

Article

Evaluation of Diversification Techniques for Legal Information Retrieval [†]

Marios Koniaris ^{1,*}, Ioannis Anagnostopoulos ² and Yannis Vassiliou ¹

¹ Knowledge and Database Systems Laboratory, Division of Computer Science, School of Electrical and Computer Engineering, National Technical University of Athens, Iroon Polytechniou 9, Zographou Campus, 15780 Athens, Greece; yv@cs.ntua.gr

² Department of Computer Science and Biomedical Informatics, School of Sciences, University of Thessaly, Papasiopoulou 2-4, 35131 Lamia, Greece; janag@dib.uth.gr

* Correspondence: mkoniari@dblab.ece.ntua.gr; Tel.: +30-210-7721602

[†] This paper is an extended version of our paper published in Koniaris, M.; Anagnostopoulos, I.; Vassiliou, Y. Diversifying the Legal Order. In Proceedings of the Artificial Intelligence Applications and Innovations: 12th IFIP WG 12.5 International Conference and Workshops, AIAI 2016, Thessaloniki, Greece, 16–18 September 2016; pp. 499–509.

Academic Editors: Katia Lida Kermanidis, Christos Makris, Phivos Mylonas and Spyros Sioutas

Received: 22 November 2016; Accepted: 19 January 2017; Published: 29 January 2017

Abstract: “Public legal information from all countries and international institutions is part of the common heritage of humanity. Maximizing access to this information promotes justice and the rule of law”. In accordance with the aforementioned declaration on free access to law by legal information institutes of the world, a plethora of legal information is available through the Internet, while the provision of legal information has never before been easier. Given that law is accessed by a much wider group of people, the majority of whom are not legally trained or qualified, diversification techniques should be employed in the context of legal information retrieval, as to increase user satisfaction. We address the diversification of results in legal search by adopting several state of the art methods from the web search, network analysis and text summarization domains. We provide an exhaustive evaluation of the methods, using a standard dataset from the common law domain that we objectively annotated with relevance judgments for this purpose. Our results: (i) reveal that users receive broader insights across the results they get from a legal information retrieval system; (ii) demonstrate that web search diversification techniques outperform other approaches (e.g., summarization-based, graph-based methods) in the context of legal diversification; and (iii) offer balance boundaries between reinforcing relevant documents or sampling the information space around the legal query.

Keywords: diversity; algorithms; legal information retrieval

1. Introduction

Nowadays, as a consequence of many open data initiatives, more and more publicly available portals and datasets provide legal resources to citizens, researchers and legislation stakeholders. Thus, legal data that were previously available only to a specialized audience and in “closed” format are now freely available on the Internet.

Portals, such as the EUR-Lex (<http://eur-lex.europa.eu/>), the European Union’s database of regulations, the on-line version of the United States Code (<http://uscode.house.gov/>), United Kingdom (<http://www.legislation.gov.uk/>), Brazil (<http://www.lexml.gov.br/>) and the Australian one (<https://www.comlaw.gov.au/>), just to mention a few, serve as an endpoint to access millions of regulations, legislation, judicial cases or administrative decisions. Such portals allow for

multiple search facilities, so as to assist users to find the information they need. For instance, the user can perform simple search operations or utilize predefined classificatory criteria, e.g., year, legal basis, subject matter, to find relevant to her/his information needs legal documents.

At the same time, the legal domain generates a huge and ever-increasing amount of open data, thus leveraging the awareness level of legislation stakeholders. Judicial ruling, precedents and interpretations of laws create a big data space with relevant and useful legal resources, while many facts and insights that could help win a legal argument usually remain hidden. For example, it is extremely difficult to search for a relevant case law, by using Boolean queries or the references contained in the judgment. Consider, for example, a patent lawyer who want to find patents as a reference case and submits a user query to retrieve information. A diverse result, i.e., a result containing several claims, heterogeneous statutory requirements and conventions, varying in the numbers of inventors and other characteristics, is intuitively more informative than a set of homogeneous results that contain only patents with similar features. In this paper, we propose a novel way to efficiently and effectively handle similar challenges when seeking information in the legal domain.

Diversification is a method of improving user satisfaction by increasing the variety of information shown to the user. As a consequence, the number of redundant items in a search result list should decrease, while the likelihood that a user will be satisfied with any of the displayed results should increase. There has been extensive work on query results' diversification (see Section 4), where the key idea is to select a small set of results that are sufficiently dissimilar, according to an appropriate similarity metric.

Diversification techniques in legal information systems can be helpful not only for citizens, but also for law issuers and other legal stakeholders in companies and large organizations. Having a big picture of diversified results, issuers can choose or properly adapt the legal regime that better fits their firms and capital needs, thus helping them operate more efficiently. In addition, such techniques can also help lawmakers, since deep understanding of legal diversification promotes evolution to better and fairer legal regulations for society [1].

In this work, extending our previous work [2], we address result diversification in the legal information retrieval (IR). To this end, in parallel with the search result diversification methods utilized in our previous work (maximal marginal relevance (MMR) [3], Max-sum [4], Max-min [4] and Mono-objective [4]), we extended the methods utilized by incorporating methods introduced for text summarization (LexRank [5] and Biased LexRank [6]) and graph-based ranking (DivRank [7] and Grasshopper [8]). We present the diversification methods utilized with an algorithmic presentation, alongside with a textual explanation of each diversification objective. We evaluate the performance of the above methods on a legal corpus annotated with relevant judgments using metrics employed in TREC Diversity Tasks, fine-tuning separately for each method, trade-off values between finding relevant to the user query documents and diverse documents in the result set.

Our findings reveal that (i) diversification methods, employed in the context of legal IR, demonstrate notable improvements in terms of enriching search results with otherwise hidden aspects of the legal query space and (ii) web search diversification techniques outperform other approaches, e.g., summarization-based, graph-based methods, in the context of legal diversification. Furthermore, our accuracy analysis can provide helpful insights for legal IR systems, wishing to balance between reinforcing relevant documents, result set similarity, or sampling the information space around the query, result set diversity.

The remainder of this paper is organized as follows: Section 2 introduces the concepts of search diversification and presents diversification algorithms, while Section 3 describes our experimental results and discuss their significance. Section 4 reviews previous work in query result diversification, diversified ranking on graphs and in the field of legal text retrieval, while it stresses the differentiation and contribution of this work. Finally, we draw our conclusions and present future work aspects in Section 5.

2. Legal Document Ranking Using Diversification

At first, we define the problem addressed in this paper and provide an overview of the diversification process. Afterwards, legal document's features relevant for our work are introduced, and distance functions are defined. Finally, we describe the diversification algorithms used in this work.

2.1. Diversification Overview

Result diversification is a trade-off between finding relevant to the user query documents and diverse documents in the result set. Given a set of legal documents and a query, our aim is to find a set of relevant and representative documents and to select these documents in such a way that the diversity of the set is maximized. More specifically, the problem is formalized as follows:

Definition 1 (Legal document diversification). *Let q be a user query and N a set of documents relevant to the user query. Find a subset $S \subseteq N$ of documents that maximize an objective function f that quantifies the diversity of documents in S .*

$$S = \underset{\substack{|S|=k \\ S \subseteq N}}{\operatorname{argmax}} f(N) \quad (1)$$

Figure 1 illustrates the overall workflow of the diversification process. At the highest level, the user submits his/her query as a way to express an information need and receives relevant documents, as shown in the left column of Figure 1, where different names and colors represent the main aspects/topics of documents. From the relevance-oriented ranking of documents, we derive a diversity-oriented ranking, produced by seeking to achieve both coverage and novelty at the same time. This list is visualized in the right column of Figure 1. Note that Document E was initially hidden in the left column, the relevance-oriented ranked result set list.

