

# Automatic Extraction of Definitions from German Court Decisions

Stephan Walter  
Department of Computational Linguistics  
Universität des Saarlandes  
66123 Saarbrücken, Germany  
stwa@coli.uni-saarland.de

Manfred Pinkal  
Department of Computational Linguistics  
Universität des Saarlandes  
66123 Saarbrücken, Germany  
pinkal@coli.uni-saarland.de

## Abstract

This paper deals with the use of computational linguistic analysis techniques for information access and ontology learning within the legal domain. We present a rule-based approach for extracting and analysing definitions from parsed text and evaluate it on a corpus of about 6000 German court decisions. The results are applied to improve the quality of a text based ontology learning method on this corpus.<sup>1</sup>

## 1 Motivation

Methods like ontology based knowledge management and information access through conceptual search have become active research topics in the general research community, with practical applications in many areas. However the use of IT in legal practice (at least in German speaking countries) is up to now mainly restricted to document preparation and management or Boolean keyword search on full-text collections. Legal ontologies have been proposed in various research projects, but they focus on an upper level of concepts and are, with only a few exceptions, small knowledge repositories that were hand-made by experts (for a summary of existing legal ontologies, cf. (Valente 2004)).

It is clear that realistically large knowledge-based applications in the legal domain will need more comprehensive ontologies incorporating e.g. up-to-date knowledge from court decisions. For this purpose an expert-based approach has to

be supplemented by automatic acquisition methods. The same is true for large-scale advanced information access: Extensive conceptual indexing of even a fraction of all court decisions published in one year seems hardly possible without automatic support. However there has been relatively little research on the use of natural language processing for this purpose (exceptions are (Lame 2005) and (Saias and Quaresma 2005)).

In this paper we look at the use of computational linguistic analysis techniques for information access and ontology learning within the legal domain. We present a rule-based method for extracting and analyzing definitions from parsed text, and evaluate this method on a corpus of about 6000 German court decisions within the field of environmental law. We then report on an experiment exploring the use of our extraction results to improve the quality of text-based ontology learning from noun-adjective bigrams. We will start however with a general discussion of the role that definitions play in legal language.

## 2 Definitions in Legal Language

Two central kinds of knowledge contained in the statutes of a code law system are *normative knowledge*, connecting legal consequences to descriptions of certain facts and situations, and *terminological knowledge*, consisting in definitions of some of the concepts used in these descriptions (Valente and Breuker 1994).

Normative content is exemplified by (1), parts of section 324a of the German criminal law. The legal consequence consisting in the specified punishment is connected to the precondition of soil pollution:

(1) *Whoever (...) allows to penetrate or releases substances into the soil and thereby pollutes it or otherwise detrimentally alters it:*

<sup>1</sup> This paper describes research within the project CORTE funded by the German Science Foundation, DFG PI 154/10-1 (<http://www.coli.uni-saarland.de/projects/corte/>)

1. in a manner that is capable of harming (...) property of significant value or a body of water (...)

shall be punished with imprisonment for not more than five years or a fine.

Terminological knowledge consists in definitions of concepts used to describe the sanctioned facts. E.g., *soil* is defined in article 2 of the German soil protection law as follows:

(2) *Soil within the meaning of this Act is the upper layer of the earth's crust (...) including its liquid components (soil solution) and gaseous components (soil air), except groundwater and beds of bodies of water.*

If the definitions contained in statutes would fully specify how the relevant concepts are to be applied, cases could be solved (once the relevant statutes have been identified) by mechanically checking which of some given concepts apply, and then deriving the appropriate legal consequences in a logical conclusion. However such a simple procedure is never possible in reality. Discussions in courts (and consequently in all legal texts that document court decisions) are in large parts devoted to pinning down whether certain concepts apply. Controversies often arise because not all relevant concepts are defined at all within statutes, and because the terms used in legal definitions are often in need of clarification themselves. For instance it may be unclear in some cases what exactly counts as the *bed of a body of water* mentioned in Example (2). Additionally, reality is complex and constantly changing, and these changes also pertain to the applicability of formerly clear-cut concepts. While this is especially true of social reality, rather physical concepts may also be affected. An often cited example is a case where the German *Reichsgericht* had to decide whether *electricity* was to be counted as a *thing*.

