

Case Frames as Contextual Mappings to Case Law in BestPortal

Rinke HOEKSTRA ^{a,b} Arno LODDER ^c Frank VAN HARMELEN ^a

^a *AI Department, VU University Amsterdam, The Netherlands*

^b *Faculty of Law, University of Amsterdam, The Netherlands*

^c *Computer Law Institute, Faculty of Law, VU University Amsterdam, The Netherlands*

Abstract. This paper introduces *case frames* as a way to provide a more meaningful structure to vocabulary mappings used to bridge the gap between laymen and legal descriptions of court proceedings. Case frames both reduce the ambiguity of queries, and improve the ability of users to formulate good quality queries. We extend the BestMap ontology with a formalisation of case frame based mappings in OWL 2, present a new version of BestPortal, and show how case frames impact retrieval results compared to simple contextual mappings and a direct fulltext search.

1. Introduction

The linked open data movement is currently experiencing a surge of attention from government organisations and businesses alike. Metadata about published information is increasingly being made available as machine-processable semantic annotations in RDF or as embedded RDFa (e.g. the Data.gov.uk portal). Explicit metadata in the form of Linked Data, allows for a much more flexible way of accessing and searching the information that available online. Raw data is given a context by embedding it in a network of interlinked datasets. By connecting a metadata vocabulary to those of other parties, information can become disclosed through a multitude of different metadata standards. Together with natural language processing technology for automatically annotating documents or extracting vocabularies, this can increase the quality and accessibility of information. However, publishing metadata as Linked Data does not readily accommodate the fact that different communities may, and most often will, use different vocabularies to describe the same data. Arguably, this phenomenon is most evident in the domain of Law, where citizens and businesses face a constant struggle to interpret the regulations they are expected to live by.

In previous work, we have argued for a *translation approach* to legal information disclosure (see section 2). Through the years this approach has taken various forms. First, as direct ontology-based classification of cases to types of liability in [9] (a simplified variant of [8]). Then, in [7, a.o.], by abstracting from a layman ontology to a legal ontology (similar to [5,4]). And finally, in [3], we argued for a lightweight approach to legal information disclosure using contextual vocabulary mappings based on the BestMap



Figure 1. Accessing court proceedings in BestPortal.

ontology of [2]. This resulted in the first version of BestPortal; a search engine for court proceedings.¹

Experiences in creating contextual mappings led us to introduce *case frames* as a way to provide a more meaningful structure to the mappings. In this paper we describe how case frames both reduce ambiguity and improve the ability of users to adequately formulate their query. We first characterise the general approach (in section 2), we then introduce the case frames (Section 3) and then discuss the implementation of these mappings in BestPortal (Section 4). Section 5 presents the latest incarnation of the system, concluding with a discussion and initial evaluation results in section 6.

2. Translation Approach

The BestPortal is intended to improve the ability of citizens to determine their legal position.² It is a portal that allows a reasonably educated layman to find relevant case law, and assess whether his case has any chance in court. To improve access and transparency of the legal system, the Dutch government publishes its laws and a selection of court proceedings online (<http://www.wetten.nl> and <http://www.rechtspraak.nl> respectively). The Rechtspraak.nl website currently offers search on several fields, including full text search on the body of proceedings.

Finding the proper document using full text search is difficult because of the way in which court proceedings are written down. This holds both for laymen and for lawyers. A layman will be unable to phrase his query in such a way that relevant documents are retrieved. For instance, a search on *dog*, *child* and *bite* may retrieve the one case in which a dog bit a child, but not the plethora of other cases in which the owner of an animal can be held liable for any damages caused to a third party. There is a clear need for an approach where citizens are advised about their position using a language that they understand. A lawyer will also have to deal with the disparateness of terminology between the written law – legislation – that defines a legal concept and the case law that applies it [7].

To overcome these limitations, we introduced a *translation approach* to legal information retrieval [7,3]. First, users express their query in terms of common sense ‘laymen’ concepts. This description is then mapped onto a description using legal concepts. Second, the legal concepts are used to perform a search on an Apache Lucene index.³ Here, it is not the label of the concept that is used, but rather a combination of weighted phrases, a vector or *fingerprint*, that represents not only the possible lexical manifestations, but also *indicative* phrases for the concept in our corpus. This distinguishes our ap-

¹BestPortal is accessible to the public through <http://semweb.cs.vu.nl/best-portal>.

