Recognizing Cited Facts and Principles in Legal Judgements

Olga Shulayeva and **Advaith Siddharthan** and **Adam Wyner** [1]

Abstract. In common law jurisdictions, legal professionals cite facts and legal principles from precedent cases to support their arguments before the court for their intended outcome in a current case. This practice stems from the doctrine of *stare decisis*, where cases that have similar facts should receive similar decisions with respect to the principles. It is essential for legal professionals to identify such facts and principles in precedent cases, though this is a highly time intensive task. In this paper, we present studies that demonstrate that human annotators can achieve reasonable agreement on which sentences in legal judgements contain cited facts and principles (respectively, $\kappa = 0.65$ and $\kappa = 0.95$ for inter- and intra-annotator agreement). We further demonstrate that it is feasible to automatically annotate sentences containing such legal facts and principles in a supervised machine learning framework, reporting per category precision and recall figures of between 79% and 89% for classifying sentences in legal judgements as cited facts, principles or neither using a Bayesian classifier, with an overall κ of 0.72 with the human-annotated gold standard.

1 Introduction

In common law jurisdictions, legal practitioners treat existing case decisions (precedents) as a source of law. Case citations, references to legal precedents, are an important argumentation tool, enabling lawyers to formulate and present their argument persuasively. This practice stems from the doctrine of *stare decisis*, which can be translated from Latin as to 'stand by the decided cases' [2], where a case under consideration that has facts similar enough to precedent cases should receive similar decisions as the precedents. A legal professional looks to establish the relevant law in the current case; to do so, she must consult precedent cases in order to establish how similar patterns of facts were decided. Citations from existing case law are used to illustrate legal principles and facts that define the conditions for application of legal principles in the current case.

Citation analysis can help legal practitioners to identify which principles have applied in a certain case and which facts have been selected as the 'material' facts of the case, i.e. the facts that influenced the decision and which are crucial in establishing the similarity between two cases. There is no defined guide on how to identify the law embedded within common law decisions, so legal professionals are expected to make themselves familiar with as many relevant decisions as possible in order to make informed predictions about the outcome of a current case. Decisions delivered by courts are binding and can therefore provide useful information for legal professionals. The information that is embedded within the cited cases includes the legal principles and facts that are used to reason to a decision. Optimally, a legal professional finds a cited case with the same facts and legal principles, and so can argue that the decision for the current case should be that of the precedent; similarly, the opposing party may identify precedents with opposing principles to argue the decision should be otherwise. More commonly, legal professionals must consider a range of precedents, each of which highlight particular facts and legal principles that support their argument (or argue against the opposition). It is, then, essential that each side in the legal dispute identifies a relevant case base which supports the legal claims made during legal arguments. As the body of common law is continually growing, human citation analysis is complex as well as knowledge and time intensive.

To support citation analysis (discussed further in Section 2.1), existing electronic tools, such as electronic databases[3], provide one word summaries for relationships between cases (e.g.'applied'). However, it is uncommon for them to extract information about the facts and the legal principles of the cited cases. This means that on many occasions readers are required to make themselves familiar with the full text of multiple law reports in order to identify the applicable law and the correct way to apply it. Thus, citation analysis tools save some labour by providing a preliminary filter on relevant cases, yet, identification of particular cases and the essential details require further manual effort.

In the course of working on citation analysis, certain key concepts of legal theory must be scoped, given that this is a report on the computational analysis of the language of the law rather than on legal theory. In particular, cases are considered to contain *ratio decidendi*, which can be translated as a *reason for a decision*, an important piece of reasoning that is incorporated into the argumentation structure of future decisions. A variety of approaches to defining *ratio decidendi* can be identified in legal theory. As defined by [28]: '*ratio decidendi* can be identified as those statements of law which are based on the facts as found and upon which the decision is based'. [11] provides several explanations on what forms the binding part of a decision:

> '(1) the rule(s) of law that the court explicitly states, or that can reasonably be inferred, that it regarded as necessary to (or important in) its resolution of the case [. . .], (2) facts the precedent court regarded as 'material,' i.e., crucial for the court's resolution, plus the result of the case; and (3) facts the court now constrained by the precedent regards as material in the earlier case plus its result.'

