A Top-Level Language-Biased Legal Ontology

Jaspreet Shaheed
Department of Computing,
Imperial College London,
SW7 2BZ, UK
jss00@doc.ic.ac.uk

Alexander Yip
Department of Computing,
Imperial College London,
SW7 2BZ, UK
ay1@doc.ic.ac.uk

Jim Cunningham
Department of Computing,
Imperial College London,
SW7 2BZ, UK
rjc@doc.ic.ac.uk

ABSTRACT

Top-level ontologies have recently been the subject of much research, as they facilitate the representation of commonsense necessary for human-like understanding and reasoning. Almost all domains rely heavily on commonsense, law being no exception. In this paper we present a natural language biased top-level ontology extended with respect to the legal domain. We review the notion of ownership, and ultimately show how our ontology can aid in reasoning about this notion in natural language.

Keywords

Top-Level Ontology, Legal Ontology, Legal Reasoning, Natural Language Processing, Ownership

1. INTRODUCTION

Almost all domains rely heavily on commonsense, law being no exception. The desire to represent this commonsense knowledge within the legal domain led McCarty to develop LLD, a language for legal discourse (detailed in [4]). More recently, ontologies, and in particular top-level ontologies have attempted to fulfil the same aim. The design of the LRI-Core top-level ontology (described in [3]) was motivated by the experience of working with numerous ontologies for specific legal domains. It abstracts from the details presented by particular ontologies, but when sub-categorised further may subsume them.

We begin this paper by presenting an extended top-level ontology which we will call NM-L, it extends to the legal domain work done by Schneider ([7], see also [8]), seeking to found the thematic analysis of natural language in a commonsense view of reality. This suits our aim, facilitating commonsense reasoning but particularly commonsense for natural language. We conclude the section by a brief comparison with the LRI-Core top level ontology [1].

In section 3 we develop a notion of ownership, taking our lead from notions developed by Hohfeld and McCarty. Ultimately ownership is formalised in a way designed to suit our intended natural language applications.

Having described our ontology, in section 4 we aim to show that it is useful. We demonstrate, by way of example, that the distinctions and elements of our ontology facilitate a naive understanding of natural language. Natural language sentences of an example describe ownership issues, and reasoning proceeds with the aid of the formalisation developed in the previous section.

Finally, in section 5 we conclude by summarising our work, discussing its limitations and suggesting some possible directions for future work.

2. THE ONTOLOGY

NM-L is an extension of an NM core which has no legal bias. The NM ontology (described in [7] and [8]) was motivated by the descriptive metaphysics of Strawson, and Parson's thematic roles. We first describe its constituents, before our legally motivated extensions. Finally, we introduce the LRI-Core top-level ontology (detailed in [1]) as a subject for comparison.

2.1 Metaphysics

Strawson's descriptive metaphysics is concerned with describing *common sense*. Central to metaphysics is the treatment of *particulars*. These are the identifiable entities that 'ground' world knowledge, and thus are essential to its representation.

Strawson distinguishes at least three types of particulars: *basic*, which given the primacy of space-time to our conceptual framework, he describes as material bodies; *private*, such as sensations, and finally *theoretical constructs*, which 'play a central role in the general economy of thought', e.g. 'economic growth', or 'ownership'.

While the NM core ontology does not question the notion of spatially featured entities (*substances*) as basic, it does not consider they should be regarded as more basic than temporally featured entities (*occurrences*). This revision is based upon observations of natural language, in which substances depend upon occurrences for their temporal characteristics, just as occurrences rely upon substances for their spatial characteristics.

Thus, the NM ontology's basic particular-types include both substances and occurrences. Note that by substances, we refer to both physical and conceptual entities (e.g. theoretical constructs). Meanwhile, we regard experiencing sensations (instances of private particulars) as occurrences.

Having determined the basic particular-types, the task of subcategorising them remains.

Substances are subcategorised into persons and non-persons. Persons are regarded as a special type of substance, as they possess both physical characteristics (e.g. height, weight, etc), and mental characteristics (e.g. thoughts, feelings, judgements, etc). This at least partially reflects Strawson's notion of a person as 'neither an animated body nor an embodied anima, but the un-analysable subject of both'.

Occurrences are subcategorised, into *States of affairs* and *Events or processes*. Furthermore, these are then subcategorised by *Private* and *Public*. This distinction effectively splits occurrences into those which are observable, and those which are not. Private would thus refer to mental states and events. Note that a more thorough description of the ontology's treatment of occurrences is presented in [7].

These ideas coalesce to form a portion of the top level ontology, as shown in figure 1.

