't based upon deep philosophical theories about meaning.

An analogous domain is that of interactive mathematical theorem provers. These systems can help mathematicians to develop mathematical proofs, by machine checking the mathematical arguments within a formal logic. However, uptake of these systems by mathematicians is quite low, very few use such a system in their day-to-day work.

The problem is that the systems don't, or aren't perceived to, interface well with the mathematicians' workflow, and it is difficult to learn how to use the system effectively.

The arguments formalised within the system don't resemble the more abstract proofs which mathematicians discuss. So the only benefit of the system is to provide that extra level of certainty that a proof is actually valid, and the effort required for this task is seemingly not felt to be worthwhile.

This may be a sociological problem amongst mathematicians, perhaps a new generation who are familiar with the technology might make greater use of such tools. The quality of the tools is also continually improving, making them more useful. However it does give us some idea of how useful any system would be to lawyers.

It is unlikely that any system we can develop in the near future would be perceived as being useful by lawyers. The system is trying to help the lawyer with a qualitative task, but is imposing some constraints upon how that task should be performed. These constraints may hinder the lawyers when they perform the task (c.f. [116]).

We can contrast this kind of qualitative application with a quantitative one: the problem of information retrieval. Lawyers have a vast quantity of documents to search through to find the relevant material for a case. This material is too great to read and fully understand. There are useful systems which can search through the documents looking for key words or features within the document. The user just has to create appropriate queries to search through the documents.

The lawyer doesn't have the time to search through all the documents, so this kind of system is very useful. Creating arguments is something that the lawyer must do, and requires an intricate understanding of the issues involved in the case. The advantage of having software support for this task is not so clear, as the software doesn't automate the task and can only provide limited support, since it doesn't fully understand the meaning of the arguments.

The main incentive for lawyers to use such a system would be if it were mandatory, that is if lawyers had to provide their arguments in a logically formalised format. However it isn't clear why such a system would be created unless it was demanded by Parliament, and it isn't clear who would lobby Parliament for such a system.

So for the immediate future systems based upon theories of ontology evolution do not appear particularly useful for creating arguments in legal cases. There might however be an application in education. There are examples of case-based reasoning systems being used to teach lawyers to identify and apply precedents [2]. Similarly, it might be possible to develop systems to teach lawyers how to form arguments about meaning.

However, an argumentation aid would likely not be specific to forming arguments

about meaning. Any system which taught students how to form coherent arguments would likely extend to teaching students how to form arguments about meaning if it had the necessary background knowledge.

So our conclusions regarding the usefulness of applications are somewhat mixed. In the short term it seems unlikely that we can develop useful aids for lawyers, but these may be useful in the long term. We can develop education tools to help students learn how to form arguments about meaning, but these tools would likely be general argumentation tools, rather than tools specific to ontology evolution.

7.4 Further Work

We have presented various claims about ontology evolution in law, and a conceptual model of the process of ontology evolution in legal cases. In this section we discuss the possible extensions to our research.

The main areas in which further work seems productive are:

- Ontologies about meaning
- Domain-specific case-studies of ontology evolution
- Contextual logics
- Interactive tools for constructing arguments about meaning
- Acquiring commonsense knowledge

7.4.1 Ontologies about Meaning

We have argued that ontologies about meaning are important to formalising the process of ontology evolution in legal cases but we have not conducted a systematic study of the content of these ontologies.

It would be useful to survey a wide variety of legal cases and note what meta-ontological terms occur. You could then try to identify the axioms that are commonly accepted for these terms and use them to create an ontology about meaning.

You could also look for examples of ontologies about meaning outside the legal domain. For example, in literary interpretation there are extensive vocabularies used in an interpretation of a text. The law uses many similar terms, so this might lead to further insight into these meta-ontologies.

Psychological studies could be conducted to find out what axioms about meaning are generally accepted by people, and what differences exist.

7.4.2 Domain-specific examples

More extensive examples using our model of creating the discourse context could be created. These could form in-depth case studies of the ontology evolution in those legal cases.

7.4.3 Contextual logics

We have noted that there are limitations to the current contextual logics. In particular the foundations of contextual logics are uncertain and it isn't clear how to combine them with argumentation formalisms.

7.4.4 Interactive tools

We have argued that useful applications could be constructed, but we have not built any prototypes ourselves. It would be interesting to try to create an interactive tool using the discourse-context model we have presented in this thesis.