The complexities stemming from the debates surrounding the defi-

[1] Department of Computing Science, University of Aberdeen, United Kingdom, email: {olga.shulayeva,advaith,azwyner}@abdn.ac.uk

[2] Source: http://thelawdictionary.org/

[3] e.g. LexisNexis Shepard's Citations Service `http://www.lexisnexis.com/en-us/products/shepards.page`

nition of ratio are excluded from the scope of this paper. Here, *ratio* will be understood as a combination of the facts of the current case along with the legal principles that are invoked when the facts of the current case are similar enough to the facts of the case that established the precedent.

This paper makes a novel, preliminary contribution towards automated identification of legal principles and facts embedded within common law citations. A gold standard corpus is created, with sentences containing legal principles and facts manually annotated. A Bayesian Multinomial Classifier (using Weka) is then applied to the corpus using a set of linguistic features to automatically identify these sentences. The main results are a demonstration that (a) the human annotation task is feasible, i.e. human annotators can achieve reasonable agreement on which sentences in legal judgements contain cited facts and principles and (b) it is feasible to automatically annotate sentences containing such legal facts and principles to a high standard. The reported studies lay the basis for further applications, including creation of meta-data for search and retrieval purposes, compilation of automated case treatment tables containing summaries about legal principles and material facts of cases, and automated analysis of reasoning patterns and consistency applied in legal argumentation.

We first present related work in Section 2. Then there are two studies, on manual annotation in Section 3 and on automated annotation in Section 4. The paper closes with some conclusions in Section 5.

2 Related work

This research aims to apply machine learning methodology in order to automatically identify legal principles and facts in case citations. A significant amount of work has been done in the area of citation analysis in scientific literature, while only a very small amount of work has been done that focuses on studying case law citations. Most existing studies on case law citations aim to identify case treatment – the relationship between citing and cited cases (e.g. *distinguished*, *explained*, and others) – or analyse citations from the point of view of network analysis, but don't focus on fine-grained analysis of the cited information. To the best of our knowledge, there is no reported work that specifically aims to apply machine learning methodology to identify legal principles and facts of the cited cases in case citations. In the following subsections, we discuss related work on citation analysis along with relevant literature on legal argumentation.

2.1 Citation analysis

The first attempts to systematise citation information were done in the field of common law by the developers of legal citators, starting with Frank Shepard in 1873, who relied on human expertise to provide discourse-aware summaries of case law citations. More recently, citation information is presented as in LexisNexis Shepard's Citations Service.

Despite lawyers being the pioneers of citation analysis [26], the research on citation analysis in common law has not been developing as fast as citation analysis in the domain of scientific reports. Eugene Garfield is often cited as one of the pioneers and key contributors towards citation analysis in science. Garfield was inspired by the Shepard's citations and argued that similar methodologies can be useful for summarisation of scientific citations [10]. Garfield employed a bibliographic approach to create ICI Citation Indexes, and the data from citation indexes was later used for a number of bibliometric studies that "extract, aggregate and analyse quantitative aspects of bibliographic information" [22]. He believed that citation analysis could be used for evaluation of scientific performance, for example, in calculation of journal ranks based on citation frequency and impact. As noted by [22], quantitative data from bibliometric studies is widely used to assess the performance of individual scholars, scientific journals, research institutions and 'general, structural aspects of the scholarly system' (e.g. measuring trends in national publication output). [22] also concluded that ICI citation indexes do not 'capture motives of individuals, but their consequences at an aggregate level' and argued for further development of qualitative citation based indicators, thus abandoning the principle underlying most citation analyses that 'all citations are equal'. Qualitative approaches in citation analysis take into account the intentions of the person who was providing the citation. They aim to capture citation qualities that are overlooked by quantitative methodologies, for example, such as polarity and sentiment. A scientific article may be frequently cited, but it can be due to criticisms or mere acknowledgements, which distinguishes it from an article introducing an approach that is widely accepted and utilised. Several researchers can be mentioned in respect of qualitative citation based indicators in science [24, 29, 4, 32, 2]. [6] conducted a research of citation behaviours and noted that at the time there was not a universal approach in citation studies. Application of qualitative citation based indicators often relies on linguistic discourse markers to generate conclusions about citations and citing behaviours. For example, citations can be classified according to sentiment polarities: confirmative or negative [24]; positive, neutral or weak [32].