²BATNA Establishment using Semantic Web Technology, see <http://www.best-project.nl>

³See <http://lucene.apache.org>.

proach from more traditional concept-based search (c.f. [1]) where search is performed on the basis of the word senses of synonyms of the concept label. For instance, a query for “strict liability for animals” will include the phrase “the animal’s own energy”, which indicates that the volatility of the animal is discussed in a document (See Figure 1).

Mapping layman descriptions to a legal characterisation of a case is not trivial. In [3] we discussed the *reducibility problem*, the problem that legal concepts cannot generally be reduced to commonsense concepts via subsumption. We therefore proposed to reify the relation: rather than basing the mapping between the two paradigms using a *direct* relation between concepts, we *infer* the mapping given an overlap between the cases that may be *described* using those concepts [2]. Furthermore, it is the *context* in which layman concepts co-occur in the annotation of an individual case that make whether a legal concept is part of a valid characterisation of that case. The BestMap ontology [2] is a framework for specifying context-aware mappings between two (or more) vocabularies as OWL 2 class descriptions.⁴

This translation approach solves two problems. First, the reducibility problem between layman and legal concepts, and secondly, the problem of fulltext search where a single concept can occur in the corpus using a wide range of lexical manifestations. In the following section, we introduce an extension to the contextual mappings of BestMap based on *case frames*. Section 4 describes our experiences in implementing this approach in BestPortal.

3. Case Frames as Contextual Mappings

The BestMap ontology introduces several relations to be used in the description of mappings. The `:d_about` property is not transitive and is used to directly link a description to a concept. The `:about` property is its transitive counterpart and is furthermore transitive over the `skos:broader` relation.⁵ A user’s query to the BestPortal is an OWL individual with `:d_about` relations to relevant concepts in a commonsense SKOS vocabulary.⁶ This individual is then classified as one or more of the mappings defined in the repository, leading to *inferred* annotations in terms of the legal vocabulary. We applied this methodology in the development of mappings for the legal area of unlawful acts.

For example, in the case of the riding school we can define a mapping that translates from (generic) commonsense vocabulary to the legal tort vocabulary as follows:⁷

```
ex:Animal_Map ≡ :about value cv:animal ⊔ :about value cv:damage ⊔
               :about value cv:dangerous_action ⊔ :about value cv:person
⊑ :d_about value lv:animal_owner ⊔
   :d_about value lv:strict_liability_for_animals ⊔
```

⁴See <http://www.best-project.nl/owl/bestmap.owl> and <http://www.w3.org/TR/owl-overview/>.

⁵See [2] and the OWL definitions of the ontology for more details.

⁶SKOS: Simple Knowledge Organization System, see <http://www.w3.org/TR/skos-reference/>.

⁷This example uses a variant of the standard Manchester Syntax, see <http://www.w3.org/TR/owl2-manchester-syntax/>. In this example, `cv` is the prefix for the commonsense vocabulary, an `lv` is the prefix for the legal vocabulary.

:d_about value lv:volatility_of_animal_behavior \sqcap

:d_about value lv:damage_caused_by_animal

Since the concepts *cv:animal* and *cv:person* respectively have *skos:narrower* relations with *cv:dog* and *cv:child*, a case described using those concepts will be classified as an instance of the *ex:Animal_Map* class, leading to new *:d_about* relations with legal concepts.

Concept Roles In earlier work, we discussed that the co-occurrence of commonsense concepts in a case description determines the applicability of legal concepts. And indeed, the mapping above requires all commonsense concepts to be present in the description before the legal characterisation can be inferred. However, it does not take into account what *roles* these concepts play in the description. To see what this leads to, consider a case where a fourteen year old child is bitten by a dog. The above mapping will not be able to distinguish that case with a case where the dog was bitten by the child. Where the former is a case of strict liability that holds the owner of the dog liable for any damage caused, the latter is a potential case of fault concerning the parents of the child.