7.4.5 Commonsense knowledge

We have noted that commonsense knowledge is crucial to both forming and understanding arguments about meaning in hard legal cases. The argumentation in these cases relies upon the common ground of commonsense knowledge that people have. So developing software for automating tasks regarding ontology evolution requires an extensive commonsense knowledge base.

There are two problems here: one is developing commonsense knowledge bases, including creating expert knowledge that is specific to a legal system, and the other is creating simulations of how people approach basic reasoning tasks. The former problem has been attempted by systems such as CYC and the latter by research into qualitative reasoning. It would be interesting to see if this work could be included into a model of legal reasoning, such as the one we have presented in this thesis.

7.5 Conclusion

We have presented and defended the claims we introduced at the beginning of the thesis. Our main claim is that ontology evolution in legal cases occurs as a side effect of the arguments about the meaning of the law in the case.

We have discussed the kinds of theories of ontology evolution that could be created and their respective consequences. The consequences include the possibility of developing systems that help lawyers construct arguments about meaning. However, these systems would require significant human interaction and further research is required to understand how useful these systems would be and how they could be constructed.

Bibliography

- [1] V. Akman and M. Surav. Steps Toward Formalizing Context. *AI Magazine*, 17(3):55–72, 1996.
- [2] V. Aleven. Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment. *Artificial Intelligence*, 150(1-2):183–237, 2003.
- [3] K.D. Ashley. Reasoning with Cases and Hypotheticals in HYPO. *International Journal of Man-Machine Studies*, 34(6):753–796, 1991.
- [4] Kevin D. Ashley and Vincent Aleven. Reasoning symbolically about partially matched cases. In *IJCAI'97: Proceedings of the 15th international joint conference on Artificial intelligence*, pages 335–341, San Francisco, CA, USA, 1997. Morgan Kaufmann Publishers Inc.
- [5] J. Barklund and A. Hamfelt. Hierarchical representation of legal knowledge with metaprogramming in logic. *Journal of Logic Programming*, 18(1):55–80, 1994.
- [6] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.
- [7] L.W. Barsalou. Situated conceptualization. *Handbook of categorization in cognitive science*, 619:650, 2005.
- [8] J. Barwise and J. Seligman. *Information Flow: the logic of distributed systems*. Cambridge University Press, 1997.
- [9] T. Bench-Capon, JB Freeman, H. Hohmann, and H. Prakken. Computational models, argumentation theories and legal practice. *Argumentation Machines: New Frontiers in Argument and Computation*, pages 85–120.

- [10] T. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1-2):97–143, 2003.
- [11] T.J.M. Bench-Capon, S. Doutre, and P.E. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171(1):42–71, 2007.
- [12] TJM Bench-Capon and G. Sartor. A quantitative approach to theory coherence. *Legal Knowledge and Information Systems. Jurix*, pages 53–62, 2001.
- [13] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):28–37, 2001.
- [14] L. Birnbaum. Rigor mortis. *Artificial Intelligence*, 47:57–77, 1991.
- [15] B. Bix. HLA Hart and the open texture of language. *Law and Philosophy*, 10(1):51–72, 1991.
- [16] B. Bix. *Jurisprudence: theory and context*. Sweet & Maxwell, 1999.
- [17] A. Bondarenko, P.M. Dung, R.A. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1-2):63–101, 1997.
- [18] J. Breuker and R. Hoekstra. Core concepts of law: taking common sense seriously. In *Formal ontology in information systems: proceedings of the third conference (FOIS-2004)*, page 210. Ios Pr Inc, 2004.
- [19] J. Breuker, A. Valente, and R. Winkels. Legal ontologies in knowledge engineering and information management. *Artificial Intelligence and Law*, 12(4):241–277, 2004.
- [20] Joost Breuker and Nienke den Haan. Separating world and regulation knowledge: where is the logic. In *ICAAIL '91: Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 92–97, New York, NY, USA, 1991. ACM.
- [21] G. Brewka, F. Roelofsen, and L. Serafini. Contextual default reasoning. In *20th International Joint Conference on Artificial Intelligence (IJCAI07)*, pages 268–273, 2007.