Recently there has been more interest toward citation studies in law, where there appear to be two major directions: applying network analysis to citations [37, 19, 34, 20, 33, 25] and classification systems allowing one to estimate the 'treatment' status of the cited case [16, 9].

[37] developed Semantics-Based Legal Citation Network, a tool that extracts and summarises citation information, allowing the users to 'easily navigate in the citation networks and study how citations are interrelated and how legal issues have evolved in the past.' The researchers note that different parts of a case can be cited and studying the reasons for citation can provide valuable information for a legal researcher. Their approach relied on RFC (reason for citing), a patented technology that allows extracting reasons of why the case has been cited. RFC performance was summarised in the patent [15], and it explored a methodology of 'identifying sentences near a document citation (such as a court case citation) that suggest the reason(s) for citing (RFC)'. In [37], the information retrieved by RFC was further organised into semantic citation networks. The task of identifying RFC may be somewhat similar to the task that is undertaken as a part of this project due to the fact that information contained in principles and facts of cited cases can be used as a part of estimating reasons for citing.

History Assistant was designed by [16] to automatically infer direct and indirect treatment history from case reports. Direct treatment history covered historically related cases, such as appeals etc. Indirect treatment history dealt with the cited cases within a document in order to establish how the cited case has been treated. It relied on the classification methodology of Shepard's citations that combines the knowledge about sentiment and aims of legal communication with heuristic information about court hierarchy. It includes such classes as applied, overruled and distinguished. History Assistant was expected to be an aid for editorial work rather than replace the effort of the editors. The program consisted of a set of natural language modules and a prior case retrieval module. Natural language process-

ing relied on machine learning methodology and employed statistical methods over annotated corpus.

[9] created LEXA – a system that relied on RDR (Ripple Down Rules) approach to identify citations within the 'distinguished' class. This category is generally best linguistically signalled and is therefore suitable for achieving high precision and recall. The key idea underpinning RDR was that the 'domain expert monitors the system and whenever it performs incorrectly he signals the error and provides as a correction a rule based on the case which generated the error, which is added to the knowledge base' [9]. The approach employed annotators to create an initial set of rules leaving the end users to refine and further expand the set. The authors claimed that 'the user can at any stage create new annotations and use them in creating rules' which may put a more significant reliance on the user input than an end user may be equipped or expecting to provide. LEXA employed 78 rules that recognized 'distinguished' citations with a precision of 70% and recall of 48.6% on the cleaned test set, which is significantly lower than the results reported by [16] for the same category: precision (94%) and recall (90%). The difference in results suggests that a complex fine-grained analysis used by [16] that included machine-learning for language processing may help achieve better classification outcomes.

2.2 Argument extraction

There have been a variety of attempts aimed at automated extraction of argumentation structure of text and its constituents. The methodologies employed by such studies often rely on extraction and further analysis of linguistic information that is available within the text. One of the relatively recent successful examples of argumentation extraction methodology can be argumentation zoning. This approach is based on the assumption that the argumentation structure can be presented as a combination of rhetorical zones that are used to group the statements according to their rhetorical role. This approach was initially used by [30, 31] for scientific reports. [12] used argumentation zoning to create summaries for common law reports. Both studies report acceptable results for most of the categories, with some categories performing better than others.

An approach similar to argumentation zoning was taken by [8] to develop a scheme for identification of argument structure of Canadian case law and [18] to analyse the structure of German court decisions. A methodology relying on manual annotation of discourse structures and in that respect similar to argumentation zoning was used by [36] to detect case elements such as Case citation, cases cited, precedential relationships, Names of parties, judges, attorneys, court sort, Roles of parties (i.e. plaintiff or defendant), attorneys, and final decision. Whilst the methodology developed does not aim to fully reconstruct argumentation structure, the information obtained during the study can be used as a part of a wider application.

[35] conducted a study aimed at identification of argumentation parts with the use of context-free grammars. Similar to [16] the study reports the following difficulties with identifying argumentation structures in legal texts: '(a) the detection of intermediate conclusions, especially the ones without rhetorical markers, as more than 20% of the conclusions are classified as premises of a higher layer conclusion; (b) the ambiguity between argument structures.' The results reported are as follows: premises – 59% precision, 70% recall; conclusions – 61% precision, 75% recall; non-argumentative information – 89% precision, 80% recall.