To what extent does making these roles explicit interfere with our intention to develop a *lightweight* approach to legal information serving? Indeed, including concept roles in case descriptions increases expressiveness. However, it does not necessarily increase the effort required in specifying mappings and can even assist the process by providing a guide along they can be defined. The vocabularies used in specifying the mappings are taxonomies with very distinct branches that correspond to the important aspects of a case. For instance, the commonsense vocabulary that was adopted from [7] contains branches for e.g. *actions*, *actors*, *consequences*, *means* and so on. The taxonomic nature of vocabularies is the source of their power. Not taking this structure into account in the definition of mappings leads to information loss that can easily be avoided. Making the roles explicit will help the modeler in determining what aspects of the case, i.e. what branches of the taxonomy, are of potential importance.

Case Frames Explicit roles have another advantage as a means to enable structured case descriptions that guide users of the portal to provide enough information about their case. If left unguided, a prototypical user will not know what and how much information is necessary for an adequate interpretation of her case. This insight led Winkels and de Bruijn [10] to develop a set of pre-structured templates for case descriptions in labour law, *case frames*. These frames were used to guide volunteers of a labour union in selecting the proper legal interpretation for frequently occurring commonsense terms. Every ambiguous commonsense term is provided with a case frame. They use the example of an ‘on-call contract’ which is not a term in Dutch labour law. The case frame assists the user in determining that there is an ‘individual working agreement’ and whether this agreement constitutes a ‘labour contract’ in the sense of the law. The outcome of this process can then be used in the reasoning process of a more heavyweight legal knowledge based system.

Unfortunately, in the case of unlawful acts we cannot count on the presence of a small number of important commonsense concepts that deserve our individual attention. The domain is simply too diverse. Also, the case frames of [10] are rather tight straight-jackets for case descriptions. These would be a stark contrast with the diversity and ‘softness’ of search results presented by the portal. We therefore looked for a more general

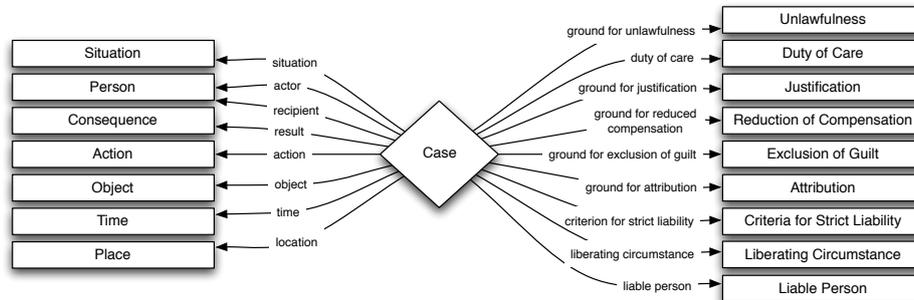


Figure 2. Frames for legal and layman case descriptions

pattern that could be applied in the description of *any* case in the domain of unlawful acts.

The unlawfulness of an action is determined by identifying the individual participants in the action, the situation in which the case takes place, and its consequences (typically some damage). Incidentally, the term ‘case frame’ is also term used in linguistics as a set of named slots for a verb. These slots, or thematic relations, describe the semantic relation of its filler with the verb. Typical thematic relations are *actor* (the agent that initiates the action), and *object* or *patient* (the entity on which the action is performed) and so on.⁸ We adopted the linguistic notion of case frame as a template for case descriptions by laymen. Thematic relations hold between a description and concepts in the vocabulary, rather than between an action and its participants directly. We selected a minimal subset of thematic relations required for these descriptions (see also Figure 2):

Action Relates a description to the action that took place (e.g. ‘hitting’, ‘threatening’).

Actor Relates a description to the agent that performed the action (e.g. ‘child’, ‘dog’).

Object Relates a description to the thing on which the action was performed (e.g. ‘house’, ‘agreement’).⁹

Location Relates a description to the location where it took place (e.g. ‘farmyard’, ‘factory’).

Result Relates a description to the result of the action (e.g. ‘financial damage’).

Situation Relates a description to a special situation in which the action took place (e.g. ‘under duress’, ‘by accident’).

These relations were integrated with the BestMap ontology by specifying them as sub properties of the :d_about property. As the descriptions of a case need to be mapped onto a legal characterisation of the case, we follow a similar procedure in defining a limited set of properties used in that characterisation. This set of properties was extracted from [6] (see Figure 2):

Ground for Unlawfulness Relates the description to the grounds on the basis of which the act is deemed unlawful, e.g. ‘violation of right’. Applicable to almost all cases.

