

Interpreting Contracts Using English Common Law Rules as stated by Lord Hoffmann

John Henderson
Beale and Company, Solicitors
London
UK

Trevor Bench-Capon
Dept. of Computer Science
University of Liverpool
UK

ABSTRACT

This paper describes a computational procedure for interpreting contracts in accordance with the English common law rules of interpretation of contract as stated by Lord Hoffmann. Our approach makes extensive use of an ontology of legal terms, specialised for the context in which the contract was made. We illustrate the approach using three examples closely based on actual cases decided by Lord Hoffmann.

1. INTRODUCTION

This paper describes a computational model of the interpretation of contract in accordance with the English common law rules of interpretation as stated by Lord Hoffman¹ in a landmark decision². It uses an ontology to provide the basic meaning of the terms in the contract, which are refined and instantiated in accordance with the background information given by the particular context in which the contract was made and the objective intention of the parties (as would be ascribed by a reasonable observer). An interpretation is considered valid if it, together with the contextually instantiated terms, allows this intention to be derived.

Experience shows that interpretation of contract is a major part of what many lawyers do, a task which is a frequent cause of disputes: [8] records that “*interpretation of contract remains the most important source of commercial litigation*”.

Legal interpretation, including interpretation of contracts, has long been the subject of study in AI and law, including the first published book in the field [18]. However, despite its importance to the legal profession and its clients, there is surprisingly little recent work on interpretation in AI and Law focussed on contracts. More recent work on legal interpretation has tended to focus on interpretation of statute ([16] and [4]). Contracts are very different from statutes, because contracts are *written* in a particular context to bind particular parties for a particular purpose, whereas statutes

¹Leonard Hubert Hoffmann, Baron Hoffmann, PC (born 8 May 1934) is a retired senior British judge. He served as a Lord of Appeal in Ordinary from 1995 to 2009.

²Investors Compensation Scheme Ltd. v West Bromwich Building Society [1997] UKHL 28

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ICAIL '17, June 12 - 16, 2017, London, UK

Copyright 2017 ACM 978-1-4503-3522-5/15/06 ...\$15.00.

<http://dx.doi.org/10.1145/2746090.2746103>.

are written to have a general application and intended to cover situations not even envisaged by the legislator, although they are, of course, *applied* to particular situations. Moreover, while the contract is taken to have a single meaning fixed at the time of signing, the meaning of a statute may evolve as it is applied by the courts to a sequence of cases. Statutes can even adapt to changing social norms.

Recent work on the problems arising out of vagueness and ambiguity in contracts has tended to focus on preventing those problems from arising in the first place either by careful drafting [1] or by making the contracts ‘computable’³. In contrast, the approach taken in this paper is to focus on how the vagueness and ambiguity inherent in natural language contracts and the errors that occur in them may be resolved by interpretation.

Our approach is based on the English common law rules of interpretation, set out below. The approach taken by the English common law is to relate the words of a contract back to the context in which the contract was formed - “language always takes meaning from context” (Lord Hoffman⁴). So, for example, the word ‘claims’ in the contractual phrase ‘all claims are settled’ is correctly interpreted to refer to those claims that would have been in the reasonable contemplation of the parties to the contract at the date when they entered into it.

We claim that this work has a direct, practical application in the jurisdiction of the United Kingdom in that experience shows that human lawyers have trouble giving priority to the contextual meaning of a contractual word over its “natural and ordinary meaning” (for example, the phrase ‘all claims’ refers naturally and ordinarily to *all* claims, not just those contemplated by the parties at the time of signing). A computational procedure which helped to overcome this reluctance⁵. would be of great assistance to UK Lawyers.

³The Stanford Computable Contracts Initiative. See <https://law.stanford.edu/projects/stanford-computable-contracts-initiative>

⁴In his decision on Investors Compensation Scheme Ltd v West Bromwich Building Society [1997].

⁵That this reluctance is not confined to lawyers is illustrated by the famous contract in Shakespeare’s *The Merchant of Venice*, by which Antonio agrees to allow Shylock to take a pound of flesh if he is unable to pay the money owed. Portia saves Antonio by telling Shylock: *Therefore prepare thee to cut off the flesh. Shed thou no blood, nor cut thou less nor more But just a pound of flesh: if thou cut'st more Or less than a just pound, be it but so much As makes it light or heavy in the substance, Or the division of the twentieth part Of one poor scruple, nay, if the scale do turn But in the estimation of a hair, Thou diest and all thy goods are confiscate..* This may be the literal meaning of the words but no reasonable observer (i.e the audience) would believe it to be the objective intention when the contract was signed. Indeed Portia’s solution comes as a complete surprise to anyone who has not previously seen the play. None the

The English common law rules of interpretation are particularly well suited to being modelled for three reasons. First, because they are prescriptive and second because they exclude from consideration the subjective intention of either of the parties to the contract: it is not the claims that were subjectively intended to be settled that are denoted in the example above, but the claims that would have been reasonably understood to be settled by the reasonable person who, at the time the contract was entered into, had access to all the background information that formed the context for the formation of the contract. Third, under the rules there is considered to be a single statement of intention, and a single set of contextual information, and therefore there should be a single interpretation. This contrasts with statutes which are supposed to be applicable to range of cases, and where the interpretation may need revision as further cases are decided [13].

In overall summary, our approach to modelling the English rules of interpretation is to treat the contractual phrase in question as a structure taken from an ontology of legal terms, refined if needed to fit the required context, which is then instantiated from the background information in accordance with the constraints imposed by the ontology. For example, the contractual word 'claims' would be represented by the legal class *Claim* taken from the ontology and instantiated using the background information so as to satisfy the definitions in the ontology.

Experience and a review of case law shows that disputes over interpretation of contract tend to fall into three (possibly overlapping) types:

- disputes over the scope of vague terms;
- disputes over ambiguous terms and
- disputes over linguistic errors.

The primary claim that we make in this paper is that these three types of linguistic uncertainty can all be explained and resolved using our model. In section 6 we will refer to leading cases which exemplify each of these three types of linguistic uncertainty and show how, in each case, our model could be used to resolve that uncertainty with the same outcome as was given in the opinions of Lord Hoffman, the judge who restated the modern English law of contract interpretation and who has done more than anyone else to develop it.