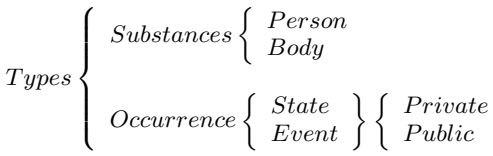

Figure 1: Types in NM

2.2 Linguistic

A way of understanding the meaning of natural language sentences is to try and understand their underlying occurrences. Tesnière suggested that typical English sentences are 'built around' their main verb, as shown in figure 2.

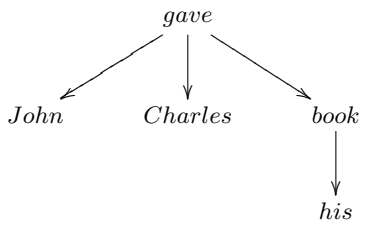

Figure 2: Tesnière's Dependency Grammar

Besides the main verb itself, Tesnière sees the rest of the sentence as fulfilling one of two roles:

- *Complement roles:* These are the 'arguments' to the verb, i.e. the semantics of the occurrence. They correspond to the subject, object, and so on, in traditional grammars.
- *Supplement roles:* These are the circumstances of the occurrence; 'on Monday', 'outside', etc. They correspond to adverbials in traditional grammar.

In his work on thematic roles, Parsons looked in more depth at the roles that parts of a sentence, in particular its noun-phrases, could acquire. He found that most noun-phrases associated with verbs could be given one of a small number of roles, these included:

- *Agent:* Person deliberately causing the event, e.g. "*John* writes a book" or "The book is signed by *John*".
- *Theme:* Entity affected by the event, or in the state, e.g. "Mary reads *a book*" or "*Mary* is wise".
- *Instrument:* Thing the event is accomplished with, e.g. "John opens the letter with *a knife*".
- Also *Goal*, *Benefactive*, *Experiencer*, *Performer*. These are detailed in [6].

Subsequently, the existence of many more roles has been posited, but most extend rather than alter, Parson's list.

It remains to describe the subcategorisation of 'Roles'. The process of determining this subcategorisation was inspired by the work of Guarino and Welty (some elements of which are described in [2]). The result is shown in figure 3.

NM is thus split into 'Types' and 'Roles'. These inherit from 'Kinds', which assumes the position of the root of the ontology.

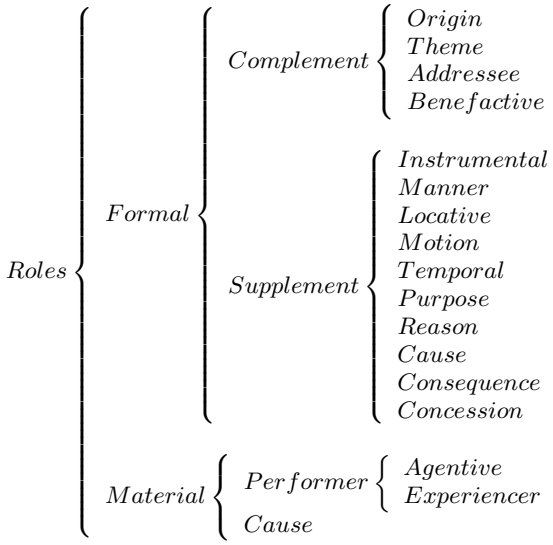

Figure 3: Roles in NM

2.3 Extensions

In order to develop NM into a form more amenable to legal reasoning (NM-L), several extensions were made.

These changes are described in more detail in [9].

2.3.1 Substances

We replace the subcategorisation *persons* and *bodys* with *animate* and *inanimate*. *Animate* is a broader characterisation than *person*, as it may also include electronic agents, and organisations. Additionally, we extend the depth of our understanding by drawing a distinction between *physical*, and *social* substances. *Social* substances are themselves subcategorised by *legal* substances. These extensions allow a more explicit categorisation of *ownership* (an *inanimate legal substance*), and the ability to capture the notion of *legal persons* (in NM, an organisation with a legal status would have been difficult to categorise).

These changes also mean that an entity can now be part of more than one category in the ontology; 'John' may be both a *person* and a *legal person*. We see these two categorisations of 'John' as representing two different senses which are suitable in different contexts. Note that *legal persons* may include organisations with a legal status, which is why we do not regard them as a subcategory of *person*.

2.3.2 Occurrences

Occurrences are separated by the categories *mental*, *environmental* and *social*. *Social* occurrences themselves are separated into *legal*, and *communicative* occurrences.