At the heart of these difficulties lies the fact that statutes are written in natural language, not in a formalized or a strongly restricted specialized language. It is widely assumed in the philosophical literature that most natural language concepts do not lend themselves to definitions fixing all potential conditions of applicability a priori. From the point of view of legal theory this open-textured character of natural language concepts is often seen as essential for the functioning of any legal system (the term open texture was

introduced into this discussion by (Hart 1961)). The use of natural language expressions allows for a continuous re-adjustment of the balance between precision and openness. This possibility is needed to provide regulations that are on the one hand reliable and on the other hand flexible enough to serve as a common ground for all kinds of social interaction. For the solution of concrete cases, the concepts made available within statute texts are supplemented by further definitions (in a wide sense, covering all kinds of modification and adaptation of concepts) given in the courts' decisions (in particular within the reasons for judgement). Such definitions for instance fix whether a certain stretch of sand counts as the *bed of a body of water* or if something is of *significant value* in the case at hand. These definitions are generally open for later amendment or revision. Still they almost always remain binding beyond the case at hand.

Easy access to definitions in decisions is therefore of great importance to the legal practitioner. Sections 3 and 4 show how computational linguistic analysis helps answering this need by enabling an accurate search for definitions in a large collection of court decisions. Accurate definition extraction is a prerequisite to building up an information system that allows for concept-centred access to the interpretational knowledge spread over tens of thousands of documents produced by courts every year.

Definitions are however not only of direct value as a source of information in legal practice. They also provide contexts that contain particularly much relevant terminology, and are therefore a good place to search for concepts to be integrated in a domain ontology. Given the importance and frequency of definitions in legal text, such an approach seems particularly promising for this domain. Section 5 describes how automatically extracted definitions improve the results of a standard ontology learning method.

### 3 Structure of Definitions

Our current work is based on a collection of more than 6000 verdicts in environmental law. As a starting point however we conducted a survey based on a random selection of 40 verdicts from various legal fields (none of them is in our present test set), which contained 130 definitions. Inspection of these definitions has shown a range of common structural elements, and has allowed us to identify typical linguistic realizations of

these structural elements. We will illustrate this with the example definition given in (3):

(3) [<sub>4</sub> Bei einem Einfamilienreihenhaus] [<sub>3</sub> liegt ein [<sub>1</sub> mangelhafter Schallschutz] [<sub>5</sub> dann] [<sub>3</sub> vor, wenn] [<sub>2</sub> die Haustrennwand einschalig errichtet wurde] (...).

(One-family row-houses have insufficient noise insulation if the separating wall is one-layered.)

This definition contains:

1. The *definiendum*, i.e. the element that is defined (*unzureichender Schallschutz - insufficient noise insulation*).

2. The *definiens*, i.e. the element that fixes the meaning to be given to the definiendum (*die Haustrennwand einschalig errichtet wurde - the separating wall is one-layered*).

Apart from these constitutive parts, it contains:

3. A *connector*, indicating the relation between definiendum and definiens (*liegt...vor, wenn, have..., if*).

4. A qualification specifying a *domain area* of applicability, i.e. a restriction in terms of the part of reality that the regulation refers to (*bei Einfamilienreihenhäusern - one-family row-houses*).

5. *Signal* words that cannot be assigned any clear function with regard to the content of the sentence, but serve to mark it as a definition (*dann - ø*).

The connector normally contains at least the predicate of the main clause, often together with further material (subjunction, relative pronoun, determiner). It not only indicates the presence of a definition. It also determines how definiens and definiendum are realized linguistically and often contains information about the type of the given definition (full, partial, by examples...). The linguistic realization of definiendum and definiens depends on the connector. One common pattern realizes the definiendum as the subject, and the definiens within a subclause. The domain area is often specified by a PP introduced by *bei* ("in the field of", *for*), as seen in the example. Further possibilities are other PPs or certain subclauses. Signal words are certain particles (*dann* in the example), adverbs (e.g. *begrifflich - conceptually*) or nominal constructions containing the definiendum (e.g. *der Begriff des..., the concept of...*).

