Citation Analysis and Beyond: in Search of Indicators Measuring Case Law Importance

Marc VAN OPIJNEN[1]
Council for the Judiciary, the Netherlands

Abstract. Could legal importance of a judicial decision be established by analyzing its position in a case law citation network? A network of nearly half a million citations was used to test a great variety of social network algorithms against three external benchmarks. All these algorithms were outperformed by a tailored variant of degree centrality. Furthermore, regression analysis was used to explore the possible relevance of eleven variables to the legal importance of case law.

Keywords. Citation networks. Legal authority. Case law databases.

Introduction

Although the increasing number of judicial decisions published on the internet is a welcome development for lawyers, scholars and citizens, available search tools can hardly keep step with this exponential growth. As a result, finding relevant decisions is becoming increasingly difficult. The term 'relevant' in the previous sentence has various denotations [1], but we focus on domain relevance, which in the legal field is also labelled 'legal importance', 'legal relevance' or 'legal authority'. For case law it can be defined as the importance of a judicial decision for the whole legal community, as distinct from, on the one hand, the influence the decision has on the parties involved, and, on the other hand, the relevance of the case for a particular user of an information system or a specific search query.

Being aware of the (potential) legal importance of a judgment can be of help when selecting cases for publication, or for the end user to filter query results or assess individual cases. The case law database of the European Court of Human Rights offers an example of such a filter. While at this court the cases are classified into four importance levels manually, at the national level, with a multiple number of cases decided, such a manual classification is beyond organizational and financial limits. Some computerized assistance might therefore be of help.

This paper gives account of research on such computer assisted establishment of legal importance. § 1 starts with an overview of previous research and defines the scope of our work. The dataset we used is described in § 2, as are the methods used to extract relevant data from the over one million documents. In § 3 we use social network analysis to explore cross-citations between cases, while in § 4 we examine the possibilities to use other variables to establish legal relevance. In § 5 we wrap up with some conclusions and suggestions for future work.

[1] Corresponding author: Marc van Opijnen, mvanopijnen@rechtspraak.nl.

1. Earlier Work and the Scope of Our Research

Due to the more prominent role of case law in their legal systems, common law countries took the lead in the development of legal citation networks. In the United States systems like Shepard's [2] have been developed decades ago, but using them in more advanced retrieval systems is only of recent date. In early research and on limited datasets Marx [3] and Tapper [4] tried using network characteristics to improve search results. Research of a more exploratory nature by Post and Eisen [5] only weakly confirmed their hypothesis that the supposedly fractal structure of the law could be measured in incoming and outgoing citations. Smith [6], using a larger dataset, came to comparable conclusions. He argued that the web of law behaves like a scale free network, i.a. with a great skewness on cited cases. Fowler c.s. [7] analyzed the citation network of all 27.000 US Supreme Court decisions, trying to identify the most central cases at any point in time. He concluded that using the HITS-algorithm (*infra*, § 3) gives better results than using simple citation counts. He also established that inward relevance and outward relevance are better predictors of future citation than publication in legal magazines or general newspapers.

In civil law systems research has been more scarce. Geist [8] analyzed editorial summaries of decisions by the Austrian *Oberster Gerichtshof* and confirmed the findings of Smith [6] concerning the power law distribution of citations. Using a basic citation count he noticed a correlation between the frequency by which cases are cited in headnotes and the publication of decisions in an influential legal magazine.

Comparable research was done by Malmgren [9] on the case law of the Court of Justice of the European Union. He used the citation information available in EUR-Lex, and took the most thoroughly discussed cases in a textbook on European law as a benchmark. Apart from a plain count of incoming citations and the HITS-algorithm, he also used the PageRank-algorithm (*infra* § 3). He found all three algorithms having a positive result when combined with a standard query, but couldn't draw any definite conclusions on the performance of any of the three measures used. Winkels c.s. [10], using both an in-degree count and PageRank on published cases of Dutch courts, also couldn't find a substantial difference between the two algorithms.