The *environmental* and *mental* categories of NM-L correspond loosely to *public* and *private* of NM. *Communicative events* are based on the theory of communicative acts, and are made separate both because of their utility in a social-legal context and the difficulty in categorising them elsewhere. They also introduce an asymmetry into NM-L, as we do not regard 'communicative states' as a category. The category *legal states* include *owns* whilst *legal events* include *buys* and *sells*.

2.3.3 Roles

Roles are extended so as to include legal roles, e.g. *represent*, which is the role a legal person assumes when representing another legal person in the context of an occurrence.

Thus NM-L is presented in figure 4. In table 1 we present elements which populate some of the categories the ontology distinguishes.

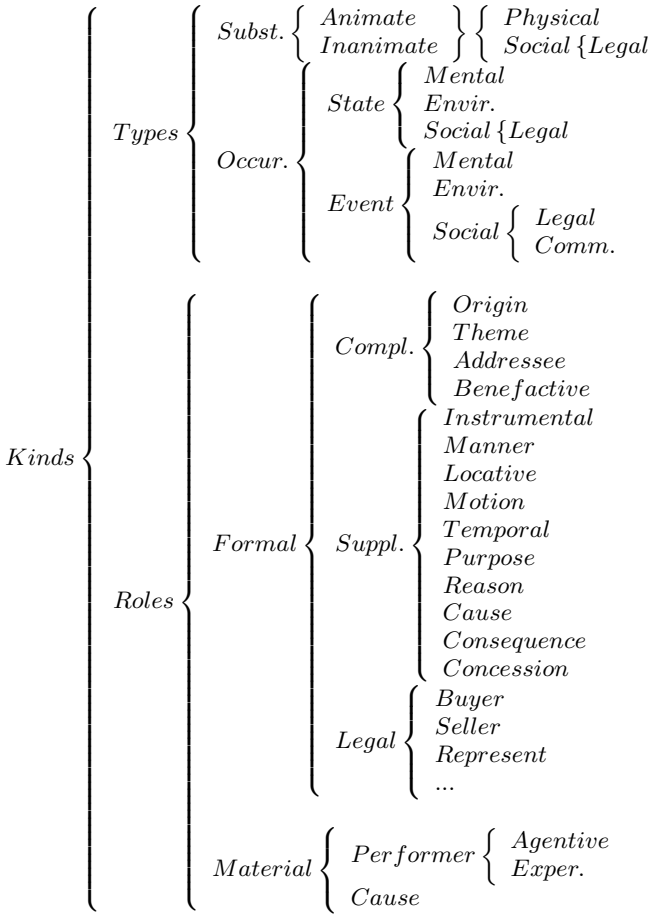

Figure 4: The NM-L Ontology

2.4 The LRI-Core Ontology

LRI-Core (detailed in [1]) is a top-level ontology. It was motivated by experiences from 'ten different legal domains over the course of a decade', and the realisation of the importance of common sense to the law. LRI-Core thus aims to 'reflect the concepts that are used to interpret the day to day world'.

Whilst LRI-Core's ultimate aim is similar to ours, its sources of inspiration are broader and include studies in cognitive science. The top two levels of the ontology are shown in figure 5, although many categories within the ontology have been subcategorised to a much greater depth.

2.4.1 Similarities

The similarities between LRI-Core and our own ontology arise from a shared 'cognitive bias' and 'commonsense objec-

Lexical Element	Semantic Category
Person	Physical animate
Legal Person	Legal animate
Jurisdiction	Legal inanimate
Object	Legal inanimate or physical inanimate
Agent	Physical animate
Money	Physical inanimate
Ownership	Legal inanimate
Right	Legal inanimate
Obligation	Legal inanimate
...	
Possesses	Physical state
Owns	Legal state
Transfers	Legal event
Delegates	Legal event
Buys	Legal event
Sells	Legal event
Borrows	Legal event
Hires	Legal event
Pays	Legal event
...	
Authority	Legal role
Represents	Legal role
Consideration	Legal role
Delegate	Legal role
Possessor	(Derived) legal role
Owner	(Derived) legal role
Seller	(Derived) legal role
Mandator	(Derived) legal role