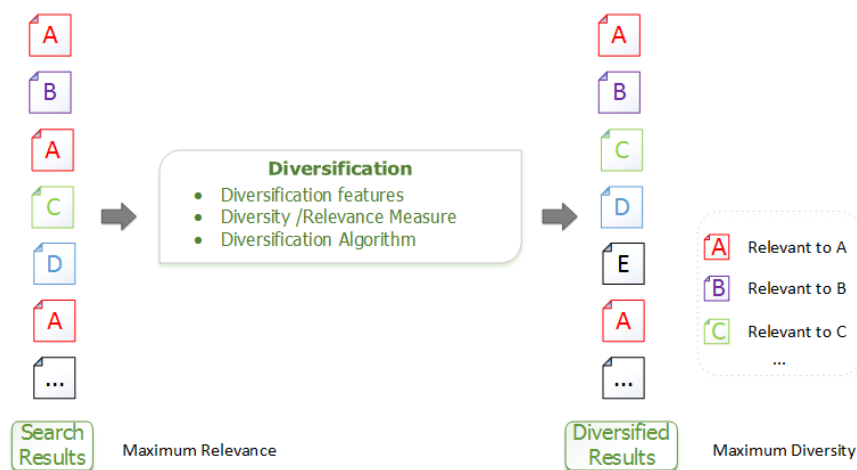


Figure 1. Diversification overview. Each document is identified by name/color representing the main aspects/topics of documents. A diversity-oriented ranked list of the documents is obtained in the right column, through diversification of the relevance-oriented ranked result set list in the left column.

Significant components of the process include:

- Ranking features, features of legal documents that will be used in the ranking process.
- Distance measures, functions to measure the similarity between two legal documents and the relevance of a query to a given document. We note that in this work, as in information retrieval in general [9], the term “distance” is used informally to refer to a dissimilarity measure derived from the characteristics describing the objects.
- Diversification heuristics, heuristics to produce a subset of diverse results.

2.2. Ranking Features/Distance Measures

Typically, diversification techniques measure diversity in terms of content, where textual similarity between items is used in order to quantify information similarity. In the vector space model [10], each document u can be represented as a term vector $U = (is_{w1u}, is_{w2u}, \dots, is_{wm_u})^T$, where w_1, w_2, \dots, w_m are all of the available terms, and is can be any popular indexing schema, e.g., $tf, tf - idf, logtf - idf$. Queries are represented in the same manner as documents.

Following, we define:

- Document similarity: Various well-known functions from the literature (e.g., Jaccard, cosine similarity, etc.) can be employed for computing the similarity of legal documents. In this work, we choose cosine similarity as a similarity measure; thus, the similarity between documents u and v , with term vectors U and V is:

$$sim(u, v) = \cos(u, v) = \frac{U \cdot V}{\|U\| \|V\|} \quad (2)$$

- Document distance: The distance of two documents is:

$$d(u, v) = 1 - sim(u, v) \quad (3)$$

- Query document similarity. The relevance of a query q to a given document u can be assigned as the initial ranking score obtained from the IR system, or calculated using the similarity measure, e.g., cosine similarity of the corresponding term vectors:

$$r(q, u) = \cos(q, u) \quad (4)$$

2.3. Diversification Heuristics

Diversification methods usually retrieve a set of documents based on their relevance scores and then re-rank the documents so that the top-ranked documents are diversified to cover more query subtopics. Since the problem of finding an optimum set of diversified documents is NP-hard, a greedy algorithm is often used to iteratively select the diversified set S .

Let N be the document set, $u, v \in N$, $r(q, u)$ the relevance of u to the query q , $d(u, v)$ the distance of u and v , $S \subseteq N$ with $|S| = k$ the number of documents to be collected and $\lambda \in [0..1]$ a parameter used for setting the trade-off between relevance and diversity. In this paper, we focus on the following representative diversification methods:

- MMR: Maximal marginal relevance [3], a greedy method to combine query relevance and information novelty, iteratively constructs the result set S by selecting documents that maximize the following objective function:

$$f_{MMR}(u, q) = (1 - \lambda) r(u, q) + \lambda \sum_{v \in S} d(u, v) \quad (5)$$

MMR incrementally computes the standard relevance-ranked list when the parameter $\lambda = 0$ and computes a maximal diversity ranking among the documents in N when $\lambda = 1$. For intermediate values of $\lambda \in [0..1]$, a linear combination of both criteria is optimized. In MMR Algorithm 1, the set S is initialized with the document that has the highest relevance to the query. Since the selection of the first element has a high impact on the quality of the result, MMR often fails to achieve optimum results.

Algorithm 1 Produce diverse set of results with MMR.**Input:** Set of candidate results N , size of diverse set k **Output:** Set of diverse results $S \subseteq N, |S| = k$ $S = \emptyset$ $N_i = \operatorname{argmax}_{v \in N} (r(v, q))$

▷ Initialize with the highest relevant to the query document

Set $S = S \cup \{i\}$ Set $N = N \setminus \{i\}$ **while** $|S| < k$ **do**Find $u = \operatorname{argmax}_{v \in N} (f_{MMR}(v, q))$

▷ Iteratively select document that maximize Equation (5)

Set $S = S \cup \{u\}$ Set $N = N \setminus \{u\}$ **end while**

- Max-sum: The Max-sum diversification objective function [4] aims at maximizing the sum of the relevance and diversity in the final result set. This is achieved by a greedy approximation, Algorithm 2, that selects a pair of documents that maximizes Equation (6) in each iteration.

$$f_{MAXSUM}(u, v, q) = (1 - \lambda) (r(u, q) + r(v, q)) + 2\lambda d(u, v) \quad (6)$$

where (u, v) is a pair of documents, since this objective considers document pairs for insertion. When $|S|$ is odd, in the final phase of the algorithm, an arbitrary element in N is chosen to be inserted in the result set S .

Max-sum Algorithm 2, at each step, examines the pairwise distances of the candidate items N and selects the pair with the maximum pairwise distance, to insert into the set of diverse items S .

Algorithm 2 Produce diverse set of results with Max-sum.**Input:** Set of candidate results N , size of diverse set k **Output:** Set of diverse results $S \subseteq N, |S| = k$ $S = \emptyset$ **for** $i = 1 \rightarrow \lfloor \frac{k}{2} \rfloor$ **do**Find $(u, v) = \operatorname{argmax}_{x, y \in N} (f_{MAXSUM}(x, y, q))$

▷ Select pair of docs that maximize Equation (6)

Set $S = S \cup \{u, v\}$ Set $N = N \setminus \{u, v\}$ **end for****if** k is odd **then** $S = S \cup \{i\}, N_i \in N$ ▷ If k is odd add an arbitrary document to S **end if**

- Max-min: The Max-min diversification objective function [4] aims at maximizing the minimum relevance and dissimilarity of the selected set. This is achieved by a greedy approximation, Algorithm 3, that initially selects a pair of documents that maximize Equation (7) and then in each iteration selects the document that maximizes Equation (8):

$$f_{MAXMIN}(u, v, q) = (1 - \lambda) (r(u, q) + r(v, q)) + \lambda d(u, v) \quad (7)$$

$$f_{MAXMIN}(u, q) = \min_{v \in S} d(u, v) \quad (8)$$

Max-min Algorithm 3, at each step, it finds, for each candidate document, its closest document belonging to S and calculates their pairwise distance d_{MIN} . The candidate document that has the maximum distance d_{MIN} is inserted into S .

Algorithm 3 Produce diverse set of results with Max-min.**Input:** Set of candidate results N , size of diverse set k **Output:** Set of diverse results $S \subseteq N, |S| = k$ $S = \emptyset$ Find $(u, v) = \operatorname{argmax}_{x, y \in N} (f_{\text{MAXMIN}}(x, y, q))$ \triangleright Initially selects documents that maximize Equation (7)Set $S = S \cup \{u, v\}$ **while** $|S| < k$ **do** Find $u = \operatorname{argmax}_{x \in N \setminus S} (f_{\text{MAXMIN}}(x, q))$ \triangleright Select document that maximize Equation (8) Set $S = S \cup \{u\}$ **end while**

- Mono-objective: Mono-objective [4] combines the relevance and the similarity values into a single value for each document. It is defined as:

$$f_{\text{MONO}}(u, q) = r(u, q) + \frac{\lambda}{|N| - 1} \sum_{v \in N} d(u, v) \quad (9)$$

Algorithm 4 approximates the Mono-objective. The algorithm, at the initialization step, calculates a distance score for each candidate document. The objective function weights each document's similarity to the query with the average distance of the document with the rest of the documents. After the initialization step, where scores are calculated, they are not updated after each iteration of the algorithm. Therefore, each step consists of selecting the document from the remaining candidates set with the maximum score and inserting it into S .