Duty of Care Relates the description to a possible violation of duty of care, e.g. knowledge of a concrete danger.

Ground for Justification Relates the description to possible excuse or grounds for justification, e.g. the right of self-defence. Applicable only in cases of guilt.

⁸See e.g. http://en.wikipedia.org/wiki/Thematic_relation for an overview of important thematic relations.

⁹Strictly speaking this relation is a combination of the thematic relations *patient* and *theme*.

Ground for Reduced Compensation Relates the description to a possible reduction of compensation, e.g. in case of an accepted risk. Applicable only in rare cases.

Ground for Exclusion of Guilt Relates the description to a possible exclusion of guilt, e.g. when the actor is a child under 14 years of age.

Ground for Attribution Relates the description to the grounds on which the act is attributed to the liable person, e.g. on the grounds of guilt, on grounds of a specific law.¹⁰

Criterion for Strict Liability Relates the description to criteria that may hold in the case of strict liability, e.g. the volatility of animal behaviour.

Liberating Circumstance Relates the description to certain circumstances that may lift strict liability, e.g. when the damage was caused as a consequence of civil war.

Liable Person The person held liable on the basis of the description.

Although these legal thematic relations cannot be used to improve search results – the query to the Lucene index is a simple weighted boolean query – they do allow us to improve on the explanation functionality of the portal. Indeed, this may seem to come dangerously close to a more heavyweight knowledge-based approach. As we argued earlier, this knowledge is already available in the vocabularies which is used in a more structured way.

Given these thematic relations, we can provide a more specific definition of the example mapping in the introduction of this section:

```

ex:Animal_Map ≡ :actor value cv:animal □ :result value cv:damage □
                :action value cv:dangerous_action □ :recipient value cv:person
□ :liable_person value lv:animal_owner □
  :ground_for_attribution value lv:strict_liability_for_animals □
  :criterion_for_strict_liability value lv:volatility_of_animal_behaviour □
  :criterion_for_strict_liability value lv:damage_caused_by_animal

```

The mapping above will not match a case description where the child was the one who bit the dog. By adopting case frames and thematic relations into the framework of BestMap, we now have a powerful way to create more distinctive descriptions of cases without a significant increase in the complexity of modelling or the expressiveness of our language.¹¹ This does not only improve the ability of users to pose their queries in a fashion that will improve their chances of intelligible results, it also increases precision for adequately formulated queries by diminishing the chances that distinct types of unlawful act have equivalent definitions.

4. Implementation

As with all things, the proof of the pudding is in the eating. This section describes our findings in implementing the approach described in the previous section. We discuss several significant adaptations to the vocabularies and fingerprints used in BestPortal.

¹⁰This relation is currently also used to express a liability attribution on grounds of strict liability, although this is not a proper use of the technical term ‘attribution’.

¹¹The expressivity of our models is that of \mathcal{RO} , which is \mathcal{ALC} with nominals and complex role inclusions.

Fingerprints The introduction of context aware mappings [2,3] required us to reconsider the way in which vocabularies and fingerprints interact. Until [7], the assumption was that a commonsense description would translate to a *single concept*: every query to the corpus was based on the fingerprint of one concept. As is evident in the example of the previous section, context aware mappings are designed to result in descriptions using *multiple* concepts from the legal vocabulary. Fingerprints of several legal concepts are combined in a query to the corpus, which can have undesirable effects.

Since context aware mappings are themselves OWL classes, they can be positioned in a subsumption hierarchy. A mapping that characterises some case as describable by *lv:strict_liability* can subsume cases of more specific forms of strict liability, e.g. *lv:strict_liability_for_animals*. This is a very powerful mechanism that we use to modularise and structure the mappings. A case that is described by a specific form of strict liability is automatically also described as a case of strict liability. The portal will combine the fingerprints of both concepts in its query to the Lucene index. It therefore requires fingerprints to be more clear-cut lexical representations of the concepts.