- [22] A. Bundy and M. Chan. Towards ontology evolution in physics. In W. Hodges, editor, *Procs. Wollic 2008*. LNCS, Springer-Verlag, July 2008.
- [23] S. Buvac, V. Buvac, and I.A. Mason. Metamathematics of contexts. *Fundamenta Informaticae*, 23(2/3/4):263–301, 1995.
- [24] D.V. Carbogim, D. Robertson, and J. Lee. Argument-based applications to knowledge engineering. *The Knowledge Engineering Review*, 15(02):119–149, 2000.
- [25] R. Carnap and R.A. George. *The logical structure of the world and pseudoproblems in philosophy*. Open Court Pub Co, 2003.
- [26] W. Ceusters. Applying evolutionary terminology auditing to the Gene Ontology. *Journal of Biomedical Informatics*, 42(3):518–529, 2009.
- [27] W. Ceusters, B. Smith, A. Kumar, and C. Dhaen. Ontology-based error detection in SNOMED-CT. In *Proceedings Medinfo 2004*, pages 482–486.
- [28] M. Chan and A. Bundy. Inconstancy: An ontology repair plan for adding hidden variables. In *Symposium on Automated Scientific Discovery, number FS-08-03 in Technical Report*, pages 10–17.
- [29] A. Chorley and T. Bench-Capon. AGATHA: Using heuristic search to automate the construction of case law theories. *Artificial Intelligence and Law*, 13(1):9–51, 2005.
- [30] H.H. Clark. *Using Language*. Cambridge University Press, 1996.
- [31] D. Davidson. A Coherence Theory of Truth and Knowledge. *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, pages 307–319, 1986.
- [32] P. De Cruz. *Comparative law in a changing world*. Cavendish Pub Ltd, 1999.
- [33] S. De Saeger and A. Shimojima. Contextual Reasoning in Agent Systems. *Lecture notes in Computer Science*, 4371:171, 2007.
- [34] J. des Rivieres and H.J. Levesque. The consistency of syntactical treatments of knowledge. In *Proceedings of the 1986 Conference on Theoretical aspects of reasoning about knowledge*, pages 115–130. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1986.

- [35] K.J. Devlin. *Logic and Information*. Cambridge University Press, 1991.
- [36] F.M. Donini, M. Lenzerini, D. Nardi, and A. Schaerf. Reasoning in description logics. *Principles of Knowledge Representation*, pages 191–236, 1996.
- [37] P.M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [38] R. Dworkin. *Law's empire*. Belknap Pr, 1986.
- [39] E. Elhauge. *Statutory default rules: how to interpret unclear legislation*. Harvard Univ Pr, 2008.
- [40] Guzior D.F. Farnsworth, W. and A. Malani. Ambiguity about Ambiguity: An Enquiry Inquiry into Legal Interpretation. *Journal of Legal Analysis (forthcoming)*, 2009.
- [41] G. Flouris, D. Manakanatas, H. Kondylakis, D. Plexousakis, and G. Antoniou. Ontology change: classification and survey. *The Knowledge Engineering Review*, 23(02):117–152, 2008.
- [42] KD Forbus. *Qualitative physics: Past, present, and future*. 1988.
- [43] G. Frege. *Begriffsschrift, Eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. English translation in *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931 - Jean van Heijenoort (ed)*. Harvard University Press, Harvard, 1879.
- [44] R.K. Gardiner. *Treaty interpretation*. Oxford University Press, 2008.
- [45] A. Gardner. *An artificial intelligence approach to legal reasoning*. MIT Press Cambridge, MA, USA, 1987.
- [46] M.R. Genesereth and N.J. Nilsson. *Logical foundations of artificial intelligence*. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1987.
- [47] C. Ghidini and F. Giunchiglia. Local Models Semantics, or contextual reasoning= locality+ compatibility. *Artificial Intelligence*, 127(2):221–259, 2001.
- [48] F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, 16:345–364, 1993.