The methodology of applying statistical tools over annotated corpus was employed by [23] to automatically detect sentences that are a part of the legal argument. The study achieved 68% accuracy for legal texts. [1] aimed to extract 'argumentation-relevant information automatically from a corpus of legal decision documents' and 'build new arguments using that information'.

A related, important distinction that should be made with regard to legal argumentation is the idea that the cited legal principles can be classed as ratio or obiter. As defined by [28]: 'ratio decidendi can be understood as those statements of law which are based on the facts as found and upon which the decision is based'. Statements that are usually included into obiter class are dissenting statements and statements that are 'based upon either nonexistent or immaterial facts of the case' [28]. From the point of view of law the main difference between ratio and obiter is that the former is binding, while the latter only possesses persuasive powers. [3] tried to automatically identify and extract ratio. [27] tried to identify obiter statements. However, the distinctions between ratio or obiter will not be used as a part of this work.

3 Manual annotation study

The manual annotation study focused on annotating the gold standard corpus and evaluating the annotation methodology. This gold standard corpus was used to extract the features necessary for the machine annotation study. Two annotators were used for the purposes of the manual annotation study: Annotator 1 and Annotator 2. Annotator 1 has legal training and Annotator 2 does not. All manual annotation was performed in GATE[4].

3.1 Method

The corpus for the gold standard was compiled from 50 common law reports that had been taken from the British and Irish Legal Institute (BAILII) website in RTF format. The length and structure of reports varied, which was most often defined by the complexity of the matter: longer and more complicated cases often had more sections. As reported by GATE Sentence Splitter (GATE 8.0.), the full corpus contained 1211012 tokens (or words) and 22617 sentences which included headings and other units that didn't form full sentences from grammatical point of view. Most reports had a section on the top introducing the court, the parties, legal representatives, case number etc. It was often the case that the legal situation was presented in the introduction and that the legal analysis was in the middle of the report. However, the reports did not follow a universal format. Conclusions were often short and situated at the end of the report. Case law citations are used to support legal argumentation and are therefore referred to as a part of legal analysis. For that reason they were rarely found in introduction or conclusion.

Annotator 1 created annotation guidelines (high level task definition, descriptions and examples for each category, and analyses of a few difficult cases) in several iterations and trained Annotator 2 on their use. The annotators were expected to identify sentences that contained the legal *principles* and *facts* of the cited cases, based on the written guidelines. Sentences associated with cited cases that are neither *principles* or *facts* are annotated as *neutral*.

The task of annotation focused on the identification of cited information within annotation areas that were defined as paragraphs having at least one citation. Citation instances had been manually annotated prior to the study. Given the discussion of the complexity

[4] GATE 8.0: https://gate.ac.uk

of jurisprudential views of legal principles, we have taken an operationalised view, based on the analysis of a legal scholar and key linguistic indicators.

All propositions that are associated with the cited case should annotated if the court deems they support the legal reasoning of the citing case. A *legal principle* is a statement which is used, along with facts, to reach a conclusion. Linguistically, a legal principle can for instance be indicated by deontic modality, e.g. expressions of *must* for obligation, *must not* for prohibition, or *may* for permission, which contrast with epistemic modalities for necessity and possibility. For example:

> As a matter of principle no order should be made in civil or family proceedings without notice to the other side unless there is a very good reason for departing from the general rule that notice must be given. (Gorbunova v Berezovsky (aka Platon Elenin) & Ors, 2013)

Legal principles can be qualified, e.g. with conditions that may limit the application of rule. It is also possible that legal principles are "active" in reasoning, yet inferred from the text, in which case, they cannot be annotated or used for further text processing.

In contrast to legal principles, there are *facts*, which are statements bearing on what uncontroversially exists, occurred, or is a piece of information. For our purposes, only sentences that refer to events which occur outside the court hearing are annotated; this excludes procedural facts. For example:

> Miss Lange was not a party to the 1965 Transfer or the 1968 Deed and she covenanted only with Mrs de Froberville (and not with Brigadier Radford) to comply with the covenants in those instruments in so far as they were still subsisting and capable of taking effect. (89 Holland Park (Management) Ltd & Ors v Hicks, 2013)

Linguistically, facts present themselves with non-modal expressions and denoting expressions, e.g. are not generic, non-actual, and indefinite.