Of course many definitions also contain further structural elements that are not present in

Example (3). For instance certain adverbials or modal verbs modify the force, validity or degree of commitment to a definition (e.g. only for typical cases). The *field of law* within which the given definition applies is often specified as a PP containing a formal reference to sections of statutes or simply the name of a statute, document, or even a complete legal field (e.g. *Umweltrecht - environmental law*). Citation information for definitions is standardly included in brackets as a reference to another verdict by date, court, and reference number.

## 4 Automatic extraction of definitions

The corpus based pilot study discussed in the last section has on the one hand shown a broad linguistic variation among definitions in reasons for judgement. No simple account, for instance in terms of keyword spotting or pattern matching, will suffice to extract the relevant information from a significant amount of occurrences. On the other hand our survey has shown a range of structural uniformities across these formulations. This section discusses computational linguistic analysis techniques that are useful to identify and segment definitions based on these uniformities.

### 4.1 Linguistic Analysis

Our current work is based on a collection of more than 6000 verdicts in environmental law that were parsed using the *Preds*-parser (*Preds* stands for *partially resolved dependency structure*), a semantically-oriented parsing system that has been developed in the Saarbrücken Computational Linguistics Department within the project COLLATE. It was used there for information extraction from newspaper text (Braun 2003, Fliedner 2004). The *Preds*-parser balances depth of linguistic analysis with robustness of the analysis process and is therefore able to provide relatively detailed linguistic information even for large amounts of syntactically complex text.

It generates a semantic representation for its input by a cascade of analysis components. Starting with a topological analysis of the input sentence, it continues by applying a phrase chunker and a named entity recognizer to the contents of the topological fields. The resulting extended topological structure is transformed to a semantic representation (called *Preds*, see above) by a series of heuristic rules. The *Preds*-format encodes semantic dependencies and modification relations within a sentence using abstract categories




Figure 1. Grammatical structure for Example (3).

such as deep subject and deep object. This way it provides a common normalized structure for various surface realizations of the same content (e.g. in active or passive voice).

The *Preds*-parser makes use of syntactic underspecification to deal with the problem of ambiguity. It systematically prefers low attachment in case of doubt and marks the affected parts of the result as default-based. Later processing steps are enabled to resolve ambiguities based on further information. But this is not necessary in general. Common parts of multiple readings can be accessed without having to enumerate and search through alternative representations. Figure 1 shows the parse for the definition in Example (3).<sup>2</sup> The parser returns an XML-tree that contains this structure together with the full linguistic information accumulated during the analysis process.

## 4.2 Search and processing

The structures produced by the *Preds* parser provide a level of abstraction that allows us to turn typical definition patterns into declarative extraction rules. Figure 2 shows one such extraction rule. It specifies (abbreviated) XPath-expressions describing definitions such as Example (3). The field *query* contains an expression characterising a sentence with the predicate *vorliegen* and a subclause that is introduced by the subjunction *wenn* (*if*). This expression is evaluated on the *Preds* of the sentences within our corpus to identify definitions. Other fields determine the locations containing the structural elements (such as *definiendum*, *definiens* and domain *area*) within the *Preds* of the identified definitions.

```
<pattern>
description=liegt vor + wenn-Nebensatz
query=sent/parse/preds/word[@stem="vorliegen"
and INDPRES and WENN]
filters=definite
definiendum=DSub
definiens=WENN/arg/word
area=PPMOD{PREP%bei}
</pattern>
```

Figure 2. Extraction rule.

The field *filters* specifies a set of XSLT-scripts used to filter out certain results. In the example we exclude definienda that are either pronominal (because we do not presently resolve anaphoric references) or definite (because these are often also anaphoric, or indicate that the sentence at hand is valid for that particular case only). Figure 3 shows how the definition in Example (3) is analyzed by this rule.

