Exploring Network Analysis in a Corpus-Based Approach to Legal Texts: a Case Study

Emilio Sulis, Llio Humphreys, Fabiana Vernero,
Ilaria Angela Amantea, Luigi Di Caro[1], Davide Audrito, Stefano Montaldo[2]

[1] Computer Science Department, University of Turin, Via Pessinetto 12, 10149, Turin, Italy sulis,humphrey,vernero,amantea,dicaro@di.unito.it
[2] Department of Law, University of Turin, Lungo Dora Siena 100, 10153, Turin, Italy stefano.montaldo@unito.it,davide.audrito557@edu.unito.it

Abstract. Automatic analysis of legal texts is increasingly of interest to address the issue of interpretation and compliance concerns. This paper proposes a two-step framework to investigate implicit relationships in legal documents, starting with a corpus-based approach. By introducing an annotation process, the goal is to obtain a gold standard corpus suitable for machine learning experiments. In a second step, we propose a set of features to perform the task of predicting relationships between parts of a norm, as a way to improve legal interpretation. We discuss our first results concerning the annotation task, as well as the adoption of graph-based measures derived from social network analysis. We perform a practical application to an European Union regulation. The proposed framework exploiting network analysis in addition to a corpus-based approach can be applied to address a binary classification task.

Keywords: Information extraction · Network analysis · Text corpora · Natural Language Processing · Legal documents · Legal informatics

1 Introduction

Computer technologies and systems are growing rapidly in different fields, including the legal domain [3]. In recent years a promising research area involves automated information extraction from legal databases and texts. Several approaches and tools focus on solutions to address this challenging task [10].

This paper addresses the specific topic of legal interpretation, defined as the fundamental activity giving meaning to legal documents. The goal of interpretation is to determine the normative messages that arise from a specific legal text. We investigate the subject by focusing on the relationship between internal parts of legislation. In fact, the identification of existing links is at the core of the interpretation process. Moreover, the automatic extraction of information is a challenging task in order to analyse inter-relationships between norm parts. Our work proposes a methodological framework to address an automated system able to consider both the explicit and implicit information in legal text. Therefore, a primary effort necessarily addresses the understanding of the meanings conveyed by the language. In our experience, this is a typical problem that exploits annotated corpora and related tools for the analysis. Furthermore, we propose the investigation of network analysis by shaping inter-relationships as a graph.

In particular, our proposed framework follows a two-step approach to address the issue. Firstly, a corpus-based approach identifies different types of inter-relationships in a norm. This initial step includes the definition of guidelines to address the human annotation effort. The output of this process includes the identification of classes or labels, in order to obtain a gold standard corpus [29] to be used with a supervised learning algorithm in a machine learning experiment [23]. We also explore card sorting to help in identifying patterns among data. The consequent classification effort may include features obtained by adopting a traditional Natural Language Processing (NLP) pipeline.

In a second step, the exploration of a graph-based analysis can help both to better understand the inter-relationships occurring in the norm and to improve the feature set that can be used in a classification task. Moreover, several metrics from social network analysis are compared to the annotation output. In this respect, our research problem is: in a general NLP framework to investigate legal text, are graph metrics helpful to investigate inter-relationships in norms? This work resumes an ongoing experience in developing NLP resources and tools for legal text [7]. In a practical manner, we started to apply the framework presented here to European Union (EU) legislation, specifically mapping related recitals and articles.

The paper first reviews background and related work in Section 2 followed by a description of our corpus-based and graph-based approach in Section 3. We describe the initial results of the proposed system to tackle the research problem in Section 4. We conclude the paper with some future work.

2 Background

Legal interpretation is a well-known mechanism from law for adapting norms to unforeseen situation [15], where several efforts exploit logic to model legal argument [9,22]. In previous work, we investigated the topic by analysing a legal document management system based on ontologies [7]. More recently, a modelling scheme for the inter-relationship between recitals and normative provisions was articulated in [18]. In this work we apply text analysis, a discipline [13] that has gained relevant attention due to its applicability to a wide set of domains. Several NLP techniques have been applied to the analysis of legal texts to address machine learning experiments [24], e.g. the classification of judgment norms [17]. A large part of these kind of techniques relies on manual annotation of text. Human annotation is a very difficult activity, both in terms of time and effort [28]. Moreover, legal texts are often vague and deliberately leave room for multiple interpretations [26]. Therefore, the setup of manual annotation is crucial to perform machine learning experiments. Some initial considerations about the problems with providing the description of the annotation scheme has been recently detailed in [1]. The exploration of machine learning approaches, NLP and network analysis to texts is a recent research subject [21]. Information extraction on legal text mostly refers to citations, such as the semantics-based citation network of [30].