The rest of this paper is organised as follows. Section 2 sets out the English common law rules of interpretation of contract as stated by Lord Hoffmann. Section 3 describes our language, in particular the nature of the ontology we use. Section 4 gives details of our model. Section 5 discusses related work. Section 6 gives three worked examples based on actual cases. Section 7 offers some concluding remarks.

2. THE RULES OF INTERPRETATION OF CONTRACT

There is a wide range of rules of interpretation, since they differ from jurisdiction to jurisdiction and within any jurisdiction they may differ as between different types of document, for example, constitution, legislation, contract, etc. They are probably all worth study in AI and Law, both in themselves and comparatively. Here we will limit ourselves to the English common law rules for the interpretation of contracts. Even within those rules there is a range of different ways in which they may be stated. However, we believe that the principles set out by Lord Hoffman in the case of *less, Portia's interpretation is not challenged, neither by the audience, nor even by Shylock.*

*Investors Compensation Scheme v West Bromwich Building Society*⁶ are foundational to any statement of those rules.

Lord Hoffman's opinion in that case states:

“(1) Interpretation is the ascertainment of the meaning which the document would convey to a reasonable person having all the background knowledge which would reasonably have been available to the parties in the situation in which they were at the time of the contract.

(2) The background was famously referred to by Lord Wilberforce as the “matrix of fact”⁷, but this phrase is, if anything, an understated description of what the background may include. Subject to the requirement that it should have been reasonably available to the parties and to the exception to be mentioned next, it includes absolutely anything which would have affected the way in which the language of the document would have been understood by a reasonable man.

(3) The law excludes from the admissible background the previous negotiations of the parties and their declarations of subjective intent. They are admissible only in an action for rectification. The law makes this distinction for reasons of practical policy and, in this respect only, legal interpretation differs from the way we would interpret utterances in ordinary life. The boundaries of this exception are in some respects unclear. But this is not the occasion on which to explore them.

(4) The meaning which a document (or any other utterance) would convey to a reasonable man is not the same thing as the meaning of its words. The meaning of words is a matter of dictionaries and grammars; the meaning of the document is what the parties using those words against the relevant background would reasonably have been understood to mean. The background may not merely enable the reasonable man to choose between the possible meanings of words which are ambiguous but even (as occasionally happens in ordinary life) to conclude that the parties must, for whatever reason, have used the wrong words or syntax. (see *Mannai Investments Co Ltd v Eagle Star Life Assurance Co Ltd*⁸).

(5) The ‘rule’ that words should be given their ‘natural and ordinary meaning’ reflects the common sense proposition that we do not easily accept that people have made linguistic mistakes, particularly in formal documents. On the other hand, if one would nevertheless conclude from the background that something must have gone wrong with the language, the law does not require judges to attribute to the parties an intention which they plainly could not have had.

Lord Diplock made this point more vigorously when he said in *The Antaios Compania Neviera SA v Salen Rederierna AB*⁹:

⁶*Investors Compensation Scheme Ltd. v West Bromwich Building Society* [1997] UKHL 28

⁷*Reardon Smith Line Limited v Yngvar Hansen-Tangen* [1976] 1 WLR 989 at 995-996.

⁸*Mannai Investments Co Ltd v Eagle Star Life Assurance Co Ltd* [1997] 2 WLR 945

⁹*Antaios Compania Neviera SA v Salen Rederierna AB* 1985 1 AC 191, 201

... if detailed semantic and syntactical analysis of words in a commercial contract is going to lead to a conclusion that flouts business common sense, it must be made to yield to business common sense.”

There are two related characteristics of these rules in particular which result in them being particularly well suited to being modelled as a computational procedure. First, under the rules, interpretation is clearly based on establishing a relationship between the phrase or word being interpreted (‘the contractual phrase’) and a definable set of information which formed the context in which the contract was formed (‘the background information’). Thus, the main task for our procedure is to establish that interpretation relation.

Second, under the rules, the interpretation relation is not established by reference to what the parties, collectively or individually, subjectively intended the relation to be. Subjective intention would be very hard if not impossible to capture and represent given that either or both the parties to the contract could assert, like Humpty Dumpty, that “when I use a word, it means just what I choose it to mean - neither more nor less” [7]¹⁰¹¹. Instead, the interpretation relation is established by reference to what the reasonable person would have understood the parties’ intention to be (their ‘objective intention’). We model this by making three assumptions. First, that the reasonable person’s understanding of the contract and its context can be represented using an ontology of legal terms, which will, as explained in the next section, provide a name for the terms and a list of attributes that define them. These definitions of legal words should be uncontroversial. Second, when the parties use a legal term in a contractual phrase, then their objective intention is to use the matching legal term from the legal ontology. Third, if the background information discloses an objective intention that a particular contractual word or phrase should refer to a particular subclass then the ontology can be refined for the purposes of the current dispute so that the equivalent legal term can refer to that particular subclass.

We can summarise our model, using Lord Hoffman’s idea of the reasonable person, as follows. We assume that the reasonable person has an understanding of law which is represented as an ontology of legal terms of the sort described below.

When the reasonable person sees legal words in a contractual phrase, such as ‘claim’, that match classes in her ontology, she constructs an instance of the appropriate class from the background information. When fully instantiated from the background information the legal words are fully interpreted and can no longer be vague - their interpretation is a list of instances of the legal terms that occur in the background information. If the contractual phrase can be considered ambiguous, then the reasonable observer “chooses between the possible meanings of the words which are ambiguous” by considering which will fulfill to the objective intention of the parties. If there is an error in the contractual phrase, ie, “something must have gone wrong with the language” (to use Lord Hoffmann’s phrase quoted above in rule 5), then the reasonable observer corrects that error by using terms taken from her legal ontology (which are correctly formulated) to represent the contractual phrase rather than the erroneous contractual phrase itself. This procedure is explained in more detail and illustrated with examples in section 6

¹⁰*Through the Looking Glass* is cited in the Court of Appeals judgement in *Investors Compensation Scheme Ltd. v West Bromwich Building Society*.

¹¹A similar problem in relying on “private semantics” arises in communication languages for multi-agent systems which make use of the subjective beliefs and attitudes of the agents concerned, such as [12]. See [20] for a discussion of the issues.

below.