## 4.3 Evaluation

We currently use 33 such extraction rules based on the connectors identified in our pilot study, together with various kinds of filters. When applied to the reasons for judgement in all 6000 decisions (containing 237935 sentences) in our environmental law corpus, these rules yield 5461 hits before filtering (since not all patterns are mutually exclusive, these hits are all within 4716 sentences). After exclusion of pronominal and in some cases definite definienda (see above), as well as definienda containing stop-words (certain very common adjectives and nouns) the number of remaining hits decreases to 1486 (in 1342 sentences).

A selection of 492 hits (in 473 sentences; all hits for rules with less than 20, at least 20 hits for others) was checked for precision by two annotators. The evaluation was based on a very inclusive concept of definition, covering many cases of doubt such as negative applicability conditions, legal preconditions or elaborations on the use of evaluative terms. Clear “no”-judgements

<sup>2</sup> Figure 1 and Figure 3 were generated using the SALSA-Tool (Burchardt et al. 2006)




Figure 3. Structural elements of the definition in Example (3).

were e.g. given for statements referring only to one particular case without any general elements, and for purely contingent statements. The overall agreement of the judgements given was relatively high, with an overall  $\kappa$  of 0.835.

Total	
33 rules	1486 hits (1342 / 237935 sent)
Annotator 1	Good: 211/473 (p = 44.6 %)
Annotator 2	Good: 230/473 (p = 48.6 %)
Best rules only	
Annotator 1	
17 rules	749 hits (749 / 1342 sent)
	Good: 176/245 (p = 71.8 %)
Annotator 2	
18 rules	764 hits (633 / 1342 sent)
	Good: 173/230 (p = 75.2 %)

Table 1. Annotation results.

Precision values within the checked hits vary considerably. However in both cases more than 50 % of all hits are by patterns that together still reach a precision of well above 70 % (Table 1).

#### 4.4 Discussion

So far, our focus in selecting rules and filters has been on optimizing precision. As our present results show, it is possible to extract definitions at an interesting degree of precision and still achieve a reasonable number of hits. However we have not addressed the issue of recall systematically yet. The assessment of recall poses greater difficulties than the evaluation of the precision of search patterns. To our knowledge no reference corpus with annotated definitions ex-

ists. Building up such a corpus is time intensive, in particular because of the large amount of text that has to be examined for this purpose. Within the 3500 sentences of the 40 decisions examined in our pilot study mentioned above, we found only about 130 definitions. While this amount is significant from the perspective of information access, it is quite small from the annotator's point of view. Moreover it has become clear in our pilot study that there is a considerable amount of definitions that cannot be identified by purely linguistic features, and that many of these are unclear cases of particular difficulty for the annotator. The proportion of such problematic cases will obviously be much higher in free text annotation than in the evaluation of our extraction results, which were generated by looking for clear linguistic cues.

Taking the ratio observed in our pilot study (130 definitions in 3500 sentences) as an orientation, the set of rules we are currently using is clearly far from optimal in terms of recall. It seems that a lot of relatively simple improvements can be made in this respect. A variety of obvious good patterns are still missing in our working set. We are currently testing a bootstrapping approach based on a seed of various noun-combinations taken from extracted definitions in order to acquire further extraction patterns. We hope to be able to iterate this procedure in a process of mutual bootstrapping similar to that described in (Riloff and Jones 1999).

Moreover all presently employed rules use patterns that correspond to the connector-parts (cf. Section 3) of definitions. Accumulations of e.g. certain *signals* and *modifiers* may turn out to indicate definitions with equal precision. We

identified a range of adverbial modifiers that are highly associated with definitions in the corpus of our pilot study, but we have not yet evaluated the effect of integrating them in our extraction patterns.

We also assume that there is great potential for more fine-grained and linguistically sensitive filtering, such that comparable precision is achieved without losing so many results.

Even with all of the discussed improvements however, the problem of definitions without clear linguistic indicators will remain. Heuristics based on domain specific information, such as citation and document structure (e.g. the first sentence of a paragraph is often a definition), may be of additional help in extending recall of our method to such cases.