Notwithstanding the significance of the pioneering work discussed above, it leaves many questions unanswered. First, only three algorithms seem to have been tested: a plain citation count, HITS and PageRank. For various reasons, other network measures might be more suitable and should be evaluated. Also the multiplicity of citations is left untouched by most of the cited studies; there might be a substantial difference between case A citing case B just once or many times.

Third, most research has been focussed on decisions from the highest court within a jurisdiction. Although the case law of these courts is quite substantial, the lower courts should not be left out, both because the volume of their published decisions is growing fast, as well as because citation behaviour of lower courts cannot be assumed to be comparable with that of supreme courts.

Fourth, in most cases attention has been limited to cross-citations between judicial decisions, leaving out citations in legal literature. Scholars often have another view on legal relevance and are not confined to cite those decisions that have relevance for the cases that are accidently on their plate. As a result, landmark cases on issues which don't reach court that often run the risk of being undervalued in comparison to less-relevant decisions which are habitually cited in routine cases.

Finally, for establishing legal relevance the focus on citations might be too one-sided. Some of the studies discussed used the publication in legal magazines as a benchmark, but it also might be an indication of legal importance on its own. Also other variables, like references to legislation might indicate legal importance, and have to be assessed.

2. Data Collection

Because it is generally recognized that network analysis is more reliable when more data are used, (nearly) all electronically available case law in the Netherlands was collected: all (174.000)[2] decisions published on the public internet site of the judiciary, nearly all decisions (from 1913 onwards) published by the two largest commercial publishers (292.000), and an internal database of the judiciary (688.000). After removing duplicates 854.000 unique decisions remained.

Scholarly writings (also referred to as 'doctrine') were taken from the two publishers mentioned: 556.000 files from 184 titles (all books, online commentaries, annotations and articles published on their online portals). Also 11.000 press releases from the judiciary's website and 13.000 entries from its internal wiki were analyzed.

After a technical conversion, case law citations were extracted. Multiplicity of citation styles was a major problem: lawyers tend to use all types of identifiers for citing cases: identifiers assigned by commercial periodicals, the normalized national case law identifier ('LJN'), or a combination of court, date and case (or docket) number. To make sure that different types of identifiers all pointed to the same underlying decisions we used a canonicalization algorithm [11]. Also all 'formal citations' (e.g. the Supreme Court citing the appealed case) were filtered out. Finally 412.000 citations in judicial decisions and 673.000 citations in scholarly writings were harvested.

Parsing references to legislation encountered a comparable problem with identifiers. For optimal results we gathered 184.000 titles, abbreviations and informal aliases from the publicly available national legislation database, EUR-Lex and an internal database of the judiciary. For references to secondary EU legislation, which are mostly – albeit hardly correctly – cited by their document number a specific parser was used [12]. Apart from the titles, the parser recognized referred articles and detected local aliases.[3] In total, 5.659.000 references to (paragraphs of) law were collected.

Furthermore, functions of varying complexity were developed to count the number of judges deciding the case, to establish the length of a decision and to extract various other metadata, like the court rendering the decision, the field of law, and the number of times a judgment is published and annotated in periodicals. For the latter, use could be made of a public register [13], containing information on 61 case law magazines.

Finally, a word of caution seems appropriate here. Although a lot of energy has been invested in improving the quality of the gathered and extracted information, some unreliability in the data remains. Partly this is due to the nature of legal language, in which references to legal sources are such an integral part of reasoning that it is not even possible to establish unambiguously which elements should be labelled as citations, even apart from incorrect or unintelligible references. Also the completely

[2] All figures in this paragraph are rounded to the nearest 1000. Data were gathered until June 2010.

[3] A local alias is a temporal name given to a law or regulation. A typical way to declare such an alias: *"Regulation XYZ (herafter referred to as 'the Regulation')"* In the subsequent text, like: *"Article 5 of the Regulation"*, 'the Regulation' has to be interpreted as 'Regulation XYZ'.

unstructured format of most judicial decisions – in which not even the case number, the date or the names of the deciding judges are properly marked – contributed to the unreliability of the research data, as did the great variety of technical formats (often being the result of many previous conversions) in which the documents from so many different sources have been supplied.