Algorithm 4 Produce diverse set of results with Mono-objective.**Input:** Set of candidate results N , size of diverse set k **Output:** Set of diverse results $S \subseteq N, |S| = k$ $S = \emptyset$ **for** $x_i \in N$ **do** $d(x_i) = f_{\text{MONO}}(x, q)$ \triangleright Calculate scores based on Equation (9)**end for****while** $|S| < k$ **do** Find $u = \operatorname{argmax}_{x_i \in N} d(x_i)$ \triangleright Sort and select *top* – k documents Set $S = S \cup \{u\}$ Set $N = N \setminus \{u\}$ **end while**

- LexRank: LexRank [5] is a stochastic graph-based method for computing the relative importance of textual units. A document is represented as a network of inter-related sentences, and a connectivity matrix based on intra-sentence similarity is used as the adjacency matrix of the graph representation of sentences.

In our setting, instead of sentences, we use documents that are in the initial retrieval set N for a given query. In this way, instead of building a graph using the similarity relationships among the sentences based on an input document, we utilize document similarity on the result set. If we consider documents as nodes, the result set document collection can be modeled as a graph by generating links between documents based on their similarity score as in Equation (2). Typically, low values in this matrix can be eliminated by defining a threshold so that only significantly similar documents are connected to each other. However, as in all discretization operations, this means information loss. Instead, we choose to utilize the strength of the similarity links. This way we use the cosine values directly to construct the similarity graph, obtaining a much denser, but weighted graph. Furthermore, we normalize our adjacency matrix B , so as to make the sum of each row equal to one.

Thus, in LexRank scoring formula Equation (10), Matrix B captures pairwise similarities of the documents, and square matrix A , which represents the probability of jumping to a random node in the graph, has all elements set to $1/|N|$, where $|N|$ is the number of documents.

$$p = [\lambda A + (1 - \lambda) B]^T p \quad (10)$$

LexRank Algorithm 5 applies a variation of PageRank [11] over a document graph. A random walker on this Markov chain chooses one of the adjacent states of the current state with probability $1 - \lambda$ or jumps to any state in the graph, including the current state, with probability λ .

Algorithm 5 Produce diverse set of results with LexRank.

Input: Set of candidate results N , size of diverse set k

Output: Set of diverse results $S \subseteq N, |S| = k$

$S = \emptyset$

$A_{|N||N|} = 1/|N|$

for $u, v \in N$ **do**

$B_{u,v} = d(u, v)$ ▷ Calculate connectivity matrix based on document similarity Equation (2)

end for

$p = f_{powermethod}(B, A)$ ▷ Calculate stationary distribution of Equation (10). (Omitted for clarity)

while $|S| < k$ **do**

Find $u = \operatorname{argmax}_{x_i \in N} p(x_i)$ ▷ Sort and select $top - k$ documents

Set $S = S \cup \{u\}$

Set $N = N \setminus \{u\}$

end while

- Biased LexRank: Biased LexRank [6] provides for a LexRank extension that takes into account a prior document probability distribution, e.g., the relevance of documents to a given query. The Biased LexRank scoring formula is analogous to LexRank scoring formula Equation (10), with matrix A , which represents the probability of jumping to a random node in the graph, proportional to the query document relevance.

Algorithm 5 is also used to produce a diversity-oriented ranking of results with the Biased LexRank method. In Biased LexRank scoring formula Equation (10), we set matrix B as the connectivity matrix based on document similarity for all documents that are in the initial retrieval set N for a given query and matrix A elements proportional to the query document relevance.

- DivRank: DivRank [7] balances popularity and diversity in ranking, based on a time-variant random walk. In contrast to PageRank [11], which is based on stationary probabilities, DivRank assumes that transition probabilities change over time; they are reinforced by the number of previous visits to the target vertex. If $p_T(u, v)$ is the transition probability from any vertex u to vertex v at time T , $p^*(d_j)$ is the prior distribution that determines the preference of visiting vertex d_j and $p_0(u, v)$ is the transition probability from u to v prior to any reinforcement, then,

$$p_T(d_i, d_j) = (1 - \lambda) \cdot p^*(d_j) + \lambda \cdot \frac{p_0(d_i, d_j) \cdot N_T(d_j)}{D_T(d_i)} \quad (11)$$

where $N_T(d_j)$ is the number of times the walk has visited d_j up to time T and,

$$D_T(d_i) = \sum_{d_j \in V} p_0(d_i, d_j) N_T(d_j) \quad (12)$$

DivRank was originally proposed in a query independent context; thus, it is not directly applicable to the diversification of search results. We introduce a query dependent prior and thus utilize DivRank as a query-dependent ranking schema. In our setting, we use documents that are in the

initial retrieval set N for a given query q , create the citation network between those documents and apply the DivRank Algorithm 6 to select $top - k$ divers' documents in S .

Algorithm 6 Produce diverse set of results with DivRank.

Input: Set of candidate results N , size of diverse set k

Output: Set of diverse results $S \subseteq N, |S| = k$

$S = \emptyset$

for $u, v \in N$ **do**

$B(u, v) = A_{u,v}$ ▷ Connectivity matrix is based on citation network adjacency matrix

end for

$p = f_{powermethod}(B)$ ▷ Calculate stationary distribution of Equation (11). (Omitted for clarity)

while $|S| < k$ **do**

Find $u = \operatorname{argmax}_{x_i \in N} p(x_i)$ ▷ Sort and select $top - k$ documents

Set $S = S \cup \{u\}$

Set $N = N \setminus \{u\}$

end while

- Grasshopper: Similar to the DivRank ranking algorithm, it is described in [8]. This model starts with a regular time-homogeneous random walk, and in each step, the vertex with the highest weight is set as an absorbing state.

$$p_T(d_i, d_j) = (1 - \lambda) \cdot p^*(d_j) + \lambda \cdot \frac{p_0(d_i, d_j) \cdot N_T(d_j)}{D_T(d_i)} \quad (13)$$

where $N_T(d_j)$ is the number of times the walk has visited d_j up to time T and,

$$D_T(d_i) = \sum_{d_j \in V} p_0(d_i, d_j) N_T(d_j) \quad (14)$$

Since Grasshopper and DivRank utilize a similar approach and will ultimately present rather similar results, we utilized Grasshopper distinctively from DivRank. In particular, instead of creating the citation network of documents belonging to the initial result set, we form the adjacency matrix based on document similarity, as previously explained in LexRank Algorithm 5.

3. Experimental Setup

In this section, we describe the legal corpus we use, the set of query topics and the respective methodology for annotating with relevance judgments for each query, as well as the metrics employed for the evaluation assessment. Finally, we provide the results along with a short discussion.

3.1. Legal Corpus

Our corpus contains 3890 Australian legal cases from the Federal Court of Australia (<http://www.fedcourt.gov.au>). The cases were originally downloaded from the Australasian Legal Information Institute (<http://www.austlii.edu.au>) and were used in [12] to experiment with automatic summarization and citation analysis. The legal corpus contains all cases from the Federal Court of Australia spanning from 2006 up to 2009. From the cases, we extracted all text and citation links for our diversification framework. Our index was built using standard stop word removal and Porter stemming, with the log-based $tf - idf$ indexing technique, resulting in a total of 3890 documents, 9,782,911 terms and 53,791 unique terms.