Table 1: Elements of the Ontology

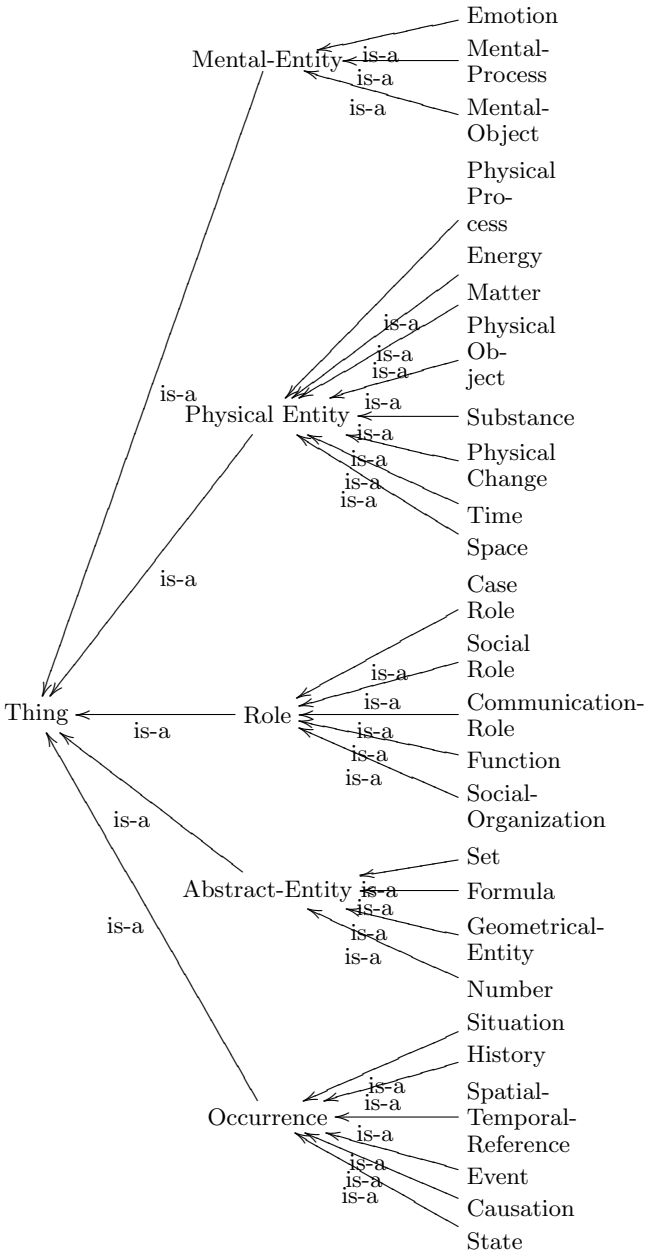

Figure 5: LRI-Core, Top Two Layers

tive'. These similarities include that, whilst the ontologies disagree precisely as to what these are, physical entities, occurrences and roles are separated at a high-level. Furthermore, occurrences are split into (amongst other things) events and states.

2.4.2 Differences

Perhaps the greatest difference between the ontologies lies in the understanding of roles. The natural language bias of our ontology means that LRI-Core is more broadly featured within this category. Another effect of our language bias is that we are more heavily focussed on categorising occurrences, since a Tesnière-like view relates occurrences to sentences. Hence, our ontology distinguishes communicative events rather than communicative roles

Another difference between LRI-Core and our own top-level ontology is our understanding of *physical*. In many respects, commonsense reifies conceptual objects, and our physical intuitions are often transferrable to them. We therefore do not regard the distinction between the physical and mental categories as being at the top level.

LRI-Core subcategorises physical entities in more ways than our ontology does. These categories include physical processes and physical change. Our ontology does not regard a physical process as a physical-entity (which we equate to our 'substance'). Our notion of physical is completely divorced of time. We regard a process as an occurrence, although, we would regard 'snapshots' of this process as substances. For instance, the process of water evaporating from a glass would be regarded as an occurrence, but at any point, one can regard the glass of water as a substance.

In comparing our ontology with LRI-Core, NM-L makes fewer distinctions. The minimalist nature has made the introduction of legal concepts relatively easy. It is not apparent how we would include more abstract or complex constructs, and the extension of our ontology to accommodate these requires further thought. However, as Breuker says in discussing LRI-Core, *'Developing a single coherent upper ontology may not be a feasible solution at all'*.

3. OWNERSHIP

This section aims to define what we mean by ownership. We first examine Hohfeld and McCarty's notions, before describing our own naive notion.

3.1 Hohfeldian Ownership

Hohfeld presented a number of legal *rights*, which regulate and direct the behaviour of the agents to whom they apply. These are outlined in table 2.

Rights	Correlatives	Opposites
claim	duty	no-claim
liberty	no-claim	duty
power	liability	no-power
immunity	disability	liability

Table 2: Hohfeldian Rights

These rights are legal relations, between two agents, viewed by the agent that will generally benefit from them. Correlatives describe the perspective of the other agent. For example, if Alex has a *claim* on Simon to wash the dishes, then Simon has a *duty* to wash them. Claims can also be made in regards to refraining from actions ('don't eat my cake'). *Power* meanwhile describes an agent's ability to bring about a change in some legal relationship, affecting another agent.