Fingerprints for generic concepts such as *lv:strict_liability* were designed to be used as stand alone queries that are able to retrieve *any* and *all* subsumed cases. Fingerprints for these concepts therefore encode hierarchical information, and include search phrases for the more specific concepts. As a result, a query that combines fingerprints of generic and specific concepts becomes a query for all other concepts subsumed by the generic one. This is clearly undesirable. Combined queries lead to more possibilities for overlapping fingerprints. As argued in the previous section, every legal case description has a number of fixed ingredients, which in combination create a unique characterisation. The stand-alone fingerprints tended to accommodate this contextual co-occurrence and oftentimes contained search phrases of contextually related concepts. This results in a query bias towards search phrases occurring in the overlap of fingerprints in the query. We checked all fingerprints for potential undesirable overlap with other concepts. Hierarchical information was pruned away from the fingerprints of non-leaf concepts, moving search phrases to lower level concepts where necessary.

A final optimisation of the set of fingerprints was made possible by the migration from a Collexis search engine to the Apache Lucene platform.¹² Lucene supports searches for words that occur within a maximum word distance inside a document. This removes the need to have fingerprints for all relevant ways in which words may be combined. It is expected that this improves recall and reduces the number of fingerprints per concept.

Vocabulary Structure We had another careful look at the legal and commonsense vocabularies in light of the extension to case frames. The organisation along thematic roles was latent in the commonsense vocabulary; even though most roles were present as concepts, the vocabulary still contained various branches that were not unified under those categories. It also contained a fair number of lexical word-like concepts such as ‘he’ and ‘she’ that were left from an earlier knowledge elicitation exercise that assumed layman descriptions would be input as natural language texts. Lastly, we enriched the vocabulary with a number of additional skos:broader relations that allow categories to play multiple roles, e.g. animals can be both actor and object in an action.

¹²See <http://www.collexis.com>.

The legal vocabulary was a large collection of concepts *related* to unlawful acts, and required more attention. The case frames depend on structured descriptions of legal cases, but the vocabulary was not designed with this structure in mind. Secondly, certain branches of the vocabulary were underdeveloped and had little or no hierarchical structure. Lastly, the vocabulary was the result of a translation from an OWL ontology to a SKOS vocabulary and did not take advantage of the lightweight semantics – and thus flexibility – of the skos:broader relation as compared to owl:subClassOf. Letting go of the strict semantics of class subsumption allows for a concise hierarchical structure that is more aligned to information retrieval.

On the basis of these insights, we developed a new legal vocabulary that is more aligned with the thematic roles of our legal case frame. The vocabulary is based on an introductory text [6] that gives a general overview of the domain of unlawful act.¹³ The number of top-level categories in the vocabulary was reduced from ten to four (liable person, liability, unlawfulness and attribution). Each lower level in the hierarchy is a palette of concepts that are to be considered before the higher level can be said to be applicable. For instance, to establish whether liability holds, it first needs to be determined what *type* of liability holds (guilt, fault or strict liability). Since all fingerprints were attached to the *original* legal vocabulary, the two vocabularies needed to be aligned. We used a simple two-way lexical alignment to determine concept similarity and bootstrap the alignment procedure. Each combination of concept labels that is the same, or has a suitable weighted Levenshtein distance, is added to an intersection of the concept labels. We then apply a Jaccard index to determine whether the two concepts are likely to be the same ($J > 0.25$): $J = \frac{|W_{old} \cap W_{new}|}{|W_{old} \cup W_{new}|}$

Running this comparison resulted in 440 out of 657 possible concept alignments (a coverage of 67%). The final merged vocabulary contains 375 mappings, between 117 (out of 180) new concepts and 174 (out of 477) concepts of the old vocabulary. Although we did not perform a proper evaluation for these proposed mappings, they formed a significant bootstrap to the alignment process (given a total of over four thousand fingerprints). For each new concept without a mapping, initial fingerprints were created based on their label. The pruning described in the previous section is applied after this step.

5. BestPortal Revisited

The BestPortal was reimplemented using client-side JavaScript technology to make the user interface more responsive and allow more flexible interaction (see Figure 3). The layman case frame is presented to the left; for each thematic role in the frame, a combo box allows the selection of a concept from an indented list. The selected concepts are then positioned in the ‘Beschrijving gebeurtenis’ bar, the role of a concept is indicated between brackets. Once a selection is made, the system checks whether any mappings apply (by querying our Sesame/SwiftOWLIM repository),¹⁴ and displays the applicable legal concepts below the layman case description. Hovering over a question mark will give additional information about the concept. If one or more legal concepts apply, the system will use their fingerprints to produce a weighted boolean query which is sent

¹³See http://semweb.cs.vu.nl/best-portal/admin/tort_vocabulary_tree.php.