Similar work focused on a network perspective concerning systems of interrelated authoritative legal texts. The detection, resolution and labeling of citations in the legal domain was addressed in [25] by exploring an automatic process of labeling citations in a legal citation graph, while [20] investigate multi-layer network on European legislation, demonstrating the two properties of temporal evolution and multi-scale structures which are very common to many real-word networks. A practical application for a deeper representation of the meaning of legislative text and the relationships between norms can be a legal knowledge management system [8].

This paper explores European legislation. Directives, and regulations to a lesser degree, are prescriptive but sufficiently general to allow addresses (usually member states) to articulate their own detailed norms and procedures as they prefer in order to achieve the goal(s) of the legislation. Not only is this kind of legal text typically goal-oriented, it is also particularly given to principle-based (balance) rather than defeasible reasoning. The initial part of the document consists of recitals, which are ostensibly explanatory and do not have the same status as normative provisions. The relationship between recitals and normative provisions is a debated topic relying on different doctrinal positions [19].

3 Framework

3.1 A corpus-based approach

Our framework combines a quite traditional approach starting from the creation of a gold standard corpus in the development of machine learning experiments based on Natural Language Processing (NLP). This is a prerequisite to perform a classification of relationships in legal texts. In our initial effort, we opted for manual annotation to produce a corpus labeled with classes of interest. Annotation is an important task for NLP, where the common pipeline includes the following steps: identification of classes or labels; the definition of detailed guidelines; training of human annotators; manual annotation of the corpus; and finally computation of inter-annotator agreement.

Identification of labels The annotation of relations in legal text includes a scheme's definition and its application to the legal document. The aim of the scheme is to clearly define the kind of information which must be annotated. This phase includes also the inventory of markers to be used, as well as the annotations granularity. This step relies on the effort of experts in the specific domain of the considered documents.

Annotation phase At this stage, it is very important that each annotator works independently. In this way there will be evidence of any situations of disagreement, which will be resolved later. In a first step, the annotation can be presented in a spreadsheet in order to facilitate the work of annotators. For instance, legal text can be split into different structural parts such as articles or paragraphs for comparison; and different Excel sheets can be used to meaningfully separate different areas of comparison. It can be useful to ask annotators to keep track of their start and end time in order to calculate the average time required for the particular annotation task, Inter-annotation agreement can be measured using metrics such as Cohen's kappa [2]. Another annotator can be involved in this stage to solve cases of disagreement.

Card sorting Commonly used by information architects, card sorting is a popular user-centred method aimed at the identification of patterns among data [27]. Participants, who are asked to work on their own, group physical or digital cards, each displaying a piece of information, based on their own mental model of the information domain. More specifically, while in closed card sorting participants are provided with a set of initial groups, in open card sorting they have no guidance, i.e., they can define the groups which they feel are the most appropriate and then they have to describe each group with a label. Groupings produced by different participants can than be merged by means of a *card * group* matrix, where $cell_{i,y}$ shows the percentage of participants who assigned $card_i$ to $group_y$. In order to obtain a shared classification, the general ratio is to assign each card to the group for which it has the highest percentage, i.e., the highest level of agreement among participants. In the case of legal text, cards can be used to display single paragraphs, the process output consisting in a series of norm groups, as in [1].

3.2 Graph-based NLP approach

Natural Language Processing By following a quite traditional NLP pipeline, the aim of this part is to start with preprocessing the legal steps to obtain stems of terms. In linguistics, stemming means reducing words into their corresponding root form. For instance, it will be possible to compare singular and plural forms of a same term occurring in different parts of the legal text. Typical passages in preprocessing include conversion to lowercase characters, and removal of punctuation marks and stop-words. The next step is tokenisation, in order to separate terms into tokens, followed by stemming. Further analysis is possible by detecting parts of speech (POS) such as verbs, nouns, adverbs, adjectives and so on. The text processing phase can be performed by using common programming languages (e.g., Python or R) with well established NLP libraries.