Apart from integrating further features in our extractors and using bootstrapping techniques for identifying new patterns, another option is to train classifiers for the identification of definitions based on parse features, such as dependency paths. This approach has for instance been used successfully for hypernym discovery (cf. Snow et al., 2005). For this task, WordNet could be used as a reference in the training and evaluation phase. The fact that no comparable reference resource is available in our case presents a great difficulty for the application of machine learning methods.

## 5 Ontology Extraction

Occurrence of a concept within a definition is likely to indicate that the concept is important for the text at hand. Moreover in court decisions, a great deal of the important (legal as well as subject domain) concepts will in fact have at least some occurrences within definitions. This can be assumed because legal argumentation (as discussed in Section 2) characteristically proceeds by adducing explicit definitions for all relevant concepts. Definition extraction therefore seems to be a promising step for identifying concepts, in particular within legal text. This section discusses how extracted definitions can be used to improve the quality of text-based ontology learning from court decisions. For this purpose we first examine the results of a standard method – identification of terms and potential class-subclass-relations through weighted bigrams – and then look at the effect of combining this method with a filter based on occurrence within definitions.

### 5.1 Bigram Extraction

Adjective-noun-bigrams are often taken as a starting point in text based ontology extraction because in many cases they contain two concepts and one relation (see e.g. Buitelaar et al. 2004). The nominal head represents one concept, while adjective and noun together represent another concept that is subordinate to the first one. There are however obvious limits to the applicability of this *concept-subconcept-rule*:

(1) It may happen that the bigram or even already the nominal head on its own do not correspond to relevant concepts, i.e. that one or both of the denoted classes are of no particular relevance for the domain.

(2) Not all adjective-noun-bigrams refer to a subclass of the class denoted by the head noun. Adjectives may e.g. be used redundantly, making explicit a part of the semantics of the head noun, or the combination may be non-compositional and therefore relatively unrelated to the class referred to by the head noun.

For these reasons, extracted bigrams generally need to be hand-checked before corresponding concepts can be integrated into an ontology. This time-intensive step can be facilitated by providing a relevance-ranking of the candidates to be inspected. Such rankings use association measures known from collocation discovery (like  $\chi^2$ , pointwise mutual information or log-likelihood-ratios). But while the elements of a collocation are normally associated in virtue of their meaning, they do not necessarily correspond to a domain concept just by this fact. Moreover, many collocations are non-compositional. An association based ranking therefore cannot solve Problem (2) just mentioned, and only partially solves Problem (1). However it seems likely that the definiendum in a definition is a domain concept, and for the reasons discussed in Section 2, it can be assumed that particularly many concepts will in fact occur within definitions in the legal domain. In order to investigate this hypothesis, we extracted all head-modifier pairs with nominal head and adjectival modifier from all parsed sentences in our corpus. We then restricted this list to only those bigrams occurring within at least one identified definiendum, and compared the proportion of domain concepts following the concept-subconcept-rule on both lists.

### 5.2 Unfiltered Extraction and Annotation

We found a total 165422 bigram-occurrences of 73319 types (in the following we use *bigrams*

to refer to types, not to occurrences) within the full corpus. From this list we deleted combinations with 53 very frequent adjectives that are mainly used to establish uniqueness for definite reference (such as *vorgenannt – mentioned above*). All types with more than 5 occurrences were then ranked by log-likelihood of observed compared to independent occurrence of the bigram elements.<sup>3</sup> The resulting list contains 4371 bigrams on 4320 ranks. Each bigram on the first 600 ranks of this list (601 bigrams, two bigrams share rank 529) was assigned one of the following five categories:

**1. Environmental domain:** Bigrams encoding concepts from the environmental domain (e.g. *unsorted construction waste*). These occur because our corpus deals with environmental law.

**2. Legal domain:** Bigrams encoding concepts from the legal domain. These range from concepts that are more or less characteristic of environmental law (e.g. various kinds of *town-planning schemes*) to very generic legal concepts (such as *statutory prerequisite*)

**3. No subconcept:** Bigrams that would be categorized as 1. or 2., but (typically for one of the reasons explained above) do not encode a subconcept of the concept associated with the head noun. An example is *öffentliche Hand* (“*public hand*”, i.e. public authorities – a non-compositional collocation).