3. Analyzing the Citation Network

An analysis of a citation network of this size can easily fill a paper by itself. Because this has been done elsewhere [14, 15], we confine ourselves to some highlights. In line with earlier research (*supra*, § 1) we found a strong power distribution: most cases don't cite other cases, and even more cases are never cited themselves. As expected, considerable differences between types of courts can be observed. Apparently most (83,6%) district court decisions are rendered without any case law reference at all, while for the highest courts this is just 73%. Published decisions though cite other cases more often, and also cite a greater variety of cases. This finding holds for all courts.

With regard to incoming citations it is not surprising to find that nearly all district court cases fall into oblivion, although 7,3% of the published district court cases is – at least once – cited by other decisions. But surprisingly, even 68% of the published decisions of the highest courts remains uncited.

Also, there is a great difference between citations by scholars and judges. While 92,7% of the published district court cases is never cited by another judge, only 75,9% of those cases goes unnoticed by scholars. Table 1 shows another difference between citation in case law and doctrine. Of those cases which are cited at least once, 48,6% is cited in literature only, and not in case law. Decisions cited in case law, but not in literature are much more rare: only 17,1%.

Table 1. The distribution of cases with regard to the number of times they have been cited in literature and/or case law. Only cases which are cited at least once have been included.

	Number of Citations in Case Law:			
Number of Citations in Literature:	**0**	**1-3**	**> 3**	**Total**
0	-	14,3%	2,8%	17,1%
1-3	35,5%	9,8%	3,4%	48,7%
> 3	13,2%	9,1%	12,0%	34,3%
Total	48,6%	33,2%	18,2%	100,0%

Finally, the distribution of citations with regard to the citing and cited court is in line with expectations: all highest courts mostly cite themselves, appellate courts cite the Supreme Court and themselves, but don't cite district courts. The latter though refer to decisions from all courts.

Our next step is to use social network analysis to examine the citation network. Terminologically, it is relevant to distinguish between two types of decisions. Cases having many incoming citations (i.e. often being cited) apparently have influence, and hence are legally relevant. They are also called 'authorities'. Cases having many outgoing citations are 'hubs', they can be said to be 'well-founded in law' [7]. Hub and authority are on a continuum, of course a decision can also be both.

Before we start evaluating various indicators on their ability to measure legal authority, we have to tackle the issue of how to assess the results. Lacking a generally agreed upon measure for legal authority, we used four different benchmarks.

First, we used the number of legal magazines a judgment is published in, for lawyers a common way to assess the authority of a decision. Because publication in magazines is often lagging behind, we shrunk our dataset, removing all decisions from the last six months.

The second benchmark is a subcollection of the first: it contains the number of annotations (i.e. commentaries on judicial decisions, written by the most distinguished scholars) which are published in conjunction with the decision. The benchmark is relevant because it is generally assumed that only landmark cases are annotated.

The third comparison was made with the citations received in legal literature, as legal scholars are assumed to comment on the most relevant decisions.[4] The final benchmark is not external, but counts all incoming citations (with no regard to multiplicity). Because some of the tested algorithms are calculated in a comparable way, a high score on (just) this benchmark should be disregarded.

To compare the performance of network algorithms with the benchmarks we use a Pearson correlation matrix (table 2). We consider correlations > 0,5 as strong, between 0,3 and 0,5 as moderate, and < 0,3 as weak.

The choice of algorithms to assess a decisions' position in the network is limited by the nature of our network, which is – in the general theory of social network analysis – 'directed' (a citation has a source and a target), 'acyclic' (cases can only cite previous cases, not future decisions) with 'weighted' relations (one case can cite another more than once).

We started with the classic centrality measures, [5] of which 'degree centrality' is the most basic: it plainly counts the number of connections to other cases. Because our network is directed, we can calculate 'in-degree' (incoming citations, which we use

Table 2. Pearson matrix showing correlations between network algorithms and external benchmarks. All values are significant with $p < 0,05$.