Following a period of training, a set of 10 reports were randomly selected (all previously unseen by the annotators) for the inter-annotator and intra-annotation agreement studies reported here. The process in short was to:

1. Use the pre-annotated citation instances to identify annotation areas – i.e. paragraphs that contain at least one citation name. Direct quotes and lists were treated as a part of the same paragraph.
2. Label each sentence in each annotation area as one of *fact*, *principle* or *neither*, following the annotation guidelines.

3.2 Results

Table 1 shows the distribution of categories in the evaluation set of 10 reports. It shows that Annotator 2, who does not have legal training, is more conservative in identifying facts and inferences than Annotator 1, who has had legal training.

The results of the inter-annotator agreement study are as follows: κ=0.65[5] (% Agreement=83.7). The intra-annotator agreement study showed that Annotator 1 (when annotating the same set of 10 reports three months apart in time) was extremely consistent: κ=0.95 (% Agreement=97.3).

[5] κ, the predominant agreement measure used in natural language processing research [5], corrects raw agreement $P(A)$ for agreement by chance $P(E)$: $\kappa = \frac{P(A)-P(E)}{1-P(E)}$

Annotator 1 proceeded to create a gold corpus of 50 reports which was used for training a machine classifier, as described next.

4 Automated annotation study

The methodology used for machine annotation employed classification of the annotation units with a Naive Bayesian Multinomial Classifier based on a set of selected features described below.

4.1 Method

The task of features selection focused on identifying the features that can help in classifying sentences. The following features were selected for extraction from the dataset:

- Unigrams
- Dependency pairs
- Length of the sentence
- Position in the text
- Part of speech tags
- Insent – a feature which indicates whether there is a citation instance in the sentence.
- Inpara – a feature which indicates sentences that were placed within annotation areas, so that sentences that were placed outside it could be filtered out.

Unigrams are widely used in text classification tasks. The performance of classifiers relying on bag-of-words approach can however be impeded by the assumption that word order and grammatical relations are not significant. To address the limitations researchers often complement unigrams by features that can capture dependencies between words. Dependency pairs derived using the Stanford Parser [7] were used to complement unigrams, creating word pairs that are grammatically linked rather than simply collocated like n-grams. Dependency features have previously been shown to be difficult to beat for a variety of text classifications tasks such as sentiment analysis [17] and stance classification [14, 21].

Part of speech tags were selected as a feature for a number of reasons. Firstly, it was expected that modal verbs and verb tense may help to classify the annotation units. Sentences that introduce facts are most often presented in the Past Indefinite tense. For example:

> The contract contained a general condition that in relation to any financial or other conditions either party could at any time before the condition was fulfilled or waived avoid the contract by giving notice.

Secondly, both epistemic and deontic modal qualifiers that use modal verbs are common in sentences containing legal principles, for example:

> It is a question which must depend on the circumstances of each case, and mainly on two circumstances, as indicating the intention, viz., the degree of annexation and the object of the annexation.

4.2 Results

Tables 2–3 report the classification performance of the Naive Bayes Multinomial classifier from the Weka toolkit [13]. The accuracy of the classifier is similar to that of the Annotator 2, who had no legal

	Annotator 1 (original annotation)	Annotator 2 (inter-annotator study)	Annotator 1 (intra-annotator study)
Principles	266 (32%)	211 (26%)	258 (31%)
Facts	56 (7%)	20 (2%)	54 (7%)
Neither	499 (61%)	590 (72%)	509 (62%)

Table 1. Distribution of categories

training in the manual study. This suggests that to the extent such annotations can be carried out based on linguistic principles alone, automated annotation can be performed to the same standard as manual annotation.