Table 1 summarizes the testing parameters and their corresponding ranges. To obtain the candidate set N , for each query sample, we keep the $top - n$ elements using cosine similarity and a log-based $tf - idf$ indexing schema. For the candidate set size $|N|$, our chosen threshold value, 100, is a typical value used in the literature [13]. Our experimental studies are performed in a two-fold

strategy: (i) qualitative analysis in terms of diversification and precision of each employed method with respect to the optimal result set; and (ii) scalability analysis of diversification methods when increasing the query parameters.

Table 1. Parameters tested in the experiments.

Parameter	Range
algorithms tested	MMR, Max-min, Max-sum, Mono-objective, LexRank, Biased LexRank, DivRank, Grasshopper
tradeoff l values	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9
candidate set size $n = N $	100
result set size $k = S $	5, 10, 20, 30
# of sample queries	298

3.2. Evaluation Metrics

As the authors of [14] claim that “there is no evaluation metric that seems to be universally accepted as the best for measuring the performance of algorithms that aim to obtain diverse rankings”, we have chosen to evaluate diversification methods using various metrics employed in TREC Diversity Tasks (<http://trec.nist.gov/data/web10.html>). In particular, we report:

- *a*-nDCG: The *a*-normalized discounted cumulative gain [15] metric quantifies the amount of unique aspects of the query q that are covered by the *top* – k ranked documents. We use $a = 0.5$, as typical in TREC evaluation.
- ERR-IA: Expected reciprocal rank-intent aware [16] is based on inter-dependent ranking. The contribution of each document is based on the relevance of documents ranked above it. The discount function is therefore not just dependent on the rank, but also on the relevance of previously ranked documents.
- S-recall: Subtopic-recall [17] is the number of unique aspects covered by the *top* – k results, divided by the total number of aspects. It measures the aspect coverage for a given result list at depth k .

3.3. Relevance Judgments

Evaluation of diversification requires a data corpus, a set of query topics and a set of relevance judgments, preferably assessed by domain experts for each query. One of the difficulties in evaluating methods designed to introduce diversity in the legal document ranking process is the lack of standard testing data. While TREC added a diversity task to the web track in 2009, this dataset was designed assuming a general web search, and so, it is not possible to adapt it to our setting. Having only the document corpus, we need to define: (a) the query topics; (b) a method to derive the subtopics for each topic; and (c) a method to annotate the corpus for each topic. In the absence of a standard dataset specifically tailored for this purpose, we looked for an objective way to evaluate and assess the performances of various diversification methods on our corpus. We do acknowledge the fact that the process of automatic query generation is at best an imperfect approximation of what a real person would do.

To this end, we have employed the following way to annotate our corpus with relevance judgments for each query:

User profiles/queries: The West American Digest System is a taxonomy of identifying points of law from reported cases and organizing them by topic and key number. It is used to organize the entire body of American law. We used the West Law Digest Topics as candidate user queries. In other words, each topic was issued as a candidate query to our retrieval system. Outlier queries, whether too specific/rare or too general, were removed using the interquartile range, below or above values $Q1$ and $Q3$, sequentially in terms of the number of hits in the result set and the score distribution for the hits,

demanding in parallel a minimum coverage of $\min|N|$ results. In total, we kept 289 queries. Table 2 provides a sample of the topics we further consider as user queries.

Table 2. West Law Digest Topics as user queries.

1: Abandoned and Lost Property	3: Abortion and Birth Control
24: Aliens Immigration and Citizenship	31: Antitrust and Trade Regulation
61: Breach of Marriage Promise	84: Commodity Futures Trading Regulation
88: Compromise and Settlement	199: Implied and Constructive Contracts
291: Privileged Communications and Confidentiality	363: Threats Stalking and Harassment

Query assessments and ground-truth: For each topic/query, we kept the $top - n$ results. An LDA [18] topic model, using an open source implementation (<http://mallet.cs.umass.edu/>), was trained on the $top - n$ results for each query. Topic modeling gives us a way to infer the latent structure behind a collection of documents since each trained topic is a set of keywords with corresponding weights. Table 3 provides a sample of the top keywords for each topic based on the result set for User Query 1: Abandoned and Lost Property .

Table 3. Top key words for topics for documents on the result set for Query 1: Abandoned and Lost Property.

Topic	Top Words
1	court applicant property respondent claim order costs company trustee trust
2	evidence agreement contract business drtime hamiltonrespondent applicant sales
3	tribunal rights land title native evidence area claim interests appellant
4	evidence agreement meeting conduct time tiltformsecond brand septembraustralia
5	evidence professor university dr property gray patent uwaclaimnrc

Based on the resulting topic distribution, with an acceptance threshold of 20%, we access a document as relevant for an topic/aspect. The chosen threshold value, 20%, in the outcome of the LDA method to annotate the corpus for each topic affects the density of our topic/aspect distribution and not the generalizability of the proposed approach or the evaluation outcome. We performed experiments with other settings, i.e., 10% and 30% achieving comparable results. We do acknowledge the fact that if we had chosen a threshold based on interquartile range, as we did in the “user profiles/queries” paragraph, the aforementioned threshold would sound less ad hoc. Nevertheless, this would come at the expense of increasing the complexity both of the paper and the system code.

Table 4 provides a sample of the topic distribution for User Query 1: Abandoned and Lost Property. For the data given in Table 4, based on the relative proportions, with an acceptance threshold of 20%, we can infer that Document Number 08_711 is relevant for Topics 1 and 2, while Document Number 07_924 is relevant for Topics 1 and 5.

Table 4. Topic composition for five random documents on the result set for Query 1: Abandoned and Lost Property.

Document	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5
08_711	0.731	0.267	0.00172	8.79×10^{-5}	4×10^{-5}
09_1395	0.467	2.25×10^{-5}	0.459	0.0746	1.04×10^{-5}
09_1383	0.99	2.85×10^{-5}	0.00944	0.000298	1.31×10^{-5}
06_1169	0.994	5.33×10^{-5}	0.00559	5.41×10^{-5}	2.46×10^{-5}
07_924	0.237	4.83×10^{-5}	4.64×10^{-5}	4.9×10^{-5}	0.763

In other words from the corresponding topic model distribution, we acquire binary assessments for each query and document in the result set. Thus, using LDA, we create our ground-truth data consisting of binary aspect assessments for each query.

We have made available our complete dataset, ground-truth data, queries and relevance assessments in standard TREC format (qrel files), so as to enhance collaboration and contribution with respect to diversification issues in legal IR (<https://github.com/mkoniari/LegalDivEval>).

3.4. Results

As a baseline to compare diversification methods, we consider the simple ranking produced by cosine similarity and log-based *tf-idf* indexing schema. For each query, our initial set N contains the *top-n* query results. The interpolation parameter $\lambda \in [0..1]$ is tuned in 0.1 steps separately for each method. We present the evaluation results for the methods employed, using the aforementioned evaluation metrics, at cut-off values of 5, 10, 20 and 30, as typical in TREC evaluations. Results are presented with fixed parameter $n = |N|$. Note that each of the diversification variations is applied in combination with each of the diversification algorithms and for each user query.

Figure 2 shows the α -normalized discounted cumulative gain (α -nDCG) of each method for different values of λ . Interestingly, web search result diversification methods (MMR, Max-sum, Max-min and Mono-objective) outperformed the baseline ranking, while text summarization methods (LexRank, Biased LexRank and Grasshopper, as it was utilized without a network citation graph) failed to improve the baseline ranking performing lower than the baseline ranking at all levels across all metrics. Graph-based methods' (DivRank) results vary across the different values of λ . We attribute this finding to the extreme sparse network of citations since our dataset covers a short time period (three years).