	Number of Publications	**Number of Annotations**	**Citations in Literature**	**In-degree**
In-degree	0,19	0,19	0,26	-
In-degree Weighted	0,21	0,21	0,29	0,98
All-degree	0,24	0,24	0,31	0,96
All-degree Weighted	0,26	0,27	0,33	0,92
Closeness	0,31	0,24	0,36	0,18
Proximity Prestige	0,30	0,23	0,36	0,18
Betweenness	0,05	0,06	0,09	0,19
HITS-authorities	0,01	0,01	0,01	0,01
PageRank	0,15	0,14	0,18	0,95
Citation Weight SPC	0,06	0,05	0,09	0,09
Citation Weight SPLC	0,06	0,05	0,10	0,09
Citation Weight SPNP	0,06	0,05	0,10	0,09
Generalized Core	0,38	0,33	0,48	0,31
Generalized Core Weighted	0,39	0,37	0,48	0,31
Marc In-degree	0,48	0,42	0,61	0,49
Marc All-degree	0,45	0,38	0,53	0,38

[4] There are different ways to calculate these benchmarks. For publication and annotation a regural count is used. For incoming citations from literature a logarithmic scale as in equation 1 was used .

[5] For the background or formulae of general algorithms see reference works like [16].

also as a benchmark), 'all-degree' and 'out-degree' (the latter is not in the matrix, *infra* § 4). In a weighted variant multiplicity is taken into account. All-degree performs better than in-degree; for both the weighted variant yields better results, although the correlations are hardly moderate.

The drawback of degree centrality is that only direct neighbours are considered: a decisions' influence in a longer line of jurisprudence is not taken into account. To overcome this we can use 'closeness centrality', which calculates the relative distance from one decision to all others. It performs slightly better than all-degree weighted, except against the annotation benchmark (and of course the in-degree).

A combination of degree and closeness was developed with 'proximity prestige'. It values the connected cases on the distance of the connection. It performs equally well as closeness. A final centrality measure is 'betweenness', establishing the degree to which a case is the single connection between two other cases. This concept is not very relevant for case law networks, and its performance is poor indeed.

The next two algorithms were developed for internet searches. The Hyperlink Induced Topic Search (HITS) calculates two vectors: 'hubs' (based on outgoing links) and 'authorities' (based on incoming links). A good hub points to good authorities and a good authority receives links from good hubs, so calculating the algorithm is an iterative process. Although HITS has been developed to be used together with a topical query, it is also used as a general algorithm (*supra* § 1).

'PageRank' is closely related to HITS, but focusses on incoming links and therefore calculates just one vector. HITS seems to be a poor performer. The hub-vector (not in table 2) is not even statistically significant, and HITS-authorities has a very weak correlation with our benchmarks. PageRank performs better than HITS, but the correlation is still weak, performing worse than most of the centrality measures.

Because our network is an (acyclic) citation network we might also invoke algorithms that have been developed for bibliographic citation analysis. We tested three related, but slightly different algorithms: 'search path count' (SPC), 'search path link count' (SPLC) and 'search path node pair' (SPNP). All three show weak correlations.

Next, we wanted to factor in the structure of the network. As was noted earlier, many courts tend to cite themselves, and also decisions within the same field of law are better connected. To test the hypothesis that the grouping of cases might be of influence, we partitioned the network, both on court type and field of law, before calculating the algorithms for each subnetwork. The results were hardly better than on the unpartitioned network, assumably because the grouping of cases is much more complex. So, we entrusted the identification of cohesive subgroups ('cores') to an algorithm. We used 'generalized core', recently developed by Batagelj and Zaveršnik [17]. It turned out to be the best performer thus far, especially in the weighted variant.

Not completely satisfied with the results thus far though, we tried several other algorithms, both well-established and of own design. Studying the network structure and the performance of the other algorithms, three things seemed to be of particular relevance. First, the network is quite incoherent, and lines of jurisprudence are not always apparent from the citations. Direct citations therefore are much more relevant than indirect citations. Second, multiplicity of citations seems to be quite relevant, and third, there has to be a mechanism to mitigate the influence of outliers (i.e. cases cited very frequently). Based on these observations we developed 'Marc degree',[6] which

[6] After the name of our project: '*M*odel for *A*utomated *R*ating of *C*ase law'.

comes in three variants: in-degree, out-degree and all-degree. Although more complex variations were tested, the most basic version turned out to achieve the best results.