- [49] F. Giunchiglia and P. Bouquet. A Context-Based Framework for Mental Representation. *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society-CogSci*, 98, 1998.
- [50] Fausto Giunchiglia, Pavel Shvaiko, and Mikalai Yatskevich. S-match: an algorithm and an implementation of semantic matching. In Y. Kalfoglou, M. Schorlemmer, A. Sheth, S. Staab, and M. Uschold, editors, *Semantic Interoperability and Integration*, number 04391 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2005. <<http://drops.dagstuhl.de/opus/volltexte/2005/37>> [date of citation: 2005-01-01].
- [51] L. Groarke. Informal logic. *Stanford Encyclopedia of Philosophy (Autumn 2007 edition)* Edward Zalta (ed.).
- [52] T.R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
- [53] N. Guarino. Formal ontology in information systems. In *Formal Ontology in Information Systems: Proceedings of the 1st International Conference*. IOS Press, 1998.
- [54] R.V. Guha. *Contexts: a formalization and some applications*. PhD thesis, Stanford University, 1991.
- [55] C.D. Hafner and D.H. Berman. The role of context in case-based legal reasoning: teleological, temporal, and procedural. *Artificial Intelligence and Law*, 10(1):19–64, 2002.
- [56] J. C. Hage. Formalizing legal coherence. In *ICAAIL '01: Proceedings of the 8th international conference on Artificial intelligence and law*, pages 22–31, New York, NY, USA, 2001. ACM.
- [57] U. Hahn and N. Chater. Understanding similarity: a joint project for psychology, case-based reasoning, and law. *Artificial Intelligence Review*, 12(5):393–427, 1998.
- [58] J. Haidt. The emotional dog and its rational tail. *Psychological Review*, 108(4):814–834, 2001.

- [59] J. Haidt and J. Graham. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1):98–116, 2007.
- [60] A. Hamfelt. Formalizing multiple interpretation of legal knowledge. *Artificial Intelligence and Law*, 3(4):221–265, 1995.
- [61] HLA Hart. *The Concept of Law*, Clarendon Law Series, 1961.
- [62] PJ Hayes. The naive physics manifesto. *Expert Systems*, pages 242–270.
- [63] John Henderson and Trevor Bench-Capon. Dynamic arguments in a case law domain. In *ICAIL '01: Proceedings of the 8th international conference on Artificial intelligence and law*, pages 60–69, New York, NY, USA, 2001. ACM.
- [64] J.R. Hobbs, M. Stickel, P. Martin, and D. Edwards. Interpretation as abduction. In *Proceedings of the 26th annual meeting on Association for Computational Linguistics*, pages 95–103. Association for Computational Linguistics Morristown, NJ, USA, 1988.
- [65] A.J.I. Jones and M. Sergot. Deontic logic in the representation of law: Towards a methodology. *Artificial Intelligence and Law*, 1(1):45–64, 1992.
- [66] Y. Kalfoglou and M. Schorlemmar. Ontology mapping: the state of the art. *The Knowledge Engineering Review*, 18(01):1–31, 2003.
- [67] Ken Kress. Coherence. In Dennis Patterson, editor, *A Companion to Philosophy of Law and Legal Theory*, pages 533–552. Blackwells, 1996.
- [68] L. Laera, V. Tamma, J. Euzenat, T. Bench-Capon, and T. Payne. Reaching agreement over ontology alignments. In *The Semantic Web - ISWC 2006*, Lecture Notes in Computer Science.
- [69] W.M. Landes and R.A. Posner. Legal Precedent: A Theoretical and Empirical Analysis. *The Journal of Law and Economics*, 19(2):249, 1976.
- [70] P. Legrand. European Legal Systems Are Not Converging. *International and Comparative Law Quarterly*, pages 52–81, 1996.
- [71] D.B. Lenat. CYC: A large-scale investment in knowledge infrastructure. 1995.

- [72] L. Lessig. *Free culture: How big media uses technology and the law to lock down culture and control creativity*. Penguin, 2004.
- [73] D.N. MacCormick and R.S. Summers. *Interpreting statutes: a comparative study*. Dartmouth, 1991.
- [74] C. Matuszek, J. Cabral, M. Witbrock, and J. DeOliveira. An introduction to the syntax and content of Cyc. In *Proceedings of the 2006 AAAI Spring Symposium on Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering*, pages 44–49, 2006.
- [75] J. McCarthy. First Order Theories of Individual Concepts and Propositions. *Machine Intelligence*, pages 129–147, 1979.
- [76] J. McCarthy. Generality in artificial intelligence. *Communications of the ACM*, 30(12):1030–1035, 1987.
- [77] J. McCarthy. Notes on formalizing context. *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1:555–560, 1993.
- [78] J. McCarthy. Modality, si! Modal logic, no! *Studia Logica*, 59(1):29–32, 1997.
- [79] J. McCarthy. Approximate objects and approximate theories. *KR2000: Principles of Knowledge Representation and Reasoning, Proceedings of the Seventh International conference*, pages 519–26, 2000.
- [80] L. Thorne McCarty and N. S. Sridharan. The representation of an evolving system of legal concepts: Ii. prototypes and deformations. In *IJCAI*, pages 246–253, 1981.
- [81] LT McCarty. A language for legal Discourse I. basic features. In *Proceedings of the 2nd international conference on Artificial intelligence and law*, pages 180–189. ACM New York, NY, USA, 1989.
- [82] L.T. McCarty. On the role of prototypes in appellate legal argument (abstract). *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 185–190, 1991.
- [83] L.T. McCarty. An implementation of *Eisner v. Macomber*. *Proceedings of the fifth international conference on Artificial intelligence and law*, pages 276–286, 1995.