3. THE ONTOLOGY USED

We have above referred to both “ambiguity” and “vagueness”. These terms have some important differences.

In the case of ambiguity, a phrase has two or more meanings, all of which seem equally plausible, and which need to be distinguished. This distinguishing can be done in terms of a more formal language. Consider the sentence *All members of my family use a telephone*. This could mean (and thirty years ago would almost certainly have meant) that there is some particular telephone that all members of my family use (i.e. a household land line). Nowadays, however, it is far more likely to mean that each member of my family has a telephone which they use (i.e. a personal mobile telephone). In first order predicate calculus the difference is expressed by the scope of the quantifiers:

$$\forall(x)\exists(y)(Fax \rightarrow Uxy)$$

$$\forall(x)(Fax \rightarrow \exists(y)Uxy)$$

In an ontology this would be represented by constraining the telephone slot of instances of my family either to 333-3333 (which is the proper name of some land line telephone, or to #MobilePhone, indicating that the slot can be filled by any instance of the class *MobilePhone*, which may differ from family member to family member. In this way ambiguity can (and *should*) be eliminated.

A vague expression, in contrast, has a single meaning, but its application may be in doubt. This is well illustrated by the *Sorites* paradox¹². The use of vague terms such “heap” is inevitable: we would certainly not want to specify the number of grains required to constitute a heap, and normally all those involved in the conversation will agree as to whether they are in the presence of a heap or not. This is the open texture referred to by Hart [10], which often needs to be resolved by taking the matter to court and getting a decision on the particular facts which can thenceforth serve as a precedent as to what can be counted as a heap. Unlike ambiguity, vagueness should not be eliminated altogether, although a decision is required for a particular situation as and when it arises. It must be recognised that not every possible situation can be envisaged, and so sometimes the application of the term must be resolved in the light of particular facts. In case law, as more and more cases are decided the open texture, the area of uncertainty, is gradually narrowed. In the interpretation of contracts, however, the vagueness is not resolved in this cumulative fashion: we have a particular context, and it is taken that the parties understood what the term meant in the context of that contract. If “heap” is used in a contract, there will be a specific contextual understanding of “heap” on the part of the reasonable observer: the contract is not meant to cover a variety of situations. In contracts we tend to resolve ambiguity by

¹²Described as follows in the Stanford Encyclopedia of Philosophy: The sorites paradox is the name given to a class of paradoxical arguments, also known as little-by-little arguments, which arise as a result of the indeterminacy surrounding limits of application of the predicates involved. For example, the concept of a heap appears to lack sharp boundaries and, as a consequence of the subsequent indeterminacy surrounding the extension of the predicate ‘is a heap’, no one grain of wheat can be identified as making the difference between being a heap and not being a heap. Given then that one grain of wheat does not make a heap, it would seem to follow that two do not, thus three do not, and so on. In the end it would appear that no amount of wheat can make a heap. We are faced with paradox since from apparently true premises by seemingly uncontroversial reasoning we arrive at an apparently false conclusion.

selecting the intended meaning, whereas we resolve vagueness by considering whether the example or examples in question do or do not fall under the contract. This is illustrated by the examples in 6.1 and 6.2. The importance of distinguishing between ambiguity and vagueness when handling uncertainty in law has long been recognised in AI and Law [3].

Our ontologies will take the form described in [17] and [2]. A class will comprise: a name, a list of superclasses, a list of subclasses, a list of attributes and a set of axioms. The subclasses are intended to be disjoint and exhaustive. Attributes comprise a triple: a name and a list of permitted values and a cardinality (although we shall not use cardinality in the discussions in this paper, since it plays no part in our examples). As one moves from class to subclass, attributes may be added, but the permitted values and the cardinality become more restrictive. Thus in moving from *plane figure* to *quadrilateral* the cardinality of *numberOfSides* is restricted to 4, and in moving from *person* to *pensioner*, the value of age is constrained to be greater than 60. A further move to *malePensioner* would constrain both *gender* and *age*. Axioms describe relations between attributes: thus for *RegularQuadrilateral*, *area* is equal to *longSide* multiplied by *shortSide*.

As [2] makes clear, ontologies cannot be universal nor be expected to be task neutral, but rather they represent the conceptualisation of a domain by a particular group of people for a particular purpose, to fulfill a particular task. Thus while we will need an ontology to represent a kind of common sense understanding of language (Wordnet [14] is often used for this purpose), this will need to be supplemented by an ontology designed to reflect the conceptualisation prevalent in the domain of the contract and the peer group of the contracting parties. As noted above, it may even be necessary to introduce subclasses required by the particular dispute. Subclasses are introduced using a *genus* and *differentia* method as in [17]. The *genus* will be the superclass of the new subclass, and the *differentia* will comprise restrictions on one or more attributes. Thus given a class *animal* (note that only the relevant attributes are shown here):

```
Animal:
  skinCover{fur, scales, feathers, hair}
  birthMethod{vivaporous, oviparous}
  legs{0, 2, 4, 6, 8}
```

we may define a subclass *bird* as a feathered biped:

```
Bird:
  skinCover{feathers}
  birthMethod{vivaporous, oviparous}
  legs{2}
  nesting{trees, ground, buildings, other}
```

As well as constraining certain attributes, we have added an attribute (applicable to birds, but not animals in general) to record where the bird nests. This will be specialised as we introduce particular species. Note that this specialisation does not define birds as egg-laying: we may, however, wish to add an axiom to *Animal*:

```
Animal.SkinCover = feathers and
Animal.legs = 2 =>
  Animal.birthMethod = oviparous.
```

This will ensure that all birds are, as a matter of fact, not by definition, egg layers. The constraints on the ontology are designed to ensure that the specialism is strict and that problems of cancellation such as those noted in [6] do not arise.

Based on this ontology we can easily produce sentences in the form of entity-attribute-value triples. Here the entity will be a class

(or instance) name, the attribute will be an attribute of that class and the value will be a valid filler for that attribute of that class: e.g. <jack, instance-of, animal> and <jack, legs, 4> will express that Jack is a quadruped. Where classes appear in entity-attribute-value triples, they are implicitly universally quantified. If an existential quantifier is needed, this can be expressed as a sub-class or, for an individual, the name of an instance. Equally we may write these triples as a logical relation: *attribute(class,value)* if we wish to use them in a logic program. This enable us, if we wish, to move smoothly to predicate calculus, or to a logic program using a language such as Prolog.