As presented in [5], using these relations we can describe ownership (of a physical object) as consisting of the following Hohfeldian rights:

1. A claim against other agents to exclusive physical control of the object, i.e. other agents have a *duty* not to use the object in any way, or take any actions that would harm or destroy it.
2. A *liberty* to use, consume or destroy the object.
3. A *power* to transfer some or all of these rights to another agent.
4. An *immunity* from the involuntary expropriation of these rights by other agents.

3.2 McCarty's Ownership

McCarty's notion of ownership (described in [5]) uses the deontic operators, permitted, forbidden, and obligatory, as well as an additional operator 'enabled', to represent Hohfeld's notion of ownership, these operators are presented below:

- *Permitted*, represented as P(ϕ,α), a 'free-choice' permission meaning that, given the condition ϕ, any way of doing the action α is permitted.
- *Forbidden*, represented as F(ϕ,α), if ϕ holds, then α is forbidden.
- *Obligatory*, represented as O(ϕ,α) meaning under the condition ϕ, the action α is obligatory
- *Enabled*; represented as E(ϕ,α) where condition ϕ enables action α to be undertaken. This operator speaks of an agent's ability to do an action as opposed to the actions permissibility.

Having decided upon the base operators, McCarty gives types to possible actions: α(Y,X) represents an action performable by Y that happens to benefit X. β(Y,X) represents an action performable by Y that happens to be detrimental to X. ΔR(Y,X) represents an action in which Y brings about a change in the legal relationship R, affecting X.

McCarty is now able to represent Hohfeld's legal conceptions, which are described by the rules below.

- *Duty*, O(ϕ,α(Y,X)), represents that Y has a duty to do $\alpha(Y, X)$. From X's point of view, if ϕ is true, X has a claim against Y for the performance of α.

- *Duty*, F(ϕ,β(Y,X)), represents that Y has a duty not to do β(Y,X). From X's point of view, X will have a claim against Y if ϕ is true, and Y performs β.
- *Liberty*, $\neg$F(ϕ,$\beta(Y, X)$), represents that Y is at liberty to perform β.
- *Liberty*, $\neg$O(ϕ,$\alpha(Y, X)$), tells us that Y is at liberty not to perform α.
- *Power*, E(ϕ,ΔR(Y,X)), tells us that Y has a power over X with respect to ΔR(Y,X), in other words, X endures the liability of Y's actions with respect to ΔR(Y,X).
- *Immunity*, $\neg$E(ϕ,ΔR(Y,X)), tells us that X has an immunity from Y with respect to Y performing ΔR(Y,X).

Note that the role of 'enabled' now becomes clear. Agents are not merely forbidden from expropriating the ownership rights of others, they are simply unable to.

Assuming that α, β and Δ now additionally pertain to, or involve the object under consideration, and 'O' and 'A' are the owner and another agent respectively, ownership can be described as below. Note that for conciseness we assume ϕ in the descriptions below.

1. F(ϕ,β(A,O)) and F(ϕ,α(A,O)): represents the exclusive physical control of the object.
2. $\neg$F(ϕ,$\beta(O, A)$) and $\neg$O(ϕ,$\alpha(O, A)$): represents the liberty to use, consume, or destroy the object.
3. E(ϕ,ΔR(O,A)): represents power over the legal relation between the object and A.
4. $\neg$E(ϕ,ΔR(A,O)): represents the inability of A to effect a change in the legal relationship between themselves and the object.

3.3 Naive Ownership

We adopt a simpler notion of ownership than McCarty for two reasons.

First, we believe it is difficult to categorise actions into those which are 'beneficial' (α) and 'detrimental' (β). Not only are these subjective analyses, but an action may switch categorisations depending upon context.

Secondly, we envisage applications in which we wish to reason about ownership in excerpts from natural language. In these excerpts, we can imagine an occurrence in which an agent sells an object they don't own. The buying agent may well perceive no difference between this occurrence, and one in which the selling agent truly owns the object. However, employing McCarty's representation, there is a fundamental difference and the former is not possible. This is correct from a legal standpoint as no legal relationships have been changed. However, this does not help us reason about the occurrence (and in particular, the bogus seller's culpability).

We are concerned with commonsense and describing rather than prescribing behaviour. Given this perspective, we would prefer to revert to a definition McCarty earlier adopted, in which the operator 'enabled' was instead 'not forbidden'. Using this definition, actions that were previously not possible are now regarded as merely forbidden.

This change leads to problems in that it effectively creates a group of legal occurrences which effect no change in legal relationships (legal states). In the longer term, there is no doubt that a finer grained representation is necessary. However, for our present purposes, this suffices and greatly simplifies our presentation.