If C_i is the number of incoming citations (multiplicity taken into account), than Marc in-degree (M_i) is:

$$M_i = 1 + \log_2 (C_i) \quad (1)$$

Marc out-degree (M_o) is calculated in a comparable way, and Marc all-degree (M_a) is:

$$M_a = M_i + M_o \quad (2)$$

Marc in-degree performs better against all others (disregarding the performance of some algorithms against in-degree, due to their relatedness), and is the only one scoring a 'strong' correlation on one of the benchmarks (and close to it on the others). Marc all-degree performs quite well too, but cannot beat the Marc in-degree.

The next goal of our research is to test the algorithms on different types of court. While table 2 is based on all cases, in table 3 the correlations between the best performing algorithms and the benchmarks are calculated per type of court. The courts are grouped by competence: (nineteen) district courts, (five) courts of appeal, and the highest courts, of which there are four in the Netherlands: Supreme court, Council of State, Central Appeals Tribunal, and the Trade and Industry Appeal Tribunal.

Since we noted earlier – at the beginning of this paragraph – that the higher the court the more incoming and outgoing citations their decisions have, it doesn't come as a surprise that the higher the competence level, the better the performance of the algorithms. A second conclusion to be drawn from table 3 is that the court type doesn't influence the mutual performance of the indicators substantially: Marc in-degree performs best on ten out of twelve; only at the courts of appeal it is (just slightly) beaten on the publication and annotation benchmarks.

Table 3. Pearson matrix showing correlations between network algorithms and external benchmarks, calculated by court type. All values are significant with $p < 0,05$.

	Number of Publications	Number of Annotations	Citations in Literature	In-degree
District Courts				
Closeness	0,12	0,06	0,09	0,07
Proximity Prestige	0,12	0,06	0,09	0,07
Generalized Core Weighted	0,14	0,07	0,10	0,13
Marc In-degree	0,24	0,14	0,23	0,48
Marc All-degree	0,16	0,08	0,11	0,19
Courts of Appeal				
Closeness	0,27	0,14	0,26	0,21
Proximity Prestige	0,27	0,14	0,26	0,21
Generalized Core Weighted	0,27	0,15	0,30	0,42
Marc In-degree	0,26	0,14	0,36	0,67
Marc All-degree	0,31	0,17	0,32	0,43
Highest Courts				
Closeness	0,32	0,28	0,50	0,23
Proximity Prestige	0,32	0,28	0,50	0,23
Generalized Core Weighted	0,37	0,36	0,56	0,31
Marc In-degree	0,45	0,39	0,66	0,48
Marc All-degree	0,44	0,37	0,61	0,40

4. Assessing Other Indicators for Establishing Legal Importance

Publications, annotations and citations by scholars and judges, discussed in the previous paragraph, could be defined as 'exogenous variables': they are not in the decision itself but in its environment. We've tested various algorithms for citations in case law against the other three, but we haven't tested the relationship between those three benchmarks themselves. Our hypothesis is that they will reinforce each other, since (generally) only published cases will be annotated and cited, and citation might lead to (more) publications. But on establishing legal authority, other indicators might be relevant too. This is not only important to discriminate between cases without incoming or outgoing citations, but also too assess the (potential) authority of current decisions: the exogenous variables are useless there, because those cases haven't been published or cited yet. We examined seven other variables that might hint at legal importance. They are (more or less) 'endogenous' – in the decision itself or in its rendering:

- The competence level of the court (according to the classification of table 3;
- The number of judges deciding the case; taking into account the differences between courts regarding the number of judges deciding basic or complex cases;
- Whether or not a news item was published on the courts' website. Decisions with a news item on the homepage of the common portal of the judiciary were considered to be even more important;
- Length of the decision (established via a word count);
- Whether or not the case has a connotation with EU-law. Although in this analysis we refrain from (other) substantive issues, the relative importance of these cases has been established elsewhere [15]. The connotation is assumed if EU-law or a decision of the Court of Justice of the EU has been cited;
- Outgoing case law citations. Lacking a benchmark to test various algorithms, Marc out-degree is used;
- References to legislation. Every single referred element has been counted, on a regular scale and with multiplicity taken into account.