Identifying features The aim of the second step of the framework is to automatically extract a set of features to be used in machine learning experiments. In particular, by following a Bag-of-Words model, several features can be considered to represent the text. In addition to traditional N-gram models (sequences of N words), several other items can be extracted from text, e.g. the frequency of different parts of speech, the sentiment polarity of different words (e.g. whether they belong to ‘positive or ‘negative dictionaries), word and character length, term frequency (using TF-IDF).

Network analysis We perform a Social Network Analysis on two kinds of graphs. A first graph connects norm types depending on the co-occurrence of terms. If the stem of a term is detected in different parts (e.g., Recital 2 and Article 6) of a legal text, then an edge will connect the two parts. The weight of the edge represents the total amount of co-occurrences. This kind of graph will describe inter-relationships between norms from a linguistic perspective. A second kind of graph explores the role of terms in the document, by linking different stems if they co-occur in the same norm type. The number of times they both appear in the same part of the document will be the weight of the edge. E.g., an edge between Stem X and Stem Y weighted 3 indicates that Stem X and Stem Y co-occur in the same norm type three times. Social network metrics from this graph can be adopted as individual features for the classification effort, e.g. degree, betweenness centrality [14], or versatility [12]. In particular, some metrics will describe the role of the vertex in the graph with respect to the relationships with other vertices.

Weighted multi-layer networks A further investigation of the topic can benefit from the analysis of a multidimentional or multi-layer network. These kinds of network have been receiving increasing attention in the last few years. The opportunity to shape inter-relationships in legal text can be related to the classification effort. As there exist different labels for different kinds of relationship, each label can be a dimension to distinguish a network. In this respect, a legal document may include different networks containing multiple connections between any pair of vertices. Then, each label denotes a different layer, both in the case of a network of stems and norm type. Recent studies introduced a framework to investigate such complex networks having an additional degree of complexity provided by multidimensionality. For instance, a recent work [5] proposed a set of basic concepts and analytical measures concerning multidimensional networks. Specific metrics of this kind of network can be adopted, e.g., the percentage of vertices or edges that belong only to a specific dimension (dimension connectivity) [5].

Classification The contribution of this phase is to formulate an experimental setting in terms of a classification task. The existence of a relation between two parts of the document can be explored with a supervised machine learning experiment carried out by using a model trained with the annotation results. In particular, the framework includes several binary classification tasks, where the classes (or labels) are: i) the existence of a relation of any kind; ii) the existence of a relation of a particular kind (one for each label identified at the beginning of the corpus-based approach. Several classification algorithms can be adopted, e.g. nave Bayes, logistic regression, decision tree, support vector machines. For instance, in this phase the gold standard corpus can be used to train a binary support vector machines with the labelled relationships.

Evaluation Performance of the classifier can be evaluated by computing the F-measure, which provides information on accuracy based on the ratio between precision and recall. Cross-validation can be applied to evaluation. For instance, in ten-fold cross validation the process involves repeating ten times the following steps: i) break the training data into 10 equally-sized partitions; ii) apply the learning algorithm on nine parts, while testing on the remaining folds. The final measure is the performance average of the ten parts. Once an estimation of the model performance has been obtained, this can be applied to new data.

4 First results

4.1 Guidelines

First of all, two legal experts identified eight different kinds of relations between recitals and (sub-)articles based on their analysis of Directive 2004/23/EC in the Italian language. These are: Conceptually Similar (whether using the same or different wording), Constitutive (linking norms containing definitions of legal terms to norms containing those terms), Motivation (where one norm provides the principle or goal that motivates another norm), Impact (in terms of conflicting goals that may restrict one or both of the norms, or norms for enforcement or monitoring that impact on the efficacy of classic deontic norms), Indirect Internal (norms A and C are linked indirectly where norm A cites another internal norm B which is related in another way to norm C), Via Other Law (a norm that is related to another norm which cites another law and cannot be understood without reference to that law), Procedural (linking a norm describing a procedure by an EU institution to support the goal of another norm), and Contextual (linking deontic or other norms to norms that provide contextual information such as jurisdiction and entry into force), and Norm Group (where two or more norms are connected due to being part of the same general requirement). For space reasons, we refer the interested reader to the full descriptions and examples provided in [1]. Annotators were asked to apply these labels to Regulation 141/2000.

An initial result of the corpus-based approach is the creation of a document including the annotation scheme provided by an expert of the domain. This document offers a clear definition for each case as well as a corresponding example. These indications become the shared guidelines for independent annotators.