4. A DESCRIPTION OF OUR MODEL

The model consists of a legal ontology and a procedure for constructing and instantiating data structures taken from that legal ontology to represent the contractual phrase, its background and the objective intention of the contracting parties. The input is a contractual phrase, the contract from which the contractual phrase is taken and some background information. All of this will be expressed as entity-attribute-value triples: we presuppose the use of some parser (many are available) to extract the information in this form.

We describe each of the elements of the model in more detail as follows.

The contractual phrase is the phrase that is to be interpreted. For example, 'all claims' in the first worked example, below.

The input contract is the contract from which the contractual phrase is taken. For example 'Alice and Bob agree that all claims are settled' is the contract in the first worked example, below.

The ontology comprises a collection of interrelated legal terms that the parties agree may exist and the definitions of which are agreed between the parties. As described in the previous section, the ontology takes the form described in [2] and represents a specific conceptualization of a domain for a task-specific purpose.

While we would expect it to start from a set of general and uncontroversial descriptions of common law terms, we assume the use of only a selected fragment drawn from a larger ontology but sufficient to represent the terms relevant to the dispute. Moreover, we accept that it may need to be refined for each specific dispute: and these refinements need to be accepted by all the parties to the dispute. Perhaps this could best be achieved by embedding the procedure in a dialogue game and have the parties construct the ontology by moves in that game (similar to ontology reconciliation in multi agent systems [9], which is often effected by a negotiation dialogue [19]). We leave this for future work.

Legal Sentences are composed from the ontology and are used to reconstruct the contractual phrases and to assert the information about the contact and the background. As explained above, in our procedure, the structure of a legal sentence will be based on the commonly used format of entity-attribute-value triples.

An example of a class in the ontology is:

```
Claim:
  ClaimName: String
  CauseOfAction: #CauseOfAction
  Remedy: [damages, rescission],
  Claimant: #Person
  Defendant: #Person
```

Thus *Claim* has four attributes, one of which, *Remedy* can take one of two values, *damages* or *rescission*. Thus given an instance of *Claim*, *C1*, <*C1* remedy damages> and <*C1* remedy rescission>

will be legal sentences, where an attribute is associated with the values it is permitted to take for that class. This may be a standard datatype (string, integer, real, etc), or a list of values, or an instance of some particular class (indicated by a '#'). Classes may also be associated with axioms to describe relations between attributes.

The background information describes the context in which the contract was entered into. This includes both the objective intention of the parties as it would have been understood by a reasonable observer and the information expressed as entity-attribute-value triples required to instantiate terms taken from the ontology.

The procedure establishes an interpretative relation between the contractual phrase and the background information by first using the ontology (modified if necessary) to define the contractual phrase, and state the objective intention which the contract would be supposed by the reasonable observer to be designed to achieve. The legal terms pertinent to the situation are then instantiated using the ontology and background, and it is considered whether the intention does indeed follow. This can broadly be seen as forming four steps:

- Step 1. Select from the ontology the classes required to express the contractual phrase, the putative interpretation and the background information. Where required specialise the classes to subclasses using the *genus* and *differentia* method.
- Step 2. Identify the objective intention from the background information. For example, if there was a statement of objective intention that 'claims' refer to 'claims for rescission', then the result should be that correct interpretation will restrict the *remedy* slot of the class *Claim*, to rescission for claims covered by the contract.
- Step 3. Use the background information to instantiate the classes identified in Step 1, so that the information can be expressed in entity-attribute-value triples.
- Step 4. Check that the contract, using the putative interpretation and the background information, allows the intention to be shown to hold.

The ability to derive the desired intention confirms the acceptability of the putative interpretation. Equally if the intention cannot be derived, the putative interpretation cannot be correct.

5. RELATED WORK

Related work falls broadly into two categories, that which addresses the problem of interpretation at the drafting stage by aiming to minimise the need for it and that which represents statutory interpretation in terms of ontological change.

In the first category is the work on logic based drafting such as [1] and the work on computable contracts, such as that carried out at Stanford¹³. Both these approaches are valuable in improving the drafting of contracts and reducing the amount of ambiguity and mistakes. However, our view is that these approaches cannot (and do not purport to) provide a complete answer to the uncertainties of language for at least two reasons. First, because many natural language words and phrases are inherently vague. Well-known examples are the vague descriptions such as 'heap of sand' or 'bald people' which may give rise to a Sorites paradox as described above. Another example is absolute prohibitions, for example, 'no cars allowed' which may give rise to what Lord Hoffmann describes as

¹³<https://law.stanford.edu/projects/stanford-computable-contracts-initiative>

'all or nothing arguments'¹⁴. He illustrates the absurdity of all or nothing arguments as follows (paragraph 65):

"The following conversation may be imagined. A motorist is stopped by a park warden driving down a road which is signposted 'no cars allowed'. He says 'but I am driving a green car'. The warden points out that it is nevertheless a car. The motorist says 'but the words cannot be read literally. Do you suggest they forbid children's toy cars?'. The warden concedes that the context suggests a prohibition for the protection of pedestrians frequenting the park and that it does not apply to toy cars¹⁵. 'And what about police cars going to an emergency? Surely there is an implied exception for emergency vehicles?'. 'Yes, perhaps there is'. 'Well then' says the motorist 'if it cannot be read literally, why should it apply to green cars?'"

Second, we would argue that there is great practical benefit in the vagueness of words. If I want to draw your attention to a particular quantity of sand, it is much more convenient to point and say 'that heap of sand' than it is to specify the number of grains that I am referring to. Equally, ambiguity may arise out of a beneficial economy of language, in that the same word can carry out several tasks, but then it is important that the context leaves no room for doubt as to which sense of the word is intended.

Therefore, we believe that there is a role to play for a procedure of interpretation of natural language contracts however carefully they are drafted. Infelicity and error may be reduced, but it will never be eliminated altogether.