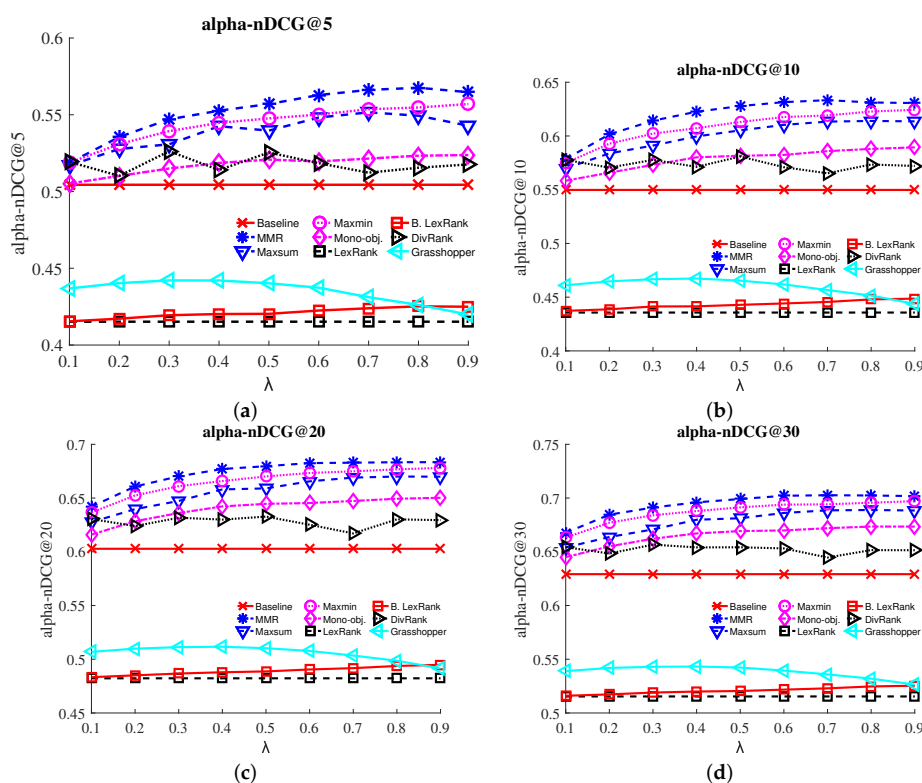


Figure 2. α -Normalized discounted cumulative gain (α -nDCG) at various levels @5, @10, @20, @30 for baseline, MMR, Max-sum, Max-min, Mono-objective, LexRank, Biased LexRank, DivRank and Grasshopper methods. (a) α -nDCG@5; (b) α -nDCG@10; (c) α -nDCG@20; (d) α -nDCG@30. (Best viewed in color.)

The trending behavior of MMR, Max-min and Max-sum is very similar, especially at levels @10, and @20, while at level @5, Max-min and Max-sum presented nearly identical α -nDCG values in many

λ values (e.g., 0.1, 0.2, 0.4, 0.6, 0.7). Finally, MMR constantly achieves better results with respect to the rest of the methods, followed by Max-min and Max-sum. Mono-objective, despite the fact that it performs better than the baseline in all λ values, still always presents a lower performance when compared to MMR, Max-min and Max-sum. It is clear that web search result diversification approaches (MMR, Max-sum, Max-min and Mono-objective) tend to perform better than the selected baseline ranking method. Moreover, as λ increases, preference for diversity, as well as a-nDCG accuracy increase for all tested methods.

Figure 3 depicts the normalized expected reciprocal rank-intent aware (nERR-IA) plots for each method with respect to different values of λ . It is clear that web search result diversification approaches (MMR, Max-sum, Max-min and Mono-objective) tend to perform better than the selected baseline ranking method. Moreover, as λ increases, preference for diversity, as well as nERR-IA accuracy increase for all tested methods. Text summarization methods (LexRank, Biased LexRank) and Grasshopper once again failed to improve the baseline ranking at all levels across all metrics, while as in a-nDCG plots, DivRank results vary across the different values of λ . MMR constantly achieves better results with respect to the rest of the methods. We also observed that Max-min tends to perform better than Max-sum. There were a few cases where both methods presented nearly similar performance, especially in lower recall levels (e.g., for nERR-IA@5 when λ equals 0.1, 0.4, 0.6, 0.7). Once again, Mono-objective presents a lower performance when compared to MMR, Max-min and Max-sum for the nERR-IA metric for all λ values applied.

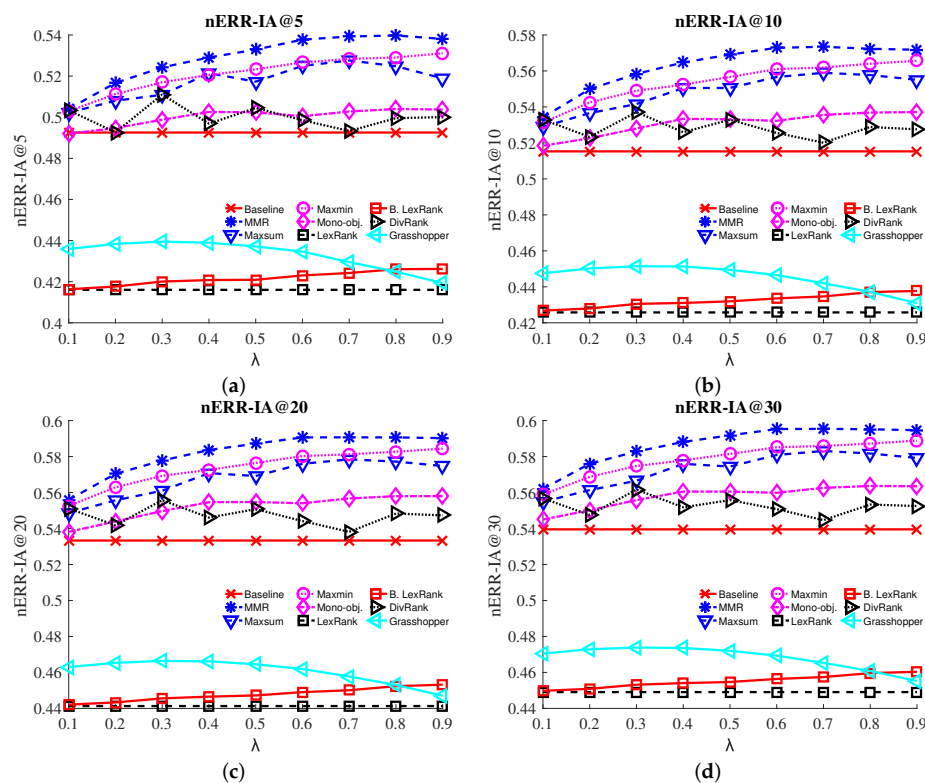


Figure 3. Normalized expected reciprocal rank-intent aware (nERR-IA) at various levels @5, @10, @20, @30 for baseline and the MMR, Max-sum, Max-min, Mono-objective, LexRank, Biased LexRank, DivRank and Grasshopper methods. (a) nERR-IA@5; (b) nERR-IA@10; (c) nERR-IA@20; (d) nERR-IA@30. (Best viewed in color.)

Figure 4 shows the subtopic-recall (S-recall) at various levels @5, @10, @20, @30 of each method for different values of λ . It is clear that the web search result diversification methods (MMR, Max-sum, Max-min and Mono-objective) tend to perform better than the baseline ranking. As λ increases, preference to diversity increases for all methods, except MMR. Subtopic-recall accuracy of all methods,

except MMR, increases when increasing λ . For lower levels (e.g., @5, @10), MMR clearly outperforms other methods, while for upper levels (e.g., @20, @30), MMR and Max-min scores are comparable. We also observe that Max-min tends to perform better than Max-sum, which in turn constantly achieves better results than Mono-objective. Finally, LexRank, Biased LexRank and Grasshopper approaches fail to improve the baseline ranking at all levels across all metrics. Overall, we noticed a similar trending behavior with the ones discussed for Figures 2 and 3.