4.2 Annotation agreement

Two annotators were involved in this initial effort. Each annotator was provided with a spreadsheet file. At the top of each sheet there are spaces to indicate the start and end times concerning the annotation activity, as well as the number of comparisons made. Each annotator was asked to indicate the time spent on multiple work "sessions".

The document (in Italian) is accessible at the following link: `www.di.unito.it/~sulis/NLPxLAW/LineeGuida_ClassificazioneTipiLegami.pdf`

To investigate the annotation phase, we looked at inter-annotation agreement on the relationship between three recitals from the considered law, and their relationships with all the articles, or sub-articles where present. Those three recitals were the first, the shortest and the longest ones. Annotators had to decide which kind of relationship exists (if any) between, for instance, the first Recital (henceforth R1) and Article 1, paragraph 1 (henceforth A1.1), and the type of relationship. If there is a relationship of any kind, a value of 1 is provided, while a value of 0 means no relationship exists. The results can be expressed as a sequence of comma separated values, e.g., a list of triplets for each type:{R1,A1.1,0; R1,A1.2,1; R1,A1.3,0; etc.}.

For analysis of the results, we shall first look at inter-annotator agreement on whether a link exists at all between the recitals and (sub-)articles studied (see Table 1). There are: 86 cases where both annotators agree on whether a link exists, 54 cases where the annotators agree on the existence of a relationship; 32 cases where both annotators agree on the absence of a relationship; 28 cases of disagreement - 22 cases where only the first annotator considers that a relationship exists; and 6 cases where only the second annotator considers that a relationship exists. This provides a percentage of agreement of 75.4%, and a Cohens kappa agreement of 0.5. However, this encouraging result is balanced with significant differences in the labelling of individual kinds of relationships. In fact, there are very different values in the distribution among classes (Table 2).

Table 1. Annotation results on the existence or not of an inter-relationship.

	Recital 1	Recital 4	Recital 8	All 3 recitals
Number of R-A pairs	38	38	38	114
Annotator 1: link exists	21	36	19	76
Annotator 2: link exists	24	25	11	60
Agreed cases of link	20	23	11	54
Agreed cases of lack of link	13	0	19	32
Agreement on whether link exists	33	23	30	86
Disagreement on whether link exists	5	15	8	28
Annotator 1 only: link exists	1	13	8	22
Annotator 2 only: link exists	4	2	0	6

While some classes have been identified with a reasonably similar number of times, others are very distant in their allocation. Qualitative analysis of the results imply that some problems arise due to differences in the directive on which the classification scheme was modelled and the regulation that was annotated. For example, Table 2 shows significant disagreement on the allocation of the Impact label. One annotator extended the Impact category to also include the impact of planned future guidelines, in consultation with member states and other parties, on the interpretation and efficacy of the stated goals of the regulation. The more legislation are annotated, the greater will be our understanding and fine-tuning of the classification scheme. A more expected source of difficulty

Table 2. Annotation results concerning the different types of inter-relationship.

	CS	CNS	MOT	IMP	II	VOL	PR	CNT
Recital 1: 1st Annotator	0	3	16	1	4	7	6	1
Recital 1: 2nd Annotator	0	2	4	9	4	8	8	1
Recital 1 Agreement	0	1	2	0	2	6	5	1
Recital 4: 1st Annotator	1	3	30	0	5	8	5	1
Recital 4: 2nd Annotator	3	2	14	10	3	7	8	1
Recital 4 Agreement	1	2	9	0	3	5	2	1
Recital 8: 1st Annotator	0	3	16	0	5	19	5	1
Recital 8: 2nd Annotator	5	1	8	0	0	3	3	1
Recital 8 Agreement	0	1	8	0	0	3	2	1
All 3 recitals: 1st Annotator	1	9	62	1	14	34	16	3
All 3 recitals: 2nd Annotator	8	5	21	19	7	18	19	3
Agreement on 3 recitals	1	4	19	0	5	14	9	3

was the Via Other Law link: determining whether it should be necessary to consult the referred legislation to fully understand the article or recital in question. A final point that is clear from the data is that the most common type of link was Motivation. This is in common with the accepted purpose of recitals providing explanatory notes to aid understanding of the substantive provisions.

The output of this phase clearly demonstrates the difficulty of the task, as well as the need to refine the labels and revise the annotation guidance. We also conclude that the provision of annotation guidelines is not sufficient to train annotators in this novel classification scheme. For further annotation exercises, we intend to conduct seminars featuring both typical and hard cases in order to ensure proper understanding of the annotation task.