**4. No concept:** All bigrams that - as a bigram - do not stand for a domain concept (although the nominal head alone may stand for a concept).

**5. Parser error:** Bigrams that were obviously misanalysed due to parser errors.

Figure 4 shows the distribution of categories among the 600 top-ranked bigrams, as well as within an additionally annotated 100 ranks towards the end of the list (ranks 3400-3500).




Figure 4. Results of log-likelihood ranking.

<sup>3</sup> The ranking was calculated by the Ngram Statistics Package described in (Bannerjee and Pedersen 2003)

For selecting the two categories of central interest, namely those of legal and environmental concepts to which the concept-subconcept rule applies, the ranking is most precise on the first few hundred ranks, and loses much of its effect on lower ranks. The percentage of such concepts decreases from 56% among the first 100 ranks to 51% among the first 200, but is roughly the same within the first 500 and 600 ranks (with even a slight increase, 45.6% compared to 46.8%). Even the segment from rank 3400 to 3500 still contains 39% of relevant terminology. There are no bigrams of the “no subconcept” category within this final segment. The explanation for this fact is probably that such bigrams (especially the non-compositional ones) are mostly established collocations and therefore show a particularly high degree of association.

It must be noted that the results of our annotation have to be interpreted cautiously. They have not yet been double-checked and during the annotation process there turned out to be a certain degree of uncertainty especially in the subclassification of the various categories of concepts (1, 2 and 3). A further category for concepts with generic attributes (e.g. *permissible*, combining with a whole range of one-word terms) would probably cover many cases of doubt. The binary distinction between concepts and non-concepts in contrast was less difficult to make, and it is surely safe to conclude about general tendencies based on our annotation.

### 5.3 Filtering and Combined Approach

By selecting only those bigrams that occur within *defienda*, the 4371 items on the original list were reduced to 227 (to allow for comparison, these were kept in the same order and annotated with their ranks as on the original list). Figure 5 shows how the various categories are distributed within the items selected from the top segments of the original list, as well as within the complete 227 filtering results.




Figure 5. Filtered results

The proportion of interesting concepts reaches about 80% and is higher than 60% on the complete selection. This is still well above the 56% precision within the top 100-segment of the original list. However the restriction to a total of 227 results on our filtered list (of which only 145 are useful) means a dramatic loss in recall. This problem can be alleviated by leaving a top segment of the original list in place (e.g. the top 200 or 500 ranks, where precision is still at a tolerably high level) and supplementing it with the lower ranks from the filtered list until the desired number of items is reached. Another option is to apply the filtering to the complete list of extracted bigrams, not only to those that occur more than 5 times. We assume that a concept that is explicitly defined is likely to be of particular relevance for the domain regardless of its frequency. Hence our definition-based filter should still work well on concept candidates that are too infrequent to be considered at all in a log-likelihood ranking, and allow us to include such candidates in our selection, too.

We investigated the effect of a combination of both methods just described. For this purpose, we first extracted all noun-adjective bigrams occurring within any of the identified definienda, regardless of their frequency within the corpus. After completing the annotation on the 627 resulting bigrams they were combined with various top segments of our original unfiltered list.

Figure 6 shows the distribution of the annotated categories among the 627 bigrams from definienda, as well as on two combined lists. *Cutoff 200/750* is the result of cutting the original list at rank 200 and filling up with the next 550 items from the filtered list. For *cutoff 500/1000* we cut the original list at rank 500 and filled up with the following 500 items from the filtered one. The distribution of categories among the original top 200 is repeated for comparison.




Figure 6. Log-likelihood and filtering combined.

Precision among the 627 filtering results is higher than among the original top 200 (almost

56% compared to 51%), and only slightly smaller even for the 1000 results in the *cutoff 500/1000* setting. Using definition extraction as an additional knowledge source, the top 1000 results retrieved are thus of a quality that can otherwise only be achieved for the top 200 results.