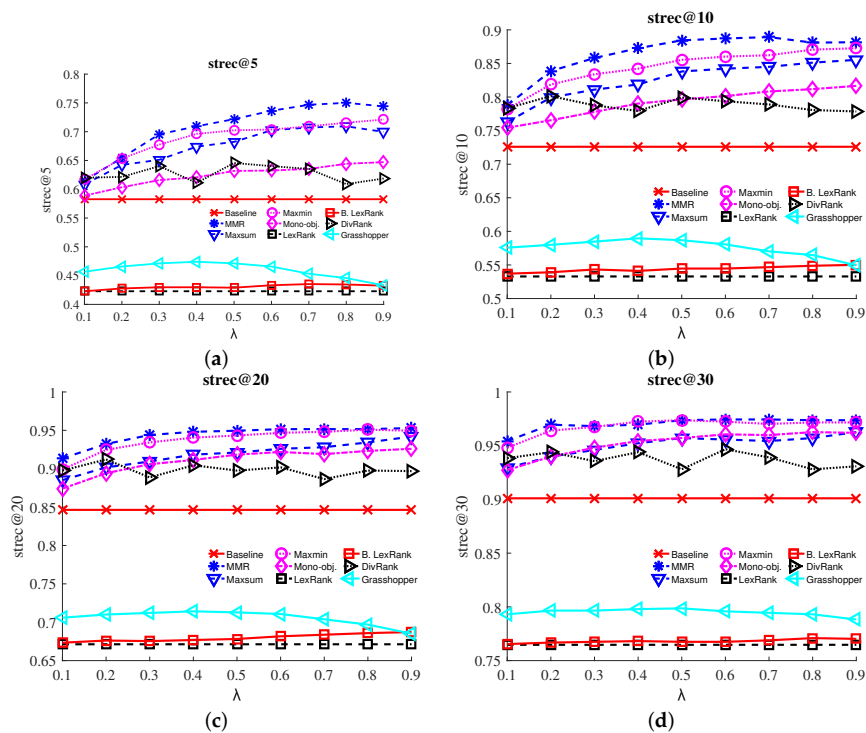


Figure 4. Subtopic recall (S-recall) at various levels @5, @10, @20, @30 for the baseline and the MMR, Max-sum, Max-min, Mono-objective, LexRank, Biased LexRank, DivRank and Grasshopper methods. (a) S-recall@5; (b) S-recall@10; (c) S-recall@20; (d) S-recall@30. (Best viewed in color.)

In summary, among all of the results, we note that the trends in the graphs look very similar. Clearly enough, the utilized web search diversification methods (MMR, Max-sum, Max-min, Mono-objective) and DivRank statistically significantly, using the paired two-sided t -test, outperform the baseline method, offering legislation stakeholders broader insight with respect to their information needs. Furthermore, trends across the evaluation metric graphs highlight balance boundaries for legal IR systems between reinforcing relevant documents or sampling the information space around the legal query.

Table 5 summarizes the average results of the diversification methods. Statistically-significant values, using the paired two-sided t -test with $p_{value} < 0.05$ are denoted with $^{\circ}$ and with $p_{value} < 0.01$ with * .

The effectiveness of diversification methods is also depicted in Table 6, which illustrates the result sets for three example queries, using our case law dataset ($|S| = 30$ and $|N| = 100$) with $\lambda = 0$ (no diversification), $\lambda = 0.1$ (light diversification), $\lambda = 0.5$ (moderate diversification) and $\lambda = 0.9$ (high diversification). Only MMR results are shown since, in almost all variations, it outperforms other approaches. Due to space limitations, we show the case title for each entry. The hyperlink for the full text for each entry can be formed using the format (<http://www.austlii.edu.au/au/cases/cth/FCA/{year}/{docNum}.html>) and substituting the given values in the Table 6. For example, for the first entry of the Table 6 the corresponding hyperlink, substituting values 2008 and 1503 for *year* and *docNum* is (<http://www.austlii.edu.au/au/cases/cth/FCA/2008/1503.html>).

Table 5. Retrieval performance of the tested algorithms with interpolation parameter $\lambda \in [0..1]$ tuned in 0.1 steps for $|N| = 100$ and $k = 30$. The highest scores are shown in bold. Statistically-significant values, using the paired two-sided t -test with $p_{value} < 0.05$ are denoted with $^\circ$ and with $p_{value} < 0.01$ with * .

Method	a-nDCG				nERR-IA				S-recall			
	@5	@10	@20	@30	@5	@10	@20	@30	@5	@10	@20	@30
$\lambda = 0.1$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5187	0.5785 *	0.642 *	0.6676 *	0.5041	0.5341	0.5559 $^\circ$	0.562 $^\circ$	0.6145 $^\circ$	0.7875 *	0.9135 *	0.9543 *
Max-sum	0.5170	0.5699 $^\circ$	0.6276 *	0.6541 *	0.5022	0.5290	0.5486	0.5549	0.6083	0.7626 $^\circ$	0.8851 *	0.9294 *
Max-min	0.5188	0.5749 *	0.6365 *	0.6633 *	0.5029	0.5313	0.5526 $^\circ$	0.5589 $^\circ$	0.6173 $^\circ$	0.7820 *	0.8990 *	0.9481 *
Mono-objective	0.5052	0.5584	0.6160	0.6450	0.4919	0.5184	0.5382	0.5450	0.5889	0.7543	0.8740 $^\circ$	0.9273 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4155 *	0.4373 *	0.4833 *	0.5160 *	0.4163 *	0.4268 *	0.4421 *	0.4498 *	0.4228 *	0.5370 *	0.6734 *	0.7654 *
DivRank	0.5195	0.5774 *	0.6304 *	0.6543 *	0.5035	0.5328	0.5511	0.5567	0.6208 $^\circ$	0.7820 *	0.8976 *	0.9384 *
Grasshopper	0.4368 *	0.4611 *	0.5069 *	0.5389 *	0.4359 *	0.4476 *	0.4630 *	0.4705 *	0.4567 *	0.5758 *	0.7059 *	0.7931 *
$\lambda = 0.2$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5356 *	0.6015 *	0.6607 *	0.6845 *	0.5167 $^\circ$	0.5499 *	0.5704 *	0.576 *	0.6547 *	0.8388 *	0.9322 *	0.9696 *
Max-sum	0.5277 $^\circ$	0.5838 *	0.6397 *	0.6637 *	0.5080	0.5366 $^\circ$	0.5557 $^\circ$	0.5613 $^\circ$	0.6422 *	0.7993 *	0.9017 *	0.9398 *
Max-min	0.5309 *	0.5929 *	0.6524 *	0.6771 *	0.5113	0.5425 *	0.5629 *	0.5687 *	0.6533 *	0.8187 *	0.9246 *	0.9640 *
Mono-objective	0.5102	0.5658	0.6284 *	0.6550 *	0.4947	0.5226	0.5439	0.5502	0.6035	0.7654 *	0.8941 *	0.9398 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4172 *	0.4387 *	0.4850 *	0.5173 *	0.4176 *	0.4280 *	0.4433 *	0.4509 *	0.4277 *	0.5391 *	0.6761 *	0.7668 *
DivRank	0.5077	0.5657	0.6209 $^\circ$	0.6454 $^\circ$	0.4902	0.5196	0.5386	0.5444	0.6159 $^\circ$	0.7931 *	0.9100 *	0.9453 *
Grasshopper	0.4403 *	0.4647 *	0.5099 *	0.5419 *	0.4384 *	0.4502 *	0.4654 *	0.4729 *	0.4657 *	0.5799 *	0.7100 *	0.7965 *
$\lambda = 0.3$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.547 *	0.6142 *	0.6702 *	0.6912 *	0.5242 *	0.5584 *	0.5778 *	0.5828 *	0.6955 *	0.8581 *	0.9439 *	0.9682 *
Max-sum	0.5308 *	0.5911 *	0.6473 *	0.6708 *	0.5109	0.5416 *	0.5610 *	0.5666 *	0.6512 *	0.8111 *	0.9093 *	0.9460 *
Max-min	0.5394 *	0.6022 *	0.6610 *	0.6840 *	0.5170 $^\circ$	0.5490 *	0.5693 *	0.5748 *	0.6775 *	0.8339 *	0.9343 *	0.9675 *
Mono-objective	0.5150	0.5731 $^\circ$	0.6361 *	0.6621 *	0.4988	0.5280	0.5497	0.5558	0.6159 $^\circ$	0.7779 *	0.9059 *	0.9481 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4194 *	0.4414 *	0.4867 *	0.5190 *	0.4201 *	0.4305 *	0.4456 *	0.4532 *	0.4298 *	0.5433 *	0.6754 *	0.7675 *
DivRank	0.5261 $^\circ$	0.5779 *	0.6316 *	0.6566 *	0.5109	0.5371 $^\circ$	0.5557 $^\circ$	0.5616 $^\circ$	0.6394 *	0.7882 *	0.8886 *	0.9356 *
Grasshopper	0.4421 *	0.4667 *	0.5113 *	0.5430 *	0.4395 *	0.4514 *	0.4664 *	0.4738 *	0.4713 *	0.5848 *	0.7121 *	0.7965 *