Ultimately then, we employ just two deontic operators, forbidden and permitted.

1. Forbidden: for other agents to use/harm/destroy the object.
2. Permitted: for the owning agent to use/consume/destroy the object.
3. Permitted: for the owner to alter the legal rights pertaining to the object.
4. Forbidden: for other agents to involuntarily expropriate (or perceive to involuntarily expropriate) these legal rights.

4. OWNERSHIP ISSUES IN LANGUAGE

In this section, we show how the NM-L may benefit simple legal reasoning[1]. We begin by presenting a simple example which features ownership issues expressed in natural language. We show how we can recognise these issues using rules composed of thematic roles and a typing of occurrences. Finally, we show how we might axiomatise our rules in Prolog.

4.1 A Simple Example

Our example monologue consists of a number of sentences of simple construction. We assume they are temporally ordered, so that each succeeds the previous. Each sentence features one 'main verb', making it amenable to analysis in the style of Tesnière. Clearly in the example presented, John has behaved illegally.

1. "John bought a car from a shop."
2. "John sold the car to Mary because he was going to Paris."
3. "John took the car from Mary."
4. "John went to Paris with the car."
5. "John sold the car to Mike."
6. "Mary went to Paris."
7. "Mary took the car from Mike."
8. "Mary was angry."

[1]An implementation of this is available for download at: `www.doc.ic.ac.uk/~jss00/legal/`

4.2 Ownership Through Roles

In section 3.3, we presented a naive notion of ownership of an object, which consisted of four statements describing what was permitted, and what was forbidden.

Other agents (i.e. not the owner), were forbidden from using/harming or destroying the object, and also involuntarily expropriating (or perceiving to involuntarily expropriate) the legal rights pertaining to it. It seems that we can often recognise occurrences of these actions by the presence of the following thematic-role and occurrence-type markers.

- Any occurrence in which the agent is not the owner and assumes the role 'agent', and the object assumes an 'instrumental' role, for example "John went to Paris *with the car*".
- Any 'environmental' occurrence in which the agent is not the owner and assumes the role 'agent', and the object assumes the role 'theme', for example "John broke the car".
- Any 'legal' occurrence in which the agent is not the owner at the beginning or end of the occurence and assumes the role 'agent', and the object assumes the role 'theme', for example "John sold the car", where John is not the owner of the car.

The first point is relatively easily explained. An 'instrumental' role generally represents use of an object.

The second point is more defeasible. Environmental occurrences relate to the physical world. They may not necessarily be harmful, but in our view, they have the potential to be construed as such. This contrasts with mental occurrences ('liked', 'thought about', etc), which it is difficult to perceive as such. This point is discussed further in section 5.

The third point aims to forbid the involuntary expropriation of rights by other agents. This is problematic as legal rights can be voluntarily ceded, for example, "John sold the car to Mary" in our example. We would like to permit Mary to buy or inherit a car, whilst forbidding her from selling or donating a car she doesn't own. We note however that at the 'end' of the 'legal' occurrences of buying, or inheriting a car (the occurrences we wish to permit), Mary *is* the owner, hence our additional check.

The other portion of the definition of ownership described those actions permitted for agents owning an object. We believe this can be recognised in a similar fashion.

- Any occurrence in which the agent is the owner and assumes the role 'agent', with the object assuming an 'instrumental' role.
- Any 'environmental' occurrence in which the agent is the owner and assumes the role 'agent', and the object assumes the role 'theme'.
- Any 'legal' occurrence in which the agent is the owner and assumes the role 'agent', with the object assuming the role 'theme'.

Note that in considering which thematic roles apply, we have only considered one form of the sentences, in which the agent of interest assumes the role 'agent'. For example, our rules recognise wrongdoing in "John sold the car to Mike", but not "Mike bought the car from John". An extension to these correlative sentences would not be difficult, but is omitted here.

Reviewing our initial example, John's violation of ownership law on a number of points now becomes clear:

- In sentence 3, he assumes an 'agent' role in an 'environmental' occurrence for a 'theme' he does not own. This indicates a possible instance of harming/destroying the car.
- In sentence 4, an object he does not own assumes an 'instrument' role in an occurrence in which he is an agent. This indicates a possible instance of use of the car.
- In sentence 5, he assumes an 'agent' role in a 'legal' occurrence, with a 'theme' that he owns neither at the beginning, nor at the end of the occurrence. This indicates a possible involuntary expropriation of rights over the object.