In the end we would like to discover how one dependent variable (authority) is influenced by all these eleven (independent) variables. Lacking such a dependent variable yet, we tried several statistical techniques (i.a. factor analysis) to explore the data, but the conclusions are most clearly shown by the exploratory use of multivariate regression analysis. With this technique one can establish the degree to which the change in a dependent variable ('regressor') can be explained by the change in one or more independent variables ('predictors') – without presupposing a causal relationship.

In table 4 the results of four separate multivariate regression analyses are shown, each one explaining how much a regressor variable (each of the four exogenous variables) is influenced by a change in the predictor variables. R^2 (on the first row) tells us how much (on a scale between 0 and 1) of the regressor can be explained by the predictor variables (i.e. all variables except the regressor itself). The (standardized) regression coefficients (in the other cells) indicate the predicted change in the regressor following a change in the predictor.

As could be expected, the four exogenous variables influence each other. However, in citing other cases judges seem most strongly influenced by literature, and not by the fact that cases are published or annotated. On the other hand, the fact that judges cite a

Table 4. Four separate multivariate regression analyses showing the influence of endogenous and exogenous variables on one regressor variable. The values shown (except for R^2) are the standardized regression coefficients ('b*'). All values are significant with $p < 0,05$, except those between ( )..

	Regressors:			
	Publications	**Annotations**	**Citations in Literature**	**Marc In-degree**
R^2 .	0,64	0,53	0,64	0,43
Predictors:				
Endogenous variables				
Type of Court	0,22	-0,04	-0,04	0,13
Number of Judges	0,13	-0,04	-0,01	0,02
News item	0,07	-0,04	-0,01	(0,00)
Length Decision	0,07	-0,03	0,01	-0,03
EU-connotation	0,02	0,00	(0,00)	-0,02
Marc Out-degree	0,02	-0,01	0,04	0,09
Legislation Cited	0,02	0,04	-0,03	0,04
Exogenous variables				
Publications	-	0,41	0,34	0,08
Annotations	0,30	-	0,31	0,02
Citations in Literature	0,33	0,41	-	0,49
Marc In-degree	0,05	0,02	0,31	-

decision doesn't seem to be a reason for periodicals to publish or annotate it, but it does offer scholars a cause to discuss it.

The endogenous variables all seem to have low impact on the regressor variables. The best performer is the court type, but given the fact that most published and cited decisions are from higher courts, this doesn't come as a surprise. Also, outgoing case law citations seem to have an, albeit weak, influence on (later) incoming citations. This supports the proposition of [7] that strong hubs are potentially strong authorities.

5. Conclusions and Future Work

By publishing, annotating and citing judicial decisions, legal publishers, scholars and judges have been rating case law for over decades. After having made this information computer readable we can use it to calculate legal authority. Our analysis shows that network algorithms that have been used in previous research, especially in-degree, HITS and PageRank, do not seem to be the most appropriate to measure legal authority. From the available algorithms closeness, proximity prestige and especially generalized core perform better. Best results though are obtained by 'Marc in-degree', which – and this is a bit of a surprise – does not take into account more distant relations. These findings should be verified by testing the algorithms on other case law collections.

A second conclusion might be that all our exogenous variables (citations in case law and in doctrine, publications and annotations) are relevant to establish the legal relevance of case law. For optimal results they should be combined, i.a. because the citation behaviour of scholars and judges seems to indicate that they have different opinions on the relevance of cases.

A third conclusion is that using the examined endogenous variables as (sole) predictors of authority won't lead to reliable results.