¹⁴OpenRDF Sesame is a triplestore, see <http://www.openrdf.org>. SwiftOWLIM is a lightweight OWL reasoner that can operate on top of Sesame, see <http://www.ontotext.com/owlim/>.

The screenshot displays the JustFind/BestPortal interface. On the left, there is a sidebar with various filters: 'Factoren', 'Wat werd er gedaan?', 'Wie of wat deed het?', 'Wat was er bij betrokken?', 'Tegen wie is het gedaan?', 'Speciale situatie?', 'Wat was het gevolg?', 'Waar gebeurde het?', and 'Wanneer of hoelang?'. The main content area shows search results for 'LJN B3703' and 'LJN BD6519'. Each result includes a date, court name, and a brief description. A map on the right shows search results on a geographical map of the Netherlands, with red icons indicating the location of the courts mentioned in the results.

Figure 3. The BestPortal interface

to the Solr Lucene index, and displays a synopsis of the results in the center frame of the page. Attributes, such as the type of court, the place of court, and the legal area are clickable links that operate as filters on the search results. For instance, selecting the “Rechtbank Groningen” will filter the results for all verdicts by the court in Groningen.

Every result is annotated with RDFa attributes and can be scraped for RDF metadata.¹⁵ It also contains explicit deeplinks to an RDF/XML representation of the metadata (the blue RDF icon), a linked data browser (Marbles)¹⁶, and a deeplink to the LJN index at the Dutch Council of the Judiciary.¹⁷ The timeline in the bottom right is a Simile widget, that displays all results according to publication date. Every Dutch court present in our court proceedings was linked to the DBPedia and Geonames linked data repositories, allowing us to present our results on a Google map (the dark red icons). We also scanned the texts of our corpus for Dutch place names taken from Geonames, and created links wherever a match was found. This allows us to position the places mentioned in court proceedings on the Google map. The entire repository is exposed via a browse-able SPARQL endpoint,¹⁸ integrates with other Linked Data resources, and is published in the CKAN Linked Data repository.¹⁹

¹⁵RDFa is a way to transparently incorporate RDF descriptions within HTML, see <http://www.w3.org/TR/rdfa-syntax/>.

¹⁶See <http://www5.wiwiw5.fu-berlin.de/marbles/>

¹⁷The LJN is the official identifier of every court proceeding.

¹⁸All URIs are dereferencable via <http://linkeddata.few.vu.nl/rechtspraak>

¹⁹See <http://ckan.net/package/rechtspraak>.

Table 1. Retrieval results for “Dog bites child” vs “Child bites dog”, top 20 hits.

	Rechtspraak.nl		w/o Case Frames		Case Frames	
	+	-	+	-	+	-
“Dog bites child”	1	3	1	19	14	6
“Child bites dog”	0	4	2	18	3	17

6. Evaluation and Discussion

A first small evaluation shows that using case frames has a significant effect on retrieval quality over the simple contextual mappings. Table 1 compares results for our example queries, “Dog bites child, resulting in damage” and “Child bites dog, resulting in damage”, achieved via a naive full text query on Rechtspraak.nl, and a query on the same corpus via simple mappings and via case frames.²⁰ Results were capped at a maximum of 20. As was expected, Rechtspraak.nl could only find a limited number of applicable cases (a total of four), and could not distinguish between the two queries. The one correct hit concerned a case where a woman was bitten by a dog. Using simple mappings improves matters somewhat – the total number of results rose to over 140 thousand – but still produces the same results for both queries, and generates high number of false positives. Case frames have a significant positive effect on the first query: false positives dropped to 30%, with no cases of animal molest. The top 6 results all concerned strict liability for animals. However, case frames do not perform as well for the second query, where most results concerned child custody cases²¹ and none of the results were about animals.

Although we are still in the process of improving and pruning our fingerprints (section 4), and the quality of search results can be much improved, even this very limited evaluation shows the potential for case frame based conceptual search as compared to the simple mappings. The difference with a traditional naive fulltext search on the same corpus is even more striking and shows that the current Rechtspraak.nl requires users to make an effort in coming up with the right formulation.