4.3 Axiomatisation

In this section, we axiomatise the rules given in the previous section using Prolog. We begin by defining the occurrences we support, and giving them appropriate types.

```
environmental(O) :- goes(O); takes(O);
  gives(O); enable(O).

legal(O) :- buys(O); sells(O);
  owns(O).
```

Next, we define the actions which affect the model of the world we are interested in. In our example these are 'buying' and 'selling', since these result in a change of ownership. We define these in terms of event-calculus, a brief reminder of some event-calculus axioms is presented in table 3.

```
initiates(E,owns(Buyer,Object),T) :-
  E = buys(Buyer,Object,Seller),
  holdsAt(owns(Seller,Object),T).
terminates(E,owns(Seller,Object),T) :-
  E = buys(Buyer,Object,Seller),
  holdsAt(owns(Seller,Object),T).

initiates(E,owns(Buyer,Object),T) :-
  E = sells(Seller,Object,Buyer),
  holdsAt(owns(Seller,Object),T).
terminates(E,owns(Seller,Object),T) :-
```

```
E = sells(Seller,Object,Buyer),
holdsAt(owns(Seller,Object),T).
```

initiates(α,β,τ)	β holds after action α at time τ
terminates(α,β,τ)	β ceases to hold after action α at τ
happens(α,τ)	Action α occurs at τ
holdsAt(β,τ)	β holds at τ

Table 3: Event Calculus Axioms

We assume that the natural language parser will identify the occurrence type ('main verb') and thematic roles associated with sentences. It associates these with an occurrence `O`, so that `agent(john,O)` means that John assumed the role 'agent' in the occurrence `O`. In the instance where an occurrence is one which affects ownership, ('buy', 'sell', etc.), the state of the world can be updated when the occurrence is recognised.

Finally, we are in a position to axiomatise our rules. In the previous section, we presented three groups of thematic-role and occurrence-type markers which identified an agent as performing an action forbidden to it. We implement these in the Prolog clause as shown below. This clause effectively allows us to check whether any agent violates our notion of ownership within that occurrence.

```
own_viol(O) :-
  % Occurrence O happens at time T
  happens(O,T), agent(Ag1,O),
  own_viol_alt(O,T,Ag1).

  own_viol_alt(O,T,Ag1) :-
    own_use_viol(O,T,Ag1) ;
    own_harm_viol(O,T,Ag1) ;
    own_right_viol(O,T,Ag1).

  own_use_viol(O,T,Ag1) :-
    instrumental(Object,O),
    holdsAt(owns(Ag2,Object),T).
    \+Ag1=Ag2.

  own_harm_viol(O,T,Ag1) :-
    environmental(O),
    theme(Object,O),
    holdsAt(owns(Ag2,Object),T),
    \+Ag1=Ag2.

  own_right_viol(O,T,Ag1) :-
    legal(O),
    theme(Object,O),
    holdsAt(owns(Ag2,Object),T),
    \+Ag1=Ag2, next(T,NextT),
    holdsAt(owns(Ag3,Object),NextT),
    \+Ag1=Ag3.
```

The first two rules are encoded in a relatively straightforward manner, the last is somewhat more difficult. As described in the previous section, the test for a 'valid' legal occurrence for an agent is whether, at the end of the occurrence, it is the owner. In the event calculus however, a *fluent*[2] is true only after the event which initiates it. We thus need to check ownership at the *next* time point. We introduce an axiom `next(T,NextT)` to provide us with this ability, although this limits the temporal models we can assume. In practice, if we wished to assume another temporal model we could subcategorise the relevant legal occurrences.

It remains to axiomatise those rules describing permitted actions. These follow in a similar fashion.

```
own_ok(O) :-
  % Occurrence O happens at time T
  happens(O,T), agent(Ag1,O),
  holdsAt(owns(Ag1,Object),T),
  own_ok_alt(O,T,Ag1).

  own_ok_alt(O,T,Ag1) :-
    own_use_ok(O,T,Ag1) ;
    own_harm_ok(O,T,Ag1) ;
    own_right_ok(O,T,Ag1).

  own_use_ok(O,T,Ag1) :-
    instrumental(Object,O).

  own_harm_ok(O,T,Ag1) :-
    environmental(O),
    theme(Object,O).

  own_right_ok(O,T,Ag1) :-
    legal(O),
    theme(Object,O).
```

5. CONCLUSIONS AND FURTHER WORK

In this paper, we have presented a top-level ontology, NM-L, an extension of the NM ontology so as to make it more suitable for the legal domain. We presented its constituents, and compared it with the LRI-Core top-level ontology. This comparison demonstrated our ontology's natural language bias. We see this as a result of our aim; to facilitate commonsense, but particularly commonsense for natural language processing.