In general, further improvements in the data and more variables are needed. We'll name just a few possible enhancements. First, so far we rated the publication in all magazines equally, while it is well-known that some magazines are more prestigious

than others. Second, the variable 'references to legislation' could be improved by discounting the more habitual or procedural citations. And third, by factoring in various aspects of time, the model should be able to account for the gradually declining relevance of cases. Fowler [7] did take time into account, but although using a huge set of citing cases, he only used five well-known and often cited cases for his test. In our network, of all 68.285 decisions cited in case law, 88,7% is cited by less than ten other decisions, which might influence the reliability of a time variable. More research on this variable is therefore needed.

A last conclusion might be that statistics by itself probably won't suffice completely to establish legal authority. Although statistical input certainly can be of help, a more 'mathematical' approach that also uses 'legal common sense', might yield better results.

References

[1] T. Saracevic, *Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II: nature and manifestations of relevance*, Journal of the American Society for Information Science and Technology 58 (2007): 1915-1933.

[2] J. F. Spriggs and T. G. Hansford, *Measuring Legal Change: The Reliability and Validity of Shepard's Citations*, Political Research Quarterly 53 (2000): 327-341.

[3] S. M. Marx, *Citation Networks in the Law*, Jurimetrics Journal 10 (1970): 121-137.

[4] C. Tapper, *The use of citation vectors for legal information retrieval*, Journal of Law & Information Science 1 (1982): 131-161.

[5] D. G. Post and M. B. Eisen, *How Long Is the Coastline of the Law? Thoughts on the Fractal Nature of Legal Systems*, The Journal of Legal Studies 29 (2000): 545-584.

[6] T. A. Smith, *The Web of Law*, San Diego Law Review 44 (2007): 309-354.

[7] J. H. Fowler, T. R. Johnson, J. F. Spriggs II, S. Jeon and P. J. Wahlbeck, *Network Analysis and the Law: Measuring the Legal Importance of Precedents at U.S. Supreme Court*, Political Analysis 15 (2007): 324-346.

[8] A. Geist, *Using Citation Analysis Techniques For Computer-Assisted Legal Research In Continental Jurisdictions*, Graduate thesis, University of Edinburgh, 2009.

[9] S. Malmgren, *Towards a theory of jurisprudential relevance ranking. Using link analysis on EU case law*, Graduate thesis, Stockholm University, 2011.

[10] R. Winkels, J. de Ruyter and H. Kroese, *Determining Authority of Dutch Case Law*, in K. M. Atkinson, ed., *Legal Knowledge and Information Systems. JURIX 2011: The Twenty-Fourth International Conference.*, IOS Press, Amsterdam, 2011, pp. 103-112.

[11] M. van Opijnen, *Canonicalizing Complex Case Law Citations*, in R. Winkels, ed., *Legal Knowledge and Information Systems - JURIX 2010: The Twenty-Third Annual Conference*, IOS Press, Amsterdam, 2010, pp. 97-106.

[12] M. van Opijnen, *Searching for References to Secondary EU Legislation*, in S. Tojo, ed., *Fourth International Workshop on Juris-informatics (JURISIN 2010)*, 2010.

[13] M. van Opijnen, *A Public Index of Case Law References – the End of Multiple and Complex Citations*, in T. M. van Engers, ed., *Legal Knowledge and Information Systems - JURIX 2006: The Nineteenth Annual Conference*, IOS Press, Amsterdam, 2006, pp. 71-80.

[14] M. van Opijnen, *Jurisprudentiestatistiek - over de publicatie, annotatie en citatie van rechterlijke uitspraken*, NJB 2011/1658.

[15] M. van Opijnen, *Het communautaire geteld – een kwantitatief onderzoek naar Europeesrechtelijke verwijzingen in rechterlijke uitspraken en rechtswetenschappelijke literatuur*, SEW, tijdschrift voor Europees en economisch recht 2011/98.

[16] W. de Nooy, A. Mrvrar and V. Batagelj, *Exploratory Social Network Analysis with Pajek*, Cambridge University Press, 2011.

[17] V. Batagelj and M. Zaveršnik, *Fast algorithms for determining (generalized) core groups in social networks*, Advances in Data Analysis and Classification 5 (2010): 129-145.