The second category of related work is work on interpretation of statute by ontology revision in [4]. The general idea underlying this work is that the concepts referred to in a statute can be defined in an ontology and, when the statute comes to be applied, the ontological concept may be extended or restricted by reference to the purpose to which the statute was directed. To illustrate, one of the examples in [4] is a law that bans vehicles from a park¹⁶. The purpose of the law is stated to be to reduce pollution. A person is fined for riding a bicycle into a park. This is because, without revision, a bicycle falls within the ontological definition of vehicle. However, under the approach described in [4] that definition can be restricted: in this case as a bicycle does not increase pollution, and so was not intended to be covered by the law. There is strong and helpful, foundational similarity between this work and our procedure in that they both found interpretation on an ontology. In [4] the ontology may be modified by reference to the purposes of the legislation represented in the ontology. In our approach the ontology may be modified according to the statement of objective intention contained in the background information.

[16] does not fit into our approximate categorisation of related work. It sets out a theoretical framework in which the many different types of argument used in the interpretation of statute (eg, argument from precedent, argument from analogy, etc.) can be represented. One of the key advantages we have in representing interpretation of contract as a fixed procedure is that it excludes from consideration many of these types of argument that might otherwise be thought to arise. For example, the interpretation of one

¹⁴Bank of Credit and Commerce SA v Munawar Ali and Others [2001]. UKHL 8

¹⁵The warden's concession is too easy: it is well known in philosophy that a toy train is not a train, any more than toy money is money.

¹⁶The fountainhead of this family of examples is usually given as [10].

case is not influenced by what has happened in previous cases or by analogies between the current case and precedent cases. Where previous cases are cited, we see them as providing examples of the usage of terms (in the manner of the historical quotations in the *Oxford English Dictionary*) and so should not be used as the basis of analogical reasoning. Thus our procedure does not use (or need) any form of case based reasoning.

Ontological change has been discussed in [5], in relation to case law in [15] and in relation to reasoning with cases in AI and Law in [11].

6. THREE EXAMPLES

This section illustrates the operation of our model on three common types of linguistic uncertainty: vagueness, ambiguity and error. Each type of uncertainty is illustrated by reference to a leading English case that concerned that type of uncertainty and the output of our model is measured against the opinion of Lord Hoffman in each of those cases.

For the purposes of this paper, we have somewhat simplified the actual cases in our examples (whilst preserving the type of uncertainty they illustrate).

6.1 Vagueness. Bank of Credit and Commerce SA v Munawar Ali and Others

An actual example of vagueness is found in *Bank of Credit and Commerce SA v Munawar Ali and Others* [2001]¹⁷. This case concerned the interpretation of a settlement agreement under which a bank (BCCI) had settled the claims of its employees who had been made redundant. The agreement stated:

The Applicant ... agrees to accept the terms set out in the documents attached in full and final settlement of all or any claims whether under statute, Common Law or Equity of whatsoever nature that exists or may exist and, in particular, **all or any claims**, rights or applications of whatever nature that the Applicant has or may have or has made [**emphasis added**].

The question at issue in the case was, very broadly speaking, whether that settlement agreement, which is a type of contract, settled a claim for damage to the reputation (stigma) of the Applicant: that claim was not contemplated by the parties at the time they entered into the settlement agreement. Thus there was a dispute about the scope of the term ‘claims’, and so it was vague, as it was not clear whether or not it should include claims for stigma.

Lord Hoffmann found that the settlement agreement did settle the Applicant’s future claims for stigma because, again very broadly speaking, under the settlement agreement the bank paid all the known claims in full and on top of that included an additional payment, in respect of any other claims that might subsequently arise. Lord Hoffmann’s opinion was that the reasonable person would conclude that this additional payment was in release of any other claims, even those not contemplated at the time of the judgement (see paragraph 47 of Lord Hoffman’s judgement).

Contractual phrase. We will simplify the disputed phrase to: ‘all claims between Alice and Bob’.

Input contract. ‘Alice and Bob agree that all claims between Alice and Bob are settled’.

Legal ontology fragment. The fragment of the ontology required will include *LegalRelation*, which states that a legal relationship exists between two people.

¹⁷UKHL 8.

```
LegalRelation
  PartyA: #Person,
  PartyB: #Person
```

LegalRelation will have a subclass, *Contract*, which will inherit the parties, constrain its members to a particular *Type* drawn from a list of different contract types and indicate the date on which the contract was made.

```
Contract
  Type: {employment,
         mortgage,
         investment,
         settlementAgreement}
  DateMade: #Date
```

We also need classes for *Claim* and *BreachOfContract*:

```
Claim
  LegalRelation: #LegalRelation
  CauseOfAction: {#BreachOfContract,
                  #Stigma}
  Remedy {Damages, Recission},
  Claimant: LegalRelation.PartyA
  Respondent: LegalRelation.PartyB
  Settled: {True, False}
```

```
BreachOfContract
  Contract: #Contract
  Claimant: Contract.PartyA
  Respondent: Contract.PartyB
```

Background information. The background information is:

- Alice and Bob intend to settle all their claims;
- there was an employment contract between Alice and Bob. For clarity we indicate instances with the prefix “*”,

```
*Contract1
  Type: employment
  PartyA: *Alice
  PartyB: *Bob
  DateMade: *Date1
```

- Alice made a claim against Bob for breach of contract seeking the remedy of damages in respect of this contract.

```
*Claim1
  LegalRelation: *Contract1
  CauseOfAction: *BreachOfContract1
  Remedy: Damages
  Claimant: *Alice
  Respondent: *Bob
  Settled: {True, False}
```

When the claim is made we cannot say whether it is settled.

Procedure. The procedure would apply as follows:

Step1 Select the relevant classes from the ontology: *Claim*, *Contract*, and *BreachOfContract*.

Step 2: represent the statement of intention ‘Alice and Bob intend to settle all their claims’. In the statement of the objective intention the reasonable observer will take into account the fact that no mention has been made of the cause of action, and that the amount of compensation is more than would be expected from the claims known to the observer at the time at which the contract was made.

```
(<claim claimant Alice> and
 <claim respondent Bob>) =>
 <claim, settled, true>.
```

Note the implicit universal quantification over claim.