4.3 Timing the annotation process

Our annotation scheme includes recording the start and end time of annotation sessions by each annotator. The average execution time for each annotation session was an average of 62 minutes. Considering the above-mentioned three Recitals, one annotator was faster then the other but split his effort in multiple work sessions: 6 sessions with an average time of 24 minutes, instead of 3 sessions with an average time of 99 minutes.

4.4 Grouping norms with card sorting

The card sorting exercise was carried out by the same people who acted as annotators. Even if this task has not yet been completed, some interesting insights have emerged from the material preparation phase, highlighting the need for methodological adjustments in comparison with traditional card sorting. Firstly, legal experts who acted as our consultants for this activity have pointed out that single parts of legal text can in principle be part of several norm groups. Thus, we had to prepare multiple cards for each recital/paragraph considered. Secondly, we observed that the number of different cards we obtained was higher than the number of cards commonly used in card sorting, and the content in each card was longer as well. Therefore, we plan to assess the annotators' cognitive load, by using both quantitative measures such as NASA-TLX [16] and in-depth, qualitative interviews, in order to adapt card sorting methodology to novel application areas beyond Information Architecture.

4.5 Graph analysis

For processing legal text, we applied typical techniques to extract useful information. For instance, by performing POS tagging on our EU regulation, we computed 1,636 words, with the prevalence of Nouns (1,239 occurrences), Adjectives (171) and Adverbs (58). In addition to typical NLP analysis, we reduced our document (recitals and (sub-)articles) by transforming words into their corresponding stems. We adopted the Python NLTK library [6] to carry out the above NLP processing, storing results into a database (MySql), to finally obtain our edgelist. Finally, we adopted the Gephi [4] open-source software for exploring and manipulating networks. In this step, two kinds of graphs help in the investigation of the inter-relationship between norm types. Figure 1 shows the representations of our legal text. To better focus on the visualisation of more promising vertices, in both cases the ones with low degrees are pruned. In both graphs, the size of each vertex is related to degree, while the size of the label indicates the betweenness centrality.

Graph of Recitals and Articles A first graph investigates the relationship between recitals and articles considering each of them as a vertex. This kind of representation is the one on the left in Figure 1, by including only the vertices with a degree higher than 10 to improve the readability of the network. Edges are proportional to their weight, i.e. the number of terms co-occurring between the two corresponding vertices. These values may represent the strength of the relationship between the two parts of the law. Graph metrics indicates here the strength of the relation. Once the annotation effort of implicit links has been completed, it would be interesting to correlate the results about inter-relationships of a legal text with the corresponding graph metrics, i.e. degree or centrality measures. We hypothesize a certain similarity between the lexical graph-based measures and the existence of a some kinds of implicit relationships (e.g. Conceptually Similar, Motivation) between different parts of the text.

Graph of Terms A second graph analysis explores the relationships among terms, after the stemming process. By considering stems as vertices, the edges represent the co-occurrences in the legal text under consideration. Figure 1 (b) describes a sample of this kind of graph by considering only vertices having a degree higher than 100. The weight of an edge represents the co-occurrences of stems in the document. Several metrics are computed for each vertex to measure the corresponding relevance in the graph. With the aim of assessing the relevance of the vertex in the graph we report two measures: degree and centrality. Degree varies in a range from 1 to 68, with an average value of 52. Stems with higher betweenness centrality refer to both the topic of the current law (i.e., "Medicinal", 21.8; "Pazient", 16.9) and common terms (i.e., "Tal", 21.9; "Scop", 17.1).