The ultimate test of an ontology is how useful it is. We aimed to show that our top-level ontology can aid in simple legal reasoning. We defined our own notion of ownership, and presented an axiomatisation which employed thematic roles and a typing of occurrences as its basic constituents. We presented a natural language monologue in which a sequence of occurrences was described, and showed that our axiomatisation managed to identify ownership violations within it.

Whilst we looked at ownership in particular, we believe that norms from other legal domains can also be axiomatised in a similar way. More generally, we hoped to show that the distinctions our ontology makes can be combined effectively for commonsense reasoning.

[2]A fact whose truth varies over time

Our example itself also presents many limitations. Consider the sentences :

1. "John drove the car to Paris"
2. "John missed the car with his blow"
3. "John threw himself at the car"

In the first sentence, the car assumes the role 'theme'. We would recognise this as a violation of ownership, but not because John was *using* the car, but (less appropriately), because he was *harming* the car. We know that driving a car means using it, but our ontology currently makes no provision for this kind of commonsense knowledge.

In the second sentence, we have an environmental occurrence, 'missing', which clearly demonstrates that the car was *not* harmed, yet our axiomatisation would label this as potentially doing harm to the car. This problem is representative of our distinctions being too coarse-grained for in-depth reasoning. In response to this, we would note that the presentation of this paper is intended to serve as an intuitive guide to the validity of the distinctions we have already made, not to preclude further distinctions.

The third sentence poses two difficulties. The first is, whilst John threw himself at the car, there is no evidence that the car came to harm. In practice, the current behaviour of the system is probably correct, since it is possible the car would be damaged. The second problem, more pertinent to our present discussion, is that the car assumes neither the role 'theme', or 'instrument', but clearly could have come to harm. The occurrence fails to be caught because the damage to the car is not the focus of the sentence, instead the focus is on the act of throwing (consider instead "John hit the car with his body"). Natural language, even in the very simple examples we have considered, can quickly become difficult to process. In less restricted language, it is likely that we would need to complement our axiomatisations with more complex analyses of sentences.

5.1 Further Work

There are several directions we would like to take our work in. Initially, we wish to consolidate our current work by further exploring the language examples we can 'process' with respect to our axiomatisation. Additionally, we would like to test our ability to axiomatise more complex and realistic notions of ownership, in particular making the link between the Hohfeldian or Deontic concepts more explicit.

We also wish to go beyond the mere facts of the occurrences, and begin to build an understanding of the mental states, and the reasoning behind them. These are clearly important for understanding the magnitude of an agent's culpability.

1. Reason: "John drove to Paris with the car *because of the emergency.*"
2. Manner: "John drove to Paris with the car *shamefully* because of the emergency."
3. Consequence: "John was so afraid *that he drove to Paris with the car.*"

The sentences above demonstrate how roles that our ontology distinguishes may aid us in more complex reasoning about agents. We would like to establish a link between these roles, and a more BDI-style view of agents. For instance, the sentences above contradict the notion that John took the car out of malice, or because he intended to profit by the occurrences.

More generally, we hope to take the ontology presented into a broad range of domains. In particular, we are interested in those involving elements of physicality and social interaction. We hope that an exploration of this orthogonal direction will lead to an ontology which makes better and deeper commonsense distinctions.

6. REFERENCES

[1] Joost Breuker. Constructing a legal core ontology: LRI-Core. In *Workshop on Ontologies and their applications*, 2004.

[2] N. Guarino and C. Welty. Identity and subsumption. In *Semantic Relations*, 2001.

[3] Radboud Winke Joost Breuker. Use and reuse of legal ontologies in knowledge engineering and information management. In *ICAIL Workshop on Legal Ontologies & Web Based Legal Information Management*, 2003.

[4] L. Thorne McCarty. A Language for Legal Discourse I. Basic Features. In *Second International Conference on Artificial Intelligence and Law*, 1993.

[5] L. Thorne McCarty. Ownership: A case study in the representation of legal concepts. In *Artificial Intelligence and Law*, 2002.

[6] T. Parsons. *Events in the Semantics of English: A Study in Sub-atomic Semantics*. MIT Press, 1990.

[7] Luc Schneider. Naive Metaphysics. Technical report, Department of Computing, Imperial College London, 2002.

[8] Luc Schneider and Jim Cunningham. Ontological Foundations of Natural Language Communication in Multi-Agent Systems. In *Knowledge Based Intelligent Information and Engineering Systems*, 2003.

[9] Alexander Yip and Jim Cunningham. Ontological Issues in Agent Ownership. In *LEA Workshop 2003: The Law and Electronic Agents*, 2003.