Table 5. Cont.

	a-nDCG				nERR-IA				S-recall			
Method	@5	@10	@20	@30	@5	@10	@20	@30	@5	@10	@20	@30
$\lambda = 0.4$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5527 *	0.6226 *	0.677 *	0.696 *	0.5291 *	0.5647 *	0.5836 *	0.5881 *	0.7093 *	0.8727 *	0.9481 *	0.9696 *
Max-sum	0.5425 *	0.5996 *	0.6580 *	0.6798 *	0.5214 *	0.5505 *	0.5708 *	0.5759 *	0.6740 *	0.8187 *	0.9183 *	0.9522 *
Max-min	0.5447 *	0.6073 *	0.6659 *	0.6879 *	0.5206 *	0.5524 *	0.5726 *	0.5778 *	0.6962 *	0.8422 *	0.9405 *	0.9723 *
Mono-objective	0.5186	0.5802 *	0.6422 *	0.6671 *	0.5025	0.5334	0.5546 °	0.5605 °	0.6208 °	0.7903 *	0.9114 *	0.9543 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4202 *	0.4415 *	0.4879 *	0.5199 *	0.4208 *	0.4310 *	0.4465 *	0.4541 *	0.4298 *	0.5412 *	0.6768 *	0.7682 *
DivRank	0.5140	0.5710 °	0.6298 *	0.6540 *	0.4968	0.5258	0.5460	0.5517	0.6111	0.7785 *	0.9045 *	0.9439 *
Grasshopper	0.4421 *	0.4673 *	0.5117 *	0.5432 *	0.4389 *	0.4513 *	0.4662 *	0.4736 *	0.4740 *	0.5896 *	0.7142 *	0.7979 *
$\lambda = 0.5$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.557 *	0.6278 *	0.6796 *	0.6991 *	0.5329 *	0.5691 *	0.5872 *	0.5918 *	0.7218 *	0.8844 *	0.9495 *	0.9737 *
Max-sum	0.5397 *	0.6052 *	0.6590 *	0.6812 *	0.5173 °	0.5506 *	0.5692 *	0.5744 *	0.6824 *	0.8381 *	0.9211 *	0.9571 *
Max-min	0.5477 *	0.6130 *	0.6701 *	0.6913 *	0.5233 *	0.5567 *	0.5764 *	0.5814 *	0.7024 *	0.8554 *	0.9433 *	0.9737 *
Mono-objective	0.5208	0.5816 *	0.6446 *	0.6693 *	0.5024	0.5331	0.5547 °	0.5605 °	0.6318 *	0.7965 *	0.9183 *	0.9571 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4203 *	0.4429 *	0.4887 *	0.5204 *	0.4209 *	0.4319 *	0.4472 *	0.4547 *	0.4291 *	0.5446 *	0.6782 *	0.7675 *
DivRank	0.5252 °	0.5810 *	0.6327 *	0.6542 *	0.5044	0.5329	0.5507	0.5558	0.6450 *	0.8000 *	0.8976 *	0.9280 *
Grasshopper	0.4402 *	0.4654 *	0.5103 *	0.5423 *	0.4371 *	0.4494 *	0.4645 *	0.4720 *	0.4713 *	0.5869 *	0.7128 *	0.7986 *
$\lambda = 0.6$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5628 *	0.6315 *	0.6825 *	0.7021 *	0.5377 *	0.5729 *	0.5906 *	0.5953 *	0.7363 *	0.8872 *	0.9516 *	0.9744 *
Max-sum	0.5482 *	0.6103 *	0.6656 *	0.6861 *	0.5250 *	0.5566 *	0.5760 *	0.5809 *	0.7024 *	0.8422 *	0.9260 *	0.9557 *
Max-min	0.5501 *	0.6176 *	0.6733 *	0.6939 *	0.5267 *	0.5610 *	0.5803 *	0.5852 *	0.7038 *	0.8602 *	0.9467 *	0.9723 *
Mono-objective	0.5196	0.5824 *	0.6454 *	0.6699 *	0.5005	0.5323	0.5540 °	0.5598 °	0.6325 *	0.8014 *	0.9218 *	0.9606 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4225 *	0.4442 *	0.4905 *	0.5217 *	0.4230 *	0.4336 *	0.4490 *	0.4564 *	0.4332 *	0.5446 *	0.6817 *	0.7675 *
DivRank	0.5185	0.5711 °	0.6253 *	0.6532 *	0.4986	0.5256	0.5442	0.5508	0.6401 *	0.7945 *	0.9017 *	0.9467 *
Grasshopper	0.4374 *	0.4619 *	0.5077 *	0.5394 *	0.4346 *	0.4465 *	0.4619 *	0.4693 *	0.4657 *	0.5806 *	0.7107 *	0.7958 *

Table 5. Cont.