	Precision	Recall	F-Measure
Principles	0.823	0.797	0.810
Facts	0.822	0.815	0.818
Neither	0.877	0.892	0.884

Number of Sentences	2659
Accuracy	0.85
κ	0.72

Table 2. Per category and aggregated statistics for automatic classifier

Machine/Human:	Principles	Facts	Neither
Principles	646	5	160
Facts	4	198	41
Neither	135	38	1432

Table 3. Confusion Matrix

5 Conclusions

An overall analysis suggests that the machine annotation experiment has returned good classification results with Naive Bayesian Multinomial classifier identifying 85% of instances correctly and achieving Kappa equal 0.72. Good combinations of precision and recall have been achieved for all categories (rounding): 82% precision and 80% recall (principles), 82% precision and 81% recall (facts), and 87% precision and 89% recall (neither). Such positive results suggest that the methodology employed as a part of this experiment can provide a suitable basis for further work.

This is a preliminary work on automatic identification of legal principles and facts that are associated with a case citation. To productively deploy a system, further development of a larger and more complex corpus would need to be done. Furthermore, tools to facilitate web-based access to the annotated statements would have to be designed. Such tools would, for example, allow a legal practitioner to not only search, say in Google, for citations mentioned in a case, but also the associated legal principles and facts, providing deep access to and insight into the development of the law. It would also offer the opportunity to access the law directly rather than via the edited and structured materials made available by legal service providers. Finally, we have only addressed accessing cited legal principles and facts, which is distinct from ranking and relating precedents, i.e. Shepardisation. The approach developed here offers some of the source material that could then be used to automate Shepardisation as well as to evaluate given citation analyses.

REFERENCES

[1] Kevin D Ashley and Vern R Walker, 'Toward constructing evidence-based legal arguments using legal decision documents and machine learning', in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pp. 176–180. ACM, (2013).

[2] Awais Athar and Simone Teufel, 'Context-enhanced citation sentiment detection', in *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 597–601. Association for Computational Linguistics, (2012).

[3] K Branting, 'Four challenges for a computational model of legal precedent', *THINK (Journal of the Institute for Language Technology and Artificial Intelligence)*, **3**, 62–69, (1994).

[4] Virginia Cano, 'Citation behavior: Classification, utility, and location', *Journal of the American Society for Information Science*, **40**(4), 284, (1989).

[5] Jean Carletta, 'Assessing agreement on classification tasks: The kappa statistic', *Computational Linguistics*, **22**(2), 249–254, (1996).

[6] Blaise Cronin, 'Norms and functions in citation: The view of journal editors and referees in psychology', *Social Science Information Studies*, **2**(2), 65–77, (1982).

[7] Marie-Catherine De Marneffe, Bill MacCartney, Christopher D Manning, et al., 'Generating typed dependency parses from phrase structure parses', in *Proceedings of LREC*, volume 6, pp. 449–454, (2006).

[8] Atefeh Farzindar and Guy Lapalme, 'Legal text summarization by exploration of the thematic structures and argumentative roles', in *Text Summarization Branches Out Workshop held in conjunction with ACL*, pp. 27–34, (2004).

[9] Filippo Galgani, Paul Compton, and Achim Hoffmann, 'Lexa: Building knowledge bases for automatic legal citation classification', *Expert Systems with Applications*, **42**(17), 6391–6407, (2015).

[10] Eugene Garfield, 'Citation indexes for science', *Science*, **122**, 108–111, (1955).

[11] Kent Greenawalt, 'Interpretation and judgment', *Yale Journal of Law & the Humanities*, **9**(2), 5, (2013).

[12] Ben Hachey and Claire Grover, 'Extractive summarisation of legal texts', *Artificial Intelligence and Law*, **14**(4), 305–345, (2006).

[13] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten, 'The weka data mining software: an update', *ACM SIGKDD explorations newsletter*, **11**(1), 10–18, (2009).

[14] Kazi Saidul Hasan and Vincent Ng, 'Why are you taking this stance? identifying and classifying reasons in ideological debates', in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 751–762, Doha, Qatar, (October 2014). Association for Computational Linguistics.

[15] Timothy L Humphrey, Xin Allan Lu, Afsar Parhizgar, Salahuddin Ahmed, James S Wiltshire Jr, John T Morelock, Joseph P Harmon, Spiro G Collias, and Paul Zhang. Automated system and method for generating reasons that a court case is cited, February 15 2005. US Patent 6,856,988.

[16] Peter Jackson, Khalid Al-Kofahi, Alex Tyrrell, and Arun Vachher, 'Information extraction from case law and retrieval of prior cases', *Artificial Intelligence*, **150**(1), 239–290, (2003).