- [84] L.T. McCarty. Ownership: A case study in the representation of legal concepts. *Artificial Intelligence and Law*, 10(1):135–161, 2002.
- [85] D. McDermott. Artificial Intelligence Meets Natural Stupidity. ACM SIGART Newsletter 57. Reprinted in Haugeland, J.(Ed.) 1981]. *Mind Design*, 1976.
- [86] D. McDermott. A critique of pure reason 1. *Computational Intelligence*, 3(1):151–160, 1987.
- [87] Fiona McNeill. *Dynamic Ontology Refinement*. PhD thesis, University of Edinburgh, 2006.
- [88] C. Mellish and X. Sun. The Semantic Web as a linguistic resource: opportunities for Natural Language Generation. In *Research and development in intelligent systems XXII: proceedings of AI-2005, the twenty-fifth SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence, Cambridge, UK, December 2005*, page 77. Springer-Verlag New York Inc, 2006.
- [89] P. Mika. Ontologies are us: A unified model of social networks and semantics. *ISWC*, 3729:522–536, 2005.
- [90] G.A. Miller. WordNet: a lexical database for English: New horizons in commercial and industrial AI. *Communications of the ACM*, 38(11):39–41, 1995.
- [91] E.T. Mueller. *Commonsense reasoning*. Morgan Kaufmann, 2006.
- [92] G.L. Murphy. *The Big Book of Concepts*. Bradford Books, 2002.
- [93] E. Nagel and J.R. Newman. *Godel's Proof*. Routledge, 2005.
- [94] I. Niles and A. Pease. Towards a standard upper ontology. In *Proceedings of the international conference on Formal Ontology in Information Systems-Volume 2001*, pages 2–9. ACM New York, NY, USA, 2001.
- [95] P. Norvig and S. Russell. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2003.
- [96] C. S. Peirce. *Collected papers of Charles Sanders Peirce / Vol. 2, Elements of logic*. Harvard University Press, Mass, 1932.

- [97] C. Perelman, L. Olbrechts-Tyteca, J. Wilkinson, and P. Weaver. *The new rhetoric: A treatise on argumentation*. University of Notre Dame Press South Bend, IN, 1969.
- [98] J. Perry and J. Barwise. *Situations and Attitudes*. Cambridge, Mass: MIT Press, 1983.
- [99] J.L. Pollock. *Cognitive carpentry: A blueprint for how to build a person*. Mit Press, 1995.
- [100] D. Poulin, P. Bratley, J. Frémont, and E. Mackaay. Legal interpretation in expert systems. In *Proceedings of the 4th international conference on Artificial intelligence and law*, pages 90–99. ACM New York, NY, USA, 1993.
- [101] H. Prakken. AI & Law, logic and argument schemes. *Argumentation*, 19(3):303–320, 2005.
- [102] H. Prakken. Formalising ordinary legal disputes: a case study. *Artificial Intelligence and Law*, 16(4):333–359, 2008.
- [103] H. Prakken and G. Vreeswijk. Logics for defeasible argumentation. In D.M. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic*, volume 4, pages 218–319. Kluwer Academic Publishers, 2002.
- [104] H. Putnam. *Mind, language, and reality*. Cambridge University Press Cambridge, 1975.
- [105] W.V.O. Quine and J.S. Ullian. *The Web of Belief*. Random House New York, 1978.
- [106] A.L. Rector, W.D. Solomon, W.A. Nowlan, TW Rush, PE Zanstra, and WMA Claassen. A terminology server for medical language and medical information systems. *Methods of Information in Medicine-Methodik der Information in der Medizin*, 34(1):147–157, 1995.
- [107] R. Reiter. Default logic. *Artificial Intelligence*, 13:81–132, 1980.
- [108] E. L. Rissland and D. B. Skalak. Cabaret: Statutory interpretation in a hybrid architecture. *International Journal of Man-Machine Studies*, 34:pp. 839–887, June 1991.