Step 3: Use the background to instantiate the classes: We have an instance of *Claim*, **Claim1*, referring to a particular claim relating to a particular breach of contract (**breach1*) of a particular Contract, (**Contract1*), also shown above. Thus the following entity attribute value triples all hold.

```
<*Contract1 instance-of Contract>
<*Contract1 type employment>
<*Contract1 partyA, *Alice>
<*Contract1 partyB, *Bob>

<*Claim1 instance-of claim>
<*Claim1 causeOfAction, breach1>
<*Claim1 remedy damages>
```

When **BreachOfContract1* is instantiated the following also hold.

```
<*breach1 contract, *Contract1>
<*breach1 claimant *Alice>
<*breach1, respondent *Bob>
```

Step 4: Add the contractual phrase and check that the intention is fulfilled. To represent the contract we make a subclass of claim, those between Alice and Bob, which we call *ABClaim*. Now to implement our proposed interpretation of the contract, we constrain the settled slot to true to represent the position that ‘all claims between Alice and Bob’ really does mean all claims between them, without further qualification.

```
ABClaim
  LegalRelation: #Contract
  CauseOfAction: {#BreachOfContract,
                  #Stigma}
  Remedy: Damages
  Claimant: (*Alice, *Bob)
  Respondent: {*Alice,*Bob}
  Settled: True
```

The different interpretations turn on whether we define claim as above, with the *differentia* restricting the parties and the status, or whether we leave Settled as {True, False} in *ABClaim* and add an axiom that

```
ABClaim.CauseOfAction = BreachOfContract
=> ABClaim.Settled = True.
```

The intention will be true if we constrain *Settled* in the definition of the class *ABClaim*, but not if we use the axiom to that class, which would require a third triple in the antecedent of the objective

intention to ensure that the cause of action was breach of contract. Thus the correct interpretation is to use the definition of *ABClaim* shown above rather than that the axiom.

Even though the background information available when the contract was made contains nothing which would allow us to instantiate any stigma claim, fully interpreted, ‘all claims’ refers to all claims between Alice and Bob (i.e all *ABClaims*), no matter what the cause of action. Note that this interpretation would also consider any claims made against Alice by Bob to be settled.

The approach taken by our procedure broadly follows the approach taken by Lord Hoffmann. If the parties had wanted to limit themselves to claims arising out of the employment contract, then they would have said ‘all claims arising out of the employment contract’. The parties were aware of this common law rule at the date they entered into the contract and understood their contract would be interpreted in accordance with it.

6.2 Ambiguity. ICS v WBBS

ICS v WBBS is the case in which Lord Hoffman gave his restatement of the rules of interpretation, as quoted in section 2 above. It concerned a statutory compensation scheme set up to compensate investors who had lost money in a failed investment arrangement under which they had mortgaged their homes to building societies, to secure loans which they then invested on the advice of investment advisors.

In order to receive compensation under the scheme, the investors were required to assign their rights to make claims against the various promoters of the arrangement (such as investment advisors), to a statutory body (‘the ICS’), so that the ICS could make claims against those promoters. The general idea being, very broadly speaking, that the ICS would recover money from the promoters and then distribute it among the investors. However, given that the investors might still have mortgages with the building societies, the objective intention of the parties was that the investor’s rights arising out of their mortgages should *not* be assigned to the ICS.

The assignment under which the investors intended to transfer their rights to the ICS stated that all claims were assigned to the ICS with the exception of “any claim (whether sounding in rescission for undue influence or otherwise)”. The ambiguity in this contractual phrase lies in whether the exception refers to any claim for rescission or to any claim whatsoever (that is whether the ‘otherwise’ relates to the grounds for the rescission, or to the rescission itself). Again, this is a slight simplification of the ambiguity in the case itself. Lord Hoffmann’s opinion was that the background information made it clear that the contractual phrase referred to ‘any claim for rescission’ (ie, the narrower interpretation), the purpose of “otherwise” being to ensure that the exception was not restricted to those arising from undue influence.

Contractual phrase. We will simplify the contractual phrase to: ‘any claim by for rescission for undue influence or otherwise’. The “otherwise” is handled by allowing the appropriate attribute of the subclass of claims (*rescissionClaims*), defined with *claim* as the *genus* and *remedy=rescission* as the *differentia*, which will allow several different values for the grounds (all those valid in the parent class), only one of which is *undue influence*.

Contract. We will simplify the contract to: *Alice agrees to assign to Bob all claims except any claim for rescission for undue influence or otherwise.*

Legal ontology fragment. We will reuse the ontology fragment from the first example with the addition of *RescissionClaim* as a subclass of *Claim*. This will restrict the legal relation to *Contract*, and the remedy to rescission. It will also introduce an attribute *grounds*, to indicate why rescission is appropriate. We also need to

add an attribute to *Claim* to indicate the person to whom the claim is assigned.

```
RescissionClaim
  LegalRelationship: #Contract
  Remedy: Rescission
  Grounds: {UndueInfluence,
           Deception,
           Coercion}
  Assigned-To: #Person
```

where Assigned-to is inherited from the parent class, *Claim*

Background Information The background information is:

- there was a mortgage between Alice and Charles. To record this we will need to have Mortgage as a subclass of *Contract*. We will also need a subclass of *Contract* relating to contracts other than Mortgages.

```
Mortgage
  LegalRelationship: #Contract
  type: mortgage
  property: {#building, #land}
```

Since subclasses are meant to be exhaustive we will need a sibling for Mortgage:

```
OtherContract
  LegalRelationship: #Contract
  type: {investment, employment}
```

```
<*mortgage1 instance-of Mortgage>
<*mortgage1 partyA *Alice>
<*mortgage1 partyB *Charles>
```

- there was also an investment contract between Alice and Charles:

```
<*contract2 instance-of OtherContract>
<*contract2 type investment>
<*contract2 partyA *Alice>
<*contract2 partyB *Charles>
```

- the parties intend that any claim that Alice has against Charles other than for rescission of mortgages are assigned to Bob.

Procedure. The procedure will operate as follows.

Step 1: We have instances of the mortgage, **mortgage1*, and the claim against the investment contract, **contract2*. We also suppose that Alice has a claim against Charles under both of these. We also note that the subclasses of claim are intended to be disjoint and exhaustive. We therefore make another subclass of *claim*, as a sibling to *RescissionClaim*. Moreover we wish to assign all such claims to Bob, and so we constrain the Assigned-to attribute to Bob.

```
DamagesClaim
  LegalRelationship: #Contract
  Remedy: Damages
  Assigned-To: *Bob
```

Step 2: We can state the objective intention of the parties as stated above as being to exclude claims for rescission of mortgage from those assigned to Bob:

```
<claim claimant *Alice> and
<claim respondent *Charles> and
not <claim remedy rescission> =>
  <claim assigned-to *Bob>
```

Step 3. We instantiate the two claims. We need to decide which of the two subclasses to instantiate, and this will depend on the remedy. The claim under mortgage can be an instance of *RescissionClaim* and so not assigned to any particular person, but the claim under **contract2* will be an instance of *DamagesClaim*, since its remedy is damages. The claim under this contract will thus be assigned to Bob.