We are currently in the process of preparing a larger evaluation using two groups of students: law students and computer science students.²² They will use either Rechtspraak.nl or the BestPortal to solve a small number of cases. The performance of the systems will be evaluated on the basis of the students’ ability to correctly predict the judgement of a legal expert for each case.

This paper introduces case frames as a better way of structuring mappings between two distinct and incommensurable vocabularies that concern the same domain. Although the approach is deemed applicable to many domains, we have successfully applied case frame based mappings in a revamped version of BestPortal. We discussed how case frames improve retrieval results without requiring significant additional modelling effort. We indicated our expectation that it will be easier for users to formulate answerable queries if they use the structure provided by these frames, and shown some initial – but promising – evaluation results.

²⁰Our full text search string was “dog bites child damage”.

²¹These form the bulk of court proceedings that concern children.

²²Results of this evaluation are expected sometime during the fall of this year, and will be presented at the conference.

References

- [1] Tommaso Agnoloni, Lorenzo Bacci, Enrico Francesconi, Wim Peters, Simonetta Montemagni, and Giulia Venturi. A two-level knowledge approach to support multilingual legislative drafting. In Joost Breuker, Pompeu Casanovas, Michel Klein, and Enrico Francesconi, editors, *Law, Ontologies and the Semantic Web*, Frontiers of Artificial Intelligence and Applications. IOS Press, Amsterdam, 2009.
- [2] Rinke Hoekstra. BestMap: Context-aware SKOS vocabulary mappings in OWL 2. In Rinke Hoekstra and Peter Patel-Schneider, editors, *Proceedings of OWL: Experiences and Directions (OWLED 2009)*, 2009. Technical Paper.
- [3] Rinke Hoekstra. BestPortal: Lessons learned in lightweight semantic access to court proceedings. In Guido Governatori, editor, *Proceedings of the 22nd International Conference on Legal Knowledge and Information Systems (JURIX 2009)*, Frontiers of Artificial Intelligence, Rotterdam, December 2009. Jurix, IOS Press.
- [4] Rinke Hoekstra. *Ontology Representation – Design Patterns and Ontologies that Make Sense*, volume 197 of *Frontiers of Artificial Intelligence and Applications*. IOS Press, Amsterdam, June 2009.
- [5] Rinke Hoekstra, Joost Breuker, Marcello Di Bello, and Alexander Boer. LKIF Core: Principled ontology development for the legal domain. In Joost Breuker, Pompeu Casanovas, Michel Klein, and Enrico Francesconi, editors, *Law, Ontologies and the Semantic Web*, volume 188 of *Frontiers of Artificial Intelligence and Applications*. IOS Press, Amsterdam, 2009.
- [6] J. Spier, T. hartlief, G.E. van Maanen, and R.D. Vriesendorp. *Verbintenissen uit de Wet en Schadevergoeding*. Kluwer, Deventer, 1997.
- [7] Elisabeth M. Uijttendbroek, Arno R. Lodder, Michel C.A. Klein, Gwen R. Wildeboer, Wouter van Steenberg, Rory L.L. Sie, Paul E.M. Huygen, and Frank van Harmelen. Retrieval of case law to provide layman with information about liability: Preliminary results of the best-project. In *Computable Models of the Law*, volume 4884/2008 of *LNCS*. 2008.
- [8] Saskia van de Ven, Rinke Hoekstra, Joost Breuker, Lars Wortel, and Abdallah El-Ali. Judging Amy: Automated legal assessment using OWL 2. In *Proceedings of OWL: Experiences and Directions (OWLED 2008 EU)*, October 2008.
- [9] Ronny van Laarschot, Wouter van Steenberg, Heiner Stuckenschmidt, Arno R. Lodder, and Frank van Harmelen. The legal concepts and the layman's terms – bridging the gap through ontology-based reasoning about liability. In *Proceedings of the 18th Annual Conference on legal Knowledge and Information Systems (JURIX 2005)*, Brussels, December 8-9 2005. IOS Press.
- [10] Raboud Winkels and Henk de Bruijn. Case frames: Bridging the gap between a case and the law. In I. Carr and A. Narayanan, editors, *Proceedings of the First European Conference on Computers, Law and AI*, pages 205–213, Exeter, 1996. EUCLID.