## 6 Conclusion

In this paper we argued that definitions are an important element of legal texts and in particular of court decisions. We provided a structural segmentation scheme for definitions and discussed a method of applying computational linguistic analysis techniques for their text-based extraction and automatic segmentation. We showed that a large number of definitions can in fact be extracted at high precision using this method, but we also pointed out that there is still much room for improvement in terms of recall, e.g. through the inclusion of further definition patterns.

Our future work in this area will focus on the integration of extraction results across documents (e.g. recognizing and collecting complementary definitions for the same concept) and on a user interface for structured access to this data. For this work we have access to a corpus of several million verdicts provided to us by the company *juris GmbH*, Saarbrücken. We also demonstrated how the identification of definitions can improve the results of text-driven ontology learning in the legal domain. When looking for noun-adjective bigrams encoding relevant concepts, it leads to a considerable increase in precision to restrict the search to definienda only. This method is more precise than selecting the top ranks of a log-likelihood ranking. Its great disadvantage is the very low total number of results, leading to poor recall. However by combining a log-likelihood ranking with definition-based concept extraction, recall can be improved while still achieving better precision than with a log-likelihood ranking alone. Moreover this combined method also retrieves concepts that are too infrequent to be included at all in a log-likelihood ranking.

There is however another, maybe even more relevant reason to look for definitions in ontology learning. Definitions in legal text often very explicitly and precisely determine all kinds of relational knowledge about the defined concept. For instance they specify explicit subordinations (as in the classical *definitio per genus et differen-*



*tiam*), introduce restrictions on roles inherited from a superconcept, determine the constitutive parts of the definiendum, or contain information about its causal relations to other concepts. As one focus of our future work we plan to investigate how such rich ontological knowledge can be extracted automatically.

Andre Valente and Jost Breuker. 1994. A functional ontology of law. *Towards a global expert system in law*. CEDAM Publishers, Padua, Italy

## References

- Satanjeev Banerjee and Ted Pedersen. 2003. The Design, Implementation, and Use of the Ngram Statistics Package. *CICLing 2003*: 370-381
- Christian Braun. 2003. Parsing German text for syntacto-semantic structures. In *Prospects and Advances in the Syntax/Semantics Interface*, Lorraine-Saarland Workshop Series, Nancy, France:99-102
- Paul Buitelaar, Daniel Olejnik, Michael Sintek. 2004. A Protégé Plug-In for Ontology Extraction from Text Based on Linguistic Analysis In: *Proceedings of the 1st European Semantic Web Symposium (ESWS)*, Heraklion, Greece
- Aljoscha Burchardt, Katrin Erk, Anette Frank, Andrea Kowalski and Sebastian Padó. 2006. SALTO -- A Versatile Multi-Level Annotation Tool. *Proceedings of LREC 2006*, Genoa, Italy.
- Gerhard Fliedner. 2004. Deriving FrameNet Representations: Towards Meaning-Oriented Question Answering. *Proceedings of the International Conference on Applications of Natural Language to Information Systems (NLDB)*. Salford, UK. LNCS 3136/2004. Springer. 64–75.
- Herbert L.A. Hart. 1961. *The concept of Law*. Oxford University Press, London, UK
- Guiraudé Lame. 2005. Using NLP Techniques to Identify Legal Ontology Components: Concepts and Relations, *Lecture Notes in Computer Science, Volume 3369*:169 – 184
- Ellen Riloff and Rosie Jones. 1999. Learning Dictionaries for Information Extraction Using Multi-level Bootstrapping, *Proceedings of AAAI-99*, 474 - 479
- José Saias and Paulo Quaresma. 2005. A Methodology to Create Legal Ontologies in a Logic Programming Information Retrieval System. *Lecture Notes in Computer Science, Volume 3369*:185 - 200
- Rion Snow, Daniel Jurafsky, and Andrew Y. Ng. 2005. Learning syntactic patterns for automatic hypernym discovery, *Proceedings of NIPS 2004*, Vancouver, Canada.
- Andre Valente. 2005. Types and Roles of Legal Ontologies. *Lecture Notes in Computer Science, Volume 3369*:65 - 76.