- [109] E.L. Rissland. AI and similarity. *IEEE Intelligent Systems*, 21(3):39–49, 2006.
- [110] E.L. Rissland and K.D. Ashley. A note on dimensions and factors. *Artificial Intelligence and Law*, 10(1):65–77, 2002.
- [111] E.L. Rissland, K.D. Ashley, and RP Loui. AI and law: a fruitful synergy. *Artificial Intelligence*, 150(1-2):1–15, 2003.
- [112] E.L. Rissland and M.T. Friedman. Detecting change in legal concepts. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 127–136. ACM New York, NY, USA, 1995.
- [113] E. Rosch and C.B. Mervis. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4):573–605, 1975.
- [114] G. Schreiber, H. Akkermans, and A. Anjewierden. *Knowledge engineering and management: the CommonKADS methodology*. the MIT Press, 1999.
- [115] MJ Sergot, F. Sadri, RA Kowalski, F. Kriwaczek, P. Hammond, and HT Cory. The British Nationality Act as a logic program. 1986.
- [116] F.M. Shipman and C.C. Marshall. Formality considered harmful: Experiences, emerging themes, and directions on the use of formal representations in interactive systems. *Computer Supported Cooperative Work (CSCW)*, 8(4):333–352, 1999.
- [117] P. Shvaiko, F. Giunchiglia, M. Schorlemmer, F. McNeill, A. Bundy, M. Marchese, M. Yatskevich, I. Zaihrayeu, B. Ho, V. Lopez, et al. OpenKnowledge Deliverable 3.1.: Dynamic ontology matching: a survey. Technical report, Technical report, Informatica e Telecomunicazioni, University of Trento, 38050 Povo-Trento (Italy), Via Sommarive 14, 2006.
- [118] N. Siegel, B. Shepard, J. Cabral, and M. Witbrock. Hypothesis Generation and Evidence Assembly for Intelligence Analysis: Cycorps Nooscape Application. In *Proceedings of the 2005 International Conference on Intelligence Analysis (IA 2005)*(McLean, VA, USA, 2005).
- [119] D.B. Skalak and E.L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial intelligence and Law*, 1(1):3–44, 1992.

- [120] B. Smith. Beyond concepts: Ontology as reality representation. *Proceedings of FOIS. International Conference on Formal Ontology and Information Systems.*, pages 73–84, 2004.
- [121] B. Smith. Ontology. In Luciano Floridi, editor, *Philosophy of Computing and Information*, pages 155–66. 2004.
- [122] L.M. Solan. Vagueness and Ambiguity in Legal Interpretation. In J Gotti M. Bhatia, V.K. Enberg and D. Heller, editors, *Vagueness in Normative texts*. Peter Lang Bern, 2005.
- [123] J. Sowa. Peirce’s Contributions to the 21st Century. *Lecture Notes in Computer Science*, 4068:54, 2006.
- [124] J.F. Sowa. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. MIT Press, 2000.
- [125] J.F. Sowa. The Challenge of Knowledge Soup. *Proc. First International Word-Net Conference, Goa, India*, 2004.
- [126] M.Q. Stearns, C. Price, KA Spackman, and AY Wang. SNOMED clinical terms: overview of the development process and project status. *JOURNAL OF BIOMEDICAL INFORMATICS*, 35:662–666, 2001.
- [127] L. Steels. Semiotic dynamics for embodied agents. *IEEE Intelligent Systems*, 21(3):32–38, 2006.
- [128] M. Stefik. *Introduction to knowledge systems*. Morgan Kaufmann, 1995.
- [129] M. Steup. Epistemology. *Stanford Encyclopedia of Philosophy (Autumn 2007 edition) Edward Zalta (ed.)*.
- [130] L. Stojanovic, A. Maedche, B. Motik, and N. Stojanovic. User-driven ontology evolution management. *Lecture Notes in Computer Science*, pages 285–300, 2002.
- [131] L. Stojanovic and B. Motik. Ontology evolution within ontology editors. In *Proceedings of the International Workshop on Evaluation of Ontology-based Tools (EON2002)*.