Step 4: Check that the intention is fulfilled. Since the assigned-to slot of the claim under the investment contract is constrained to Bob, then the intention is fulfilled in respect of this contract. In the case of the mortgage claim, however, the last antecedent does not hold, and so it should not be assigned to Bob. Since this claim is an instance of *RescissionClaim*, and hence the assigned-to attribute is not constrained to any particular person, this is so, and the intention is fulfilled. We have identified exactly the right set of claims, neither going too far by assigning all the claims, nor not far enough by assigning no claims at all.

Thus the contractual phrase is completely interpreted as any claim by the investor for *rescission*, whatever the grounds.

The approach taken in our procedure broadly follows the approach taken in Lord Hoffmann's opinion. He finds that an explanatory note of the contract contained in the background information "says categorically and without qualification" that the investor (Alice in our example) gives up any claim other than a claim for rescission: and thus assigning the claims arising from the investment scheme. Lord Hoffmann also notes that "no lawyer in his right mind who intended simply to say that all claims against WBBS were reserved to the investor, would have used this parenthesis". The intention of the parenthesis seems rather to make it clear (to a reasonable observer) that the grounds are not restricted to undue influence, rather than the exception was intended to cover any claim whatsoever.

The objection might be made that, in reality, the objective intention of the parties may not be stated clearly enough so as to form a statement that can be used to amend the data structure in question. Our answer to this objection is that if the background information does not provide a definitive statement of objective intention, then the uncertainty could be resolved by a further rule being introduced into the procedure that would give priority, in the absence of objective intention, to the natural and ordinary meaning of the words. This would be in accordance with the fifth of Lord Hoffmann's rules quoted in section 2.

6.3 Linguistic Error. *Chartbrook v Persimmon*

*Chartbrook v Persimmon*¹⁸ concerned the interpretation of a contractual mechanism for calculating the price to be paid to a property developer under a property development transaction which included some residential units. Part of the price was the Additional Residential Payment (ARP). The ARP would vary depending upon how much the residential units in the development were sold for. The agreement defined the ARP as "23.4% of the price achieved for each Residential Unit ('Unit Price') in excess of the Minimum Guaranteed Residential Unit Value ('MGRUV') less the Costs and Incentives ('C and I')."

¹⁸ *Chartbrook Ltd v Persimmon Homes Ltd* [2009] UKHL 38.

Expressed mathematically and following the convention that operations within brackets are carried out before operations outside brackets, then the contract seems to express:

$$A \text{ ARP} = 23.4\% (\text{Unit Price} - (\text{MGRUV} - (\text{C}+\text{I}))).$$

On the facts of the case this would give a value to ARP of some £4.5 million. Whereas broadly speaking, the background information made it clear that the purpose of the ARP was to provide that if a residential unit sold for more than its expected price, then the developer would be entitled to an additional amount by which 23.4% of the higher price exceeded the MGRUV, ie, expressed mathematically,

$$B \text{ ARP} = (23.4\%(\text{Unit Price} - (\text{C}+\text{I})) - \text{MGRUV}).$$

This would suggest that there was an error in the syntax of the ARP, since common sense would suggest that the costs and incentives should be deducted from the actual price (since the incentives are offered in the hope of increasing that price) rather than some notional minimum price.

Contractual Phase. We will simplify the input contractual phrase to that expressed in interpretation A above: ‘*Net Profit = Selling-Price plus Expenses*’: that is we will take MGRUV as \$0, so that the ARP is the whole price. Note that this phrase already suggests the something has gone wrong: as normally understood “Net Profit” would be the selling price after *deduction* of expenses.

Contract. Alice agrees to pay Bob 23.4% of the Net Profit.

Background Information The background information includes

- a statement of intention that the parties intend Bob to be paid 23.4% of the Net Profit. Suppose that subsequently, Alice has earned income of \$100 and incurred expenses of \$10. Bob claims payment of \$25.74 (using formula A, the literal reading of the contractual phrase) and Alice refuses to pay, since formula B would yield Bob only \$21.06. In a deal worth millions, the difference is significant.

The Legal Ontology. contains the object:

```
Deal
  SellingPrice: integer
  Expenses: integer
  NetProfit: integer
Axiom: NetProfit = SellingPrice - Expenses.
```

We also have a notion of debt resulting from a deal:

```
DealDebt
  Source: #Deal
  Creditor: #Person
  Debtor: #Person
  Percentage: real
  Amount: real
Axiom <#Deal Net-Profit NP> =>
  Amount = Percentage * NP.
```

With this input, our procedure would operate as follows:

Step 1: Make an instance of *deal*: *deal1.

```
<*deal1 instance-of deal>
<*deal1 SellingPrice 100>
<*deal1 Expenses 10>
<*deal1 NetProfit 90>
```

Step 2: Identify the statement of intention. The intention is to create a debt arising from a deal, ABDealDebt:

```
ABDealDebt:
  Source: #Deal
  Creditor: Bob
  Debtor: Alice
  Percentage 23.4
  Amount: 23.4 * Source.NetProfit / 100
```

Step 3: create an instance of *ABDealDebt* for *deal1.

```
*dealDebt1
  Source: *Deal1
  Creditor: *Alice
  Debtor: &Bob
  Percentage: 23.4
  Amount: 21.06
```

Step 4: Conclude that this supports interpretation B, so finding for Alice. That interpretation A cannot have been intended (at least by Alice) is shown by considering the case where, because of a collapse in the property bubble, Alice was only able to sell the units by offering incentives that ate up the whole of the selling price. Now, while Alice made no net profit at all, under interpretation A Bob would be able to claim 23.4% of the Selling Price + Incentives, which ‘flouts business common sense’ (to use Lord Diplock’s phrase quoted by Lord Hoffman in rule 5 of section 2).