Method	a-nDCG				nERR-IA				S-recall			
	@5	@10	@20	@30	@5	@10	@20	@30	@5	@10	@20	@30
$\lambda = 0.7$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5662 *	0.6333 *	0.6829 *	0.7026 *	0.5393 *	0.5734 *	0.5907 *	0.5954 *	0.7467 *	0.8893 *	0.9516 *	0.9744 *
Max-sum	0.5516 *	0.6134 *	0.6690 *	0.6886 *	0.5276 *	0.5590 *	0.5784 *	0.5831 *	0.7073 *	0.8450 *	0.9280 *	0.9543 *
Max-min	0.5536 *	0.6191 *	0.6749 *	0.6941 *	0.5283 *	0.5619 *	0.5812 *	0.5858 *	0.7093 *	0.8623 *	0.9481 *	0.9702 *
Mono-objective	0.5215	0.5859 *	0.6473 *	0.6720 *	0.5028	0.5355 °	0.5567 °	0.5626 °	0.6353 *	0.8083 *	0.9190 *	0.9599 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4239 *	0.4456 *	0.4917 *	0.5229 *	0.4242 *	0.4347 *	0.4501 *	0.4574 *	0.4353 *	0.5467 *	0.6837 *	0.7689 *
DivRank	0.5123	0.5654	0.6170	0.6446	0.4933	0.5202	0.5381	0.5446	0.6353 *	0.7896 *	0.8865 *	0.9391 *
Grasshopper	0.4312 *	0.4566 *	0.5034 *	0.5359 *	0.4295 *	0.4420 *	0.4577 *	0.4653 *	0.4533 *	0.5702 *	0.7038 *	0.7945 *
$\lambda = 0.8$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5676 *	0.6312 *	0.6834 *	0.7024 *	0.5397 *	0.5723 *	0.5906 *	0.5952 *	0.7502 *	0.881 *	0.9516 *	0.9737 *
Max-sum	0.5494 *	0.6140 *	0.6700 *	0.6889 *	0.5248 *	0.5577 *	0.5772 *	0.5817 *	0.7093 *	0.8512 *	0.9343 *	0.9571 *
Max-min	0.5547 *	0.6228 *	0.6767 *	0.6957 *	0.5291 *	0.5640 *	0.5827 *	0.5872 *	0.7156 *	0.8706 *	0.9509 *	0.9716 *
Mono-objective	0.5234	0.5880 *	0.6493 *	0.6735 *	0.5040	0.5369 °	0.5580 *	0.5636 *	0.6443 *	0.8118 *	0.9232 *	0.9626 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4252 *	0.4479 *	0.4937 *	0.5247 *	0.4261 *	0.4371 *	0.4524 *	0.4596 *	0.4346 *	0.5488 *	0.6858 *	0.7709 *
DivRank	0.5155	0.5735 *	0.6304 *	0.6517 *	0.4995	0.5289	0.5484	0.5534	0.6090	0.7806 *	0.8976 *	0.9280 *
Grasshopper	0.4262 *	0.4515 *	0.4986 *	0.5320 *	0.4249 *	0.4372 *	0.4530 *	0.4608 *	0.4457 *	0.5647 *	0.6969 *	0.7931 *
$\lambda = 0.9$												
baseline	0.5044	0.5498	0.6028	0.6292	0.4925	0.5153	0.5333	0.5395	0.5827	0.7260	0.8464	0.9010
MMR	0.5647 *	0.6306 *	0.6834 *	0.7018 *	0.5381 *	0.5718 *	0.5902 *	0.5946 *	0.7439 *	0.8817 *	0.9529 *	0.9737 *
Max-sum	0.5429 *	0.6136 *	0.6699 *	0.6884 *	0.5188 *	0.5551 *	0.5748 *	0.5792 *	0.6997 *	0.8554 *	0.9419 *	0.9626 *
Max-min	0.5570 *	0.6244 *	0.6781 *	0.6971 *	0.5311 *	0.5657 *	0.5844 *	0.5889 *	0.7211 *	0.8727 *	0.9495 *	0.9716 *
Mono-objective	0.5238	0.5894 *	0.6502 *	0.6734 *	0.5037	0.5371 °	0.5580 *	0.5635 *	0.6471 *	0.8166 *	0.9260 *	0.9619 *
LexRank	0.4152 *	0.4357 *	0.4823 *	0.5154 *	0.4160 *	0.4258 *	0.4413 *	0.4491 *	0.4228 *	0.5329 *	0.6713 *	0.7647 *
Biased LexRank	0.4250 *	0.4486 *	0.4948 *	0.5254 *	0.4262 *	0.4377 *	0.4532 *	0.4604 *	0.4325 *	0.5502 *	0.6872 *	0.7702 *
DivRank	0.5177	0.5721 °	0.6295 *	0.6511 *	0.5001	0.5275	0.5473	0.5524	0.6187 °	0.7785 *	0.8969 *	0.9280 *
Grasshopper	0.4199 *	0.4438 *	0.4915 *	0.5264 *	0.4193 *	0.4309 *	0.4470 *	0.4552 *	0.4332 *	0.5495 *	0.6851 *	0.7882 *

Table 6. Result sets (document titles for three example queries, using the dataset ($|S| = 30$ and $|N| = 100$) with $\lambda = 0$ (no diversification), $\lambda = 0.1$ (light diversification), $\lambda = 0.5$ (moderate diversification) and $\lambda = 0.9$ (high diversification)).

	Baseline ($\lambda = 0$)	MMR: Light Diversity ($\lambda = 0.1$)	MMR: Moderate Diversity ($\lambda = 0.5$)	MMR: High Diversity ($\lambda = 0.9$)
<i>Query 24: Aliens Immigration and Citizenship</i>				
1	Virgin Holdings SA v Commissioner of Taxation [2008] FCA 1503 (10 October 2008)	Virgin Holdings SA v Commissioner of Taxation [2008] FCA 1503 (10 October 2008)	Virgin Holdings SA v Commissioner of Taxation [2008] FCA 1503 (10 October 2008)	Virgin Holdings SA v Commissioner of Taxation [2008] FCA 1503 (10 October 2008)
2	Undershaft (No 1) Limited v Commissioner of Taxation [2009] FCA 41 (3 February 2009)	Fowler v Commissioner of Taxation [2008] FCA 528 (21 April 2008)	Fowler v Commissioner of Taxation [2008] FCA 528 (21 April 2008)	Soh v Commonwealth of Australia [2008] FCA 520 (18 April 2008)
3	Fowler v Commissioner of Taxation [2008] FCA 528 (21 April 2008)	Wight v Honourable Chris Pearce, MP, Parliamentary Secretary to the Treasurer [2007] FCA 26 (29 January 2007)	Coleman v Minister for Immigration & Citizenship [2007] FCA 1500 (27 September 2007)	SZJDI v Minister for Immigration & Citizenship (No. 2) [2008] FCA 813 (16 May 2008)
4	Wight v Honourable Chris Pearce, MP, Parliamentary Secretary to the Treasurer [2007] FCA 26 (29 January 2007)	Undershaft (No 1) Limited v Commissioner of Taxation [2009] FCA 41 (3 February 2009)	Charlie v Minister for Immigration and Citizenship [2008] FCA 1025 (10 July 2008)	Charlie v Minister for Immigration and Citizenship [2008] FCA 1025 (10 July 2008)
5	Coleman v Minister for Immigration & Citizenship [2007] FCA 1500 (27 September 2007)	Coleman v Minister for Immigration & Citizenship [2007] FCA 1500 (27 September 2007)	VSAB v Minister for Immigration and Multicultural and Indigenous Affairs [2006] FCA 239 (17 March 2006)	VSAB v Minister for Immigration and Multicultural and Indigenous Affairs [2006] FCA 239 (17 March 2006)
<i>Query 84: Commodity Futures Trading Regulation</i>				
1	BHP Billiton Iron Ore Pty Ltd. v The National Competition Council [2006] FCA 1764 (18 December 2006)	BHP Billiton Iron Ore Pty Ltd. v The National Competition Council [2006] FCA 1764 (18 December 2006)	BHP Billiton Iron Ore Pty Ltd. v The National Competition Council [2006] FCA 1764 (18 December 2006)	BHP Billiton Iron Ore Pty Ltd. v The National Competition Council [2006] FCA 1764 (18 December 2006)
2	Australian Securities & Investments Commission v Lee [2007] FCA 918 (15 June 2007)	Australian Securities & Investments Commission v Lee [2007] FCA 918 (15 June 2007)	Australian Securities & Investments Commission v Lee [2007] FCA 918 (15 June 2007)	Australian Competition and Consumer Commission v Dally M Publishing and Research Pty Limited [2007] FCA 1220 (10 August 2007)

Table 6. Cont.

	Baseline ($\lambda = 0$)	MMR: Light Diversity ($\lambda = 0.1$)	MMR: Moderate Diversity ($\lambda = 0.5$)	MMR: High Diversity ($\lambda = 0.9$)
3	Woodside Energy Ltd. (ABN 63 005 482 986) v Commissioner of Taxation (No 2) [2007] FCA 1961 (10 December 2007)	Woodside Energy Ltd. (ABN 63 005 482 986) v Commissioner of Taxation (No 2) [2007] FCA 1961 (10 December 2007)	Woodside Energy Ltd. (ABN 63 005 482 986) v Commissioner of Taxation (No 2) [2007] FCA 1961 (10 December 2007)	Heritage Clothing Pty Ltd. trading as Peter Jackson Australia v Mens Suit Warehouse Direct Pty Ltd. trading as Walter Withers [2008] FCA 1775 (28 November 2008)
4	BHP Billiton Iron Ore Pty Ltd. v National Competition Council (No 2) [2007] FCA 557 (19 April 2007)	Keynes v Rural Directions Pty Ltd. (No 2) (includes Corrigendum dated 16 July 2009) [2009] FCA 567 (3 June 2009)	Keynes v Rural Directions Pty Ltd. (No 2) (includes Corrigendum dated 16 July 2009) [2009] FCA 567 (3 June 2009)	Travellex Limited v Commissioner of Taxation (Corrigendum dated 4 February 2009) [2008] FCA 1961 (19 December 2008)
5	Keynes v Rural Directions Pty Ltd. (No 2) (includes Corrigendum dated 16 July 2009) [2009] FCA 567 (3 June 2009)	Queanbeyan City Council