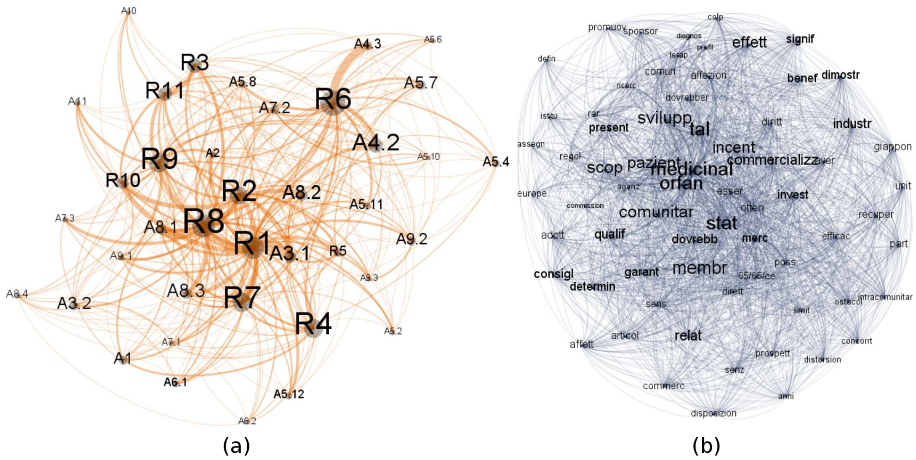

Fig. 1. Examples of graphs to investigate relations between recitals and (sub-)articles. In (a) the vertices are both Recitals (R) and (sub-)articles (A), where edges represent a link if two vertices have at least one common stem (the weight sum of all the co-occurring stems in the two parts of the document). In (b) vertices are stems, while weighted edges represent the number of parts of the text in which they both appear.

4.6 Feature set

The representation of each part of the legal text can be addressed with a bag-of-words approach, or including n-grams to capture more context around each term, as well as POS-based features (e.g., the number of nouns, adjectives, etc. in each part). The output of this step consists in a set of features collected as a numeric feature vector. The goal is to obtain a feature model to perform binary classification. We add to these traditional approaches the information derived from graph analysis. By considering the graph with stems as vertices, each vertex metric can be used to create a feature of the corresponding parts of the legal text (e.g., articles or recitals). For instance, the average value of vertices' metrics (e.g. degree, betwenness centrality, closeness centrality) can be exploited in the classification step. Similar considerations concern the graph of relations between stems co-occurring in different parts of the legal text. In this case, the final set of features may include network metrics depending on the network topology.

5 Conclusions and future work

This work proposed a framework to develop an automated system to find interrelated parts of a norm. These links may help legal professionals to interpret legal issues for specific cases. After describing the framework, we described a practical work-in-progress application carried out at the University of Turin (Computer Science Department and Law Department). In our data-centric framework we address the task of automatically identifying implicit internal links between norms in the same legislation. In particular, we introduce here the development of a corpus of annotated links mapping recitals and articles. We initially applied our framework on EU legislation. This kind of legal text is deliberately imprecise, to allow members to fulfill the objectives in their own way. The first analysis of this step carried out by two annotators clearly shows the complexity of the task: while the agreement seems quite satisfactory in the identification of inter-related parts, the disagreement on the types of relationship is very high. In addition to the corpus-based approach, we took the opportunity to operate a classification task. The experiment can exploit a set of features extracted from traditional NLP approaches. In addition, we explored here the adoption of features from a graph analysis to improve the results. In this framework, graph-based metrics concerning vertices can be of interest as a feature of a machine learning experiment. As a future work, we plan to adopt such metrics to perform binary classification tasks. The classes will be those adopted in the annotation phase, i.e. the presence or absence of a link as well as the kind of existing relationship. Special attention will be given to multidimensional network analysis [11], where the different networks of inter-relationships can be shaped depending on the label result of the annotation to explore new features of interest.

References

1. Amantea, I.A., Caro, L.D., Humphreys, L., Nanda, R., Sulis, E.: Modelling norm types and their inter-relationships in EU directives. In: Proc. of the Third Workshop ASAIL co-located at ICAIL, Montreal. CEUR-WS.org (2019)
2. Artstein, R., Poesio, M.: Inter-coder agreement for computational linguistics. Computational Linguistics **34**(4), 555–596 (2008)
3. Ashley, K.D.: Artificial intelligence and legal analytics: new tools for law practice in the digital age. Cambridge University Press (2017)
4. Bastian, M., Heymann, S., Jacomy, M.: Gephi: An open source software for exploring and manipulating networks. In: Third international AAAI conference on weblogs and social media (03 2009)
5. Berlingerio, M., Coscia, M., Giannotti, F., Monreale, A., Pedreschi, D.: Foundations of multidimensional network analysis. In: 2011 International Conference on Advances in Social Networks Analysis and Mining. pp. 485–489 (2011)
6. Bird, S., Klein, E., Loper, E.: Natural Language Processing with Python. O'Reilly Media (2009)
7. Boella, G., Caro, L.D., Humphreys, L., Robaldo, L., Rossi, P., Torre, L.: Eunomos, a legal document and knowledge management system for the web to provide relevant, reliable and up-to-date information on the law. Artificial Intelligence and Law **24**(3), 245283 (Sep 2016)
8. Boella, G., Di Caro, L., Leone, V.: Semi-automatic knowledge population in a legal document management system. Artificial Intelligence and Law **27**, 227–251 (12 2018)
9. Branting, L.K.: Data-centric and logic-based models for automated legal problem solving. Artificial Intelligence and Law **25**(1), 5–27 (2017)
10. Cheng, T.T., Cua, J.L., Tan, M.D., Yao, K.G., Roxas, R.E.: Information extraction from legal documents. In: 2009 Eighth International Symposium on Natural Language Processing. pp. 157–162 (2009)