There might be two immediate objections to this example. First, while the procedure produces the same results as would be produced by Lord Hoffmann’s opinion, it appears to have been achieved by a different approach. The approach in Lord Hoffmann’s opinion is based on the reasonable person’s understanding of the business purpose of the APR, whereas the approach in our procedure is simply to correct the error in the contractual phrase by giving priority to the legal term in the legal ontology over the words in the contractual phrase. However, we believe that this is a difference without substance and, if helpful, the reasonable person’s understanding could be re-cast to say that the purpose of referring to a percentage of the net profit was to give Bob a (reasonable) part of the income. Net Profit is in the ontology because it is an accepted business concept: in contrast, the rejected interpretation contains a concept ((MGRUV - (C+I))) which is not in the ontology because it makes no business sense: for what possible reason would incentives and costs be used to reduce the minimum value? The only effect would be to discourage any advertising and other incentives, which would be in the interests of neither party.

Second, whilst our method of error correction works on this particular error, it would not work for all types of error. For example, if there were an error in the legal word in the contractual phrase, for example, if the contractual phrase said ‘all contracts’ in error when it should have said ‘all claims’. As presently configured, the procedure would call the wrong term and either fail to produce an answer or produce the wrong interpretation. However, we believe that the procedure could be developed to meet this and other types of error.

7. CONCLUDING REMARKS

We review our procedure by reference to the claims that we made for it in the introduction to this paper: first, that it provides a computational model of interpretation of contract able to produce the outcomes found by Lord Hoffmann on three main types of linguistic uncertainty and second, that it would be of practical benefit to UK lawyers.

In respect of the extent to which our model is a computational model of the outcomes of the opinions of Lord Hoffmann, this is shown by the examples in section 6 above.

In respect to the practical value of an implementation based on our model, experience shows that even the most able and experienced practising lawyers sometimes have difficulty in letting go of the natural and ordinary meaning and giving priority to the contextual meaning of a contractual phrase. For example, in the decision of the Court of Appeal in *ICS v. WBBS*, Lord Justice Leggatt found that the interpretation subsequently preferred by Lord Hoffmann was “not an available meaning of the word” (see a reference to this in Lord Hoffmann’s opinion in the House of Lords judgement). Whereas, as Lord Hoffman had previously observed, “Words in themselves do not mean anything: it is people who use words to refer to things”¹⁹. Therefore, it is at least reasonable to suppose that a computer may be better at making a context-based interpretation since it would not suffer from any temptation to impose its own ordinary and natural meaning onto the contractual phrase, and so concentrate on the meaning of the phrase in the context of the particular contract under consideration.

8. REFERENCES

- [1] L. E. Allen and L. J. Lysaght. Modern logic as a tool for remedying ambiguities in legal documents and analyzing the structure of legal documents’ contained definitions. In *Logic in the Theory and Practice of Lawmaking*, pages 383–407. Springer, 2015.
- [2] T. Bench-Capon. Ontologies in AI and Law. In *Inteligencia Artificial Aplicada al Derecho. Memoria del Congreso Internacional de Cultura y Sistemas Jurídicos Comparados. Ciudad de México, Instituto de Investigaciones Jurídicas*, pages 65–85, 2005.
- [3] T. Bench-Capon and M. Sergot. Towards a rule-based representation of open texture in law. In C. Walter, editor, *Computer Power and Legal Reasoning*, pages 39–60. Greenwood Press, 1989.
- [4] G. Boella, G. Governatori, A. Rotolo, and L. van der Torre. A logical understanding of legal interpretation. In *Proceedings of Knowledge Representation 2010*, pages 563–565, 2010.
- [5] G. C. Bowker and S. L. Star. *Sorting things out: Classification and its consequences*. MIT press, 2000.
- [6] R. J. Brachman. I lied about the trees, or, defaults and definitions in knowledge representation. *AI magazine*, 6(3):80, 1985.
- [7] L. Carroll. *Through the looking glass and what Alice Found There*. Macmillan, 1872.
- [8] R. J. Gilson, C. F. Sabel, and R. E. Scott. Text and context: Contract interpretation as contract design. *Cornell Law Review*, 100:23, 2014.
- [9] A. Hameed, A. Preece, and D. Sleeman. Ontology reconciliation. In *Handbook on ontologies*, pages 231–250. Springer, 2004.
- [10] H. L. Hart. The concept of law. *Oxford: Clarendon Press*, 1961.
- [11] J. Henderson and T. Bench-Capon. Dynamic arguments in a case law domain. In *Proceedings of the 8th International Conference on Artificial Intelligence and Law*, pages 60–69. ACM, 2001.
- [12] Y. Labrou and T. Finin. A semantics approach for KQML –a general purpose communication language for software agents. In *Proceedings of the third international conference on Information and knowledge management*, pages 447–455. ACM, 1994.
- [13] E. H. Levi. An introduction to legal reasoning. *The University of Chicago Law Review*, 15(3):501–574, 1948.
- [14] G. A. Miller. Wordnet: a lexical database for English. *Communications of the ACM*, 38(11):39–41, 1995.
- [15] M. Radin. Case law and stare decisis: Concerning” präjudizienrecht in amerika”. *Columbia Law Review*, 33(2):199–212, 1933.
- [16] G. Sartor, D. Walton, F. Macagno, and A. Rotolo. Argumentation schemes for statutory interpretation: A logical analysis. In *Proceedings of Jurix 2014*, pages 11–20, 2014.
- [17] V. Tamma. *An Ontology Model Supporting Multiple Ontologies for Knowledge Sharing*. PhD thesis, University of Liverpool, 2002.
- [18] A. van der Lieth Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. MIT press, Cambridge, Mass., 1987.
- [19] J. van Diggelen, R.-J. Beun, F. Dignum, R. M. van Eijk, and J.-J. Meyer. Ontology negotiation in heterogeneous multi-agent systems: The anemone system. *Applied Ontology*, 2(3-4):267–303, 2007.
- [20] M. Wooldridge. Issues in agent-based software engineering. In *International Workshop on Cooperative Information Agents*, pages 1–18. Springer, 1997.

¹⁹*Mannai Investments Co Ltd v Eagle Star Life Assurance Co Ltd* [1997] 2 WLR 945