[17] Mahesh Joshi and Carolyn Penstein-Rosé, 'Generalizing dependency features for opinion mining', in *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pp. 313–316. Association for Computational Linguistics, (2009).

[18] Florian Kuhn, 'A description language for content zones of german court decisions', in *Proceedings of the LREC 2010 Workshop on the Semantic Processing of Legal Texts*, pp. 1–7, (2010).

[19] Elizabeth A Leicht, Gavin Clarkson, Kerby Shedden, and Mark EJ Newman, 'Large-scale structure of time evolving citation networks', *The European Physical Journal B*, **59**(1), 75–83, (2007).

[20] Yonatan Lupu and Erik Voeten, 'Precedent in international courts: A network analysis of case citations by the european court of human rights', *British Journal of Political Science*, **42**(02), 413–439, (2012).

[21] Angrosh Mandya, Advaith Siddharthan, and Adam Wyner, 'Scrutable feature sets for stance classification', in *Proceedings of the 3rd Work-*

shop on Argument Mining, ACL 2016, Berlin, Germany, (2016). Association for Computational Linguistics.
[22] Henk F Moed, 'Citation analysis of scientific journals and journal impact measures', *Current Science*, **89**(12), (2005).
[23] Marie-Francine Moens, Erik Boiy, Raquel Mochales Palau, and Chris Reed, 'Automatic detection of arguments in legal texts', in *Proceedings of the 11th international conference on Artificial intelligence and law*, pp. 225–230. ACM, (2007).
[24] Michael J. Moravcsik and Poovanalingan Murugesan, 'Some results on the function and quality of citations', **5**, 88–91, (1975).
[25] Thom Neale, 'Citation analysis of canadian case law', *J. Open Access L.*, **1**, 1, (2013).
[26] Patti Ogden, 'Mastering the lawless science of our law: A story of legal citation indexes', *Law Libr. J.*, **85**, 1, (1993).
[27] José Plug, 'Indicators of obiter dicta. a pragma-dialectical analysis of textual clues for the reconstruction of legal argumentation', *Artificial Intelligence and Law*, **8**(2-3), 189–203, (2000).
[28] Martin Raz, 'Inside precedents: The ratio decidendi and the obiter dicta', *Common L. Rev.*, **3**, 21, (2002).
[29] John Swales, 'Citation analysis and discourse analysis', *Applied linguistics*, **7**(1), 39–56, (1986).
[30] Simone Teufel, Jean Carletta, and Marc Moens, 'An annotation scheme for discourse-level argumentation in research articles', in *Proceedings of the ninth conference on European chapter of the Association for Computational Linguistics*, pp. 110–117. Association for Computational Linguistics, (1999).
[31] Simone Teufel, Advaith Siddharthan, and Colin Batchelor, 'Towards discipline-independent argumentative zoning: evidence from chemistry and computational linguistics', in *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 3-Volume 3*, pp. 1493–1502. Association for Computational Linguistics, (2009).
[32] Simone Teufel, Advaith Siddharthan, and Dan Tidhar, 'Automatic classification of citation function', in *Proceedings of the 2006 conference on empirical methods in natural language processing*, pp. 103–110. Association for Computational Linguistics, (2006).
[33] Marc van Opijnen, 'Citation analysis and beyond: in search of indicators measuring case law importance.', in *JURIX*, volume 250, pp. 95–104, (2012).
[34] Radboud Winkels, Jelle De Ruyter, and Henryk Kroese, 'Determining authority of dutch case law', *Legal Knowledge and Information Systems*, **235**, 103–112, (2011).
[35] Adam Wyner, Raquel Mochales-Palau, Marie-Francine Moens, and David Milward. Approaches to text mining arguments from legal cases, semantic processing of legal texts: where the language of law meets the law of language, 2010.
[36] Adam Z Wyner, 'Towards annotating and extracting textual legal case elements', *Informatica e Diritto: special issue on legal ontologies and artificial intelligent techniques*, **19**(1-2), 9–18, (2010).
[37] Paul Zhang and Lavanya Koppaka, 'Semantics-based legal citation network', in *Proceedings of the 11th international conference on Artificial intelligence and law*, pp. 123–130. ACM, (2007).