11. Dickison, M.E., Magnani, M., Rossi, L.: Multilayer social networks. Cambridge University Press (2016)
12. Domenico, M.D., Solé-Ribalta, A., Omodei, E., Gómez, S., Arenas, A.: Ranking in interconnected multilayer networks reveals versatile nodes. Nature communications **6**, 6868 (2015)
13. Fairclough, N.: Discourse and text: Linguistic and intertextual analysis within discourse analysis. Discourse & society **3**(2), 193–217 (1992)
14. Freeman, L.C.: A Set of Measures of Centrality Based on Betweenness. Sociometry **40**(1), 35–41 (Mar 1977)
15. Gordon, T.F.: The role of exceptions in models of the law. Formalisierung im Recht und Ansätze juristischer Expertensysteme pp. 52–59 (1986)
16. Hart, S.G., Staveland, L.E.: Development of nasa-tlx (task load index): Results of empirical and theoretical research. In: Hancock, P.A., Meshkati, N. (eds.) Human Mental Workload, Advances in Psychology, vol. 52, pp. 139 – 183. North-Holland (1988)
17. Howe, J.S.T., Khang, L.H., Chai, I.E.: Legal area classification: A comparative study of text classifiers on singapore supreme court judgments. arXiv preprint arXiv:1904.06470 (2019)
18. Humphreys, L., Santos, C., Di Caro, L., Boella, G., Van Der Torre, L., Robaldo, L.: Mapping recitals to normative provisions in eu legislation to assist legal interpretation. In: JURIX. pp. 41–49 (2015)
19. Klimas, T., Vaiciukaite, J.: The law of recitals in european community legislation. ILSA Journal of International & Comparative Law **15**(6), 61–93 (7 2008)
20. Koniaris, M., Anagnostopoulos, I., Vassiliou, Y.: Network analysis in the legal domain: A complex model for european union legal sources. Journal of Complex Networks **6**(2), 243–268 (2018)
21. Mihalcea, R., Radev, D.: Graph-based natural language processing and information retrieval. Cambridge university press (2011)
22. Prakken, H.: A logical framework for modelling legal argument. In: Proceedings of the 4th international conference on Artificial intelligence and law. pp. 1–9 (1993)
23. Pustejovsky, J., Stubbs, A.: Natural Language Annotation for Machine Learning: A guide to corpus-building for applications. O'Reilly Media, Inc. (2012)
24. Robaldo, L., Villata, S., Wyner, A., Grabmair, M.: Introduction for artificial intelligence and law: special issue natural language processing for legal texts (2019)
25. Sadeghian, A., Sundaram, L., Wang, D., Hamilton, W., Branting, K., Pfeifer, C.: Semantic edge labeling over legal citation graphs. In: Proceedings of the workshop on legal text, document, and corpus analytics (LTDCA-2016). pp. 70–75 (2016)
26. Sainsbury, M.: Logical forms: An introduction to philosophical logic. John Wiley & Sons (2000)
27. Spencer, D.: Card Sorting: Designing Usable Categories. Rosenfeld Media, 1 edn. (4 2009)
28. Tomanek, K., Wermter, J., Hahn, U.: An approach to text corpus construction which cuts annotation costs and maintains reusability of annotated data. In: Proceedings of the 2007 EMNLP-CoNLL. pp. 486–495 (2007)
29. Wissler, L., Almashraee, M., Díaz, D.M., Paschke, A.: The gold standard in corpus annotation. In: IEEE GSC, Passau (06 2014)
30. Zhang, P., Koppaka, L.: Semantics-based legal citation network. In: Proceedings of the 11th International Conference on Artificial Intelligence and Law. p. 123130. ICAIL 07, Association for Computing Machinery, New York, NY, USA (2007)