- [132] Ljiljana Stojanovic, Alexander Maedche, Nenad Stojanovic, and Rudi Studer. Ontology evolution as reconfiguration-design problem solving. In *K-CAP '03: Proceedings of the 2nd international conference on Knowledge capture*, pages 162–171, New York, NY, USA, 2003. ACM Press.
- [133] R. Studer, V.R. Benjamins, and D. Fensel. Knowledge engineering: principles and methods. *Data & Knowledge Engineering*, 25(1-2):161–197, 1998.
- [134] A. Tarski. The semantic conception of truth: and the foundations of semantics. *Philosophy and phenomenological research*, pages 341–376, 1944.
- [135] P. Thagard. *Conceptual Revolutions*. Princeton Univ Pr, 1992.
- [136] P. Thagard. *Coherence in Thought and Action*. Mit Pr, 2000.
- [137] P. Thagard. Causal inference in legal decision making: Explanatory coherence vs. bayesian networks. *Applied Artificial Intelligence*, 18(3):231–249, 2004.
- [138] P. Thagard and K.A. Verbeurgt. Coherence as Constraint Satisfaction. *Cognitive Science*, 22(1):1–24, 1998.
- [139] R.H. Thomason. A note on syntactical treatments of modality. *Synthese*, 44(3):391–395, 1980.
- [140] B.L.E.E. Tim. *Weaving the Web: the orginal design and ultimate destiny of the World Wide Web by its inventor*. Harper.
- [141] S.E. Toulmin et al. *The uses of argument*. Cambridge University Press Cambridge, 2003.
- [142] W. Twining and D. Miers. *How to Do Things With Rules (3rd Edition)*. Butterworths, 1991.
- [143] C. Unwin. An object model for use in oral and written advocacy. *Artificial Intelligence and Law*, 16(4):389–402, 2008.
- [144] André Valente, Joost Breuker, and Bob Brouwer. Legal modeling and automated reasoning with online. *Int. J. Hum.-Comput. Stud.*, 51(6):1079–1125, 1999.
- [145] Saskia van de Ven, Joost Breuker, Rinke Hoekstra, and Lars Wortel. Automated legal assessment in owl 2. In *Proceeding of the 2008 conference on*

- Legal Knowledge and Information Systems*, pages 170–175, Amsterdam, The Netherlands, The Netherlands, 2008. IOS Press.
- [146] P.R.S. Visser and T.J.M. Bench-Capon. A comparison of four ontologies for the design of legal knowledge systems. *Artificial Intelligence and Law*, 6(1):27–57, 1998.
- [147] F. Waismann. Verifiability. *Proceedings of the Aristotelian Society*, 19:119–150, 1945.
- [148] D.N. Walton. *Informal logic: A handbook for critical argumentation*. Cambridge University Press, 1989.
- [149] D.N. Walton. *The new dialectic: Conversational contexts of argument*. Univ of Toronto Pr, 1998.
- [150] Richard W. Weyhrauch. Prolegomena to a theory of mechanized formal reasoning. *Artif. Intell.*, 13(1-2):133–170, 1980.
- [151] Y. Wilks. Ontotherapy: or how to stop worrying about what there is. Invited presentation, Ontolex 2002. In *Workshop on Ontologies and Lexical Knowledge Bases, 27th May. Held in conjunction with the Third International Conference on Language Resources and Evaluation-LREC02*, pages 29–31, 2002.
- [152] J.M. Williams and G.G. Colomb. *The craft of argument*. Longman, 2006.
- [153] T.D. Wilson. *Strangers to ourselves*. Belknap Press of Harvard University Press, 2002.
- [154] L. Wittgenstein. *Philosophical Investigations*, tr. *GEM Anscombe (Oxford: Blackwell)*, 1953.
- [155] M. Wooldridge, P. McBurney, and S. Parsons. On the meta-logic of arguments. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 560–567. ACM New York, NY, USA, 2005.
- [156] A. Wyner, T. Bench-Capon, and K. Atkinson. Arguments, values and baseballs: Representation of Popov v. Hayashi. pages 151–160, 2007.

- [157] H. Yoshino. The systematization of legal meta-inference. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 266–275. ACM New York, NY, USA, 1995.
- [158] M. Zander. *The law-making process*. Cambridge Univ Pr, 2004.
- [159] J. Zeleznikow and A. Stranieri. The split-up system: integrating neural networks and rule-based reasoning in the legal domain. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 185–194. ACM New York, NY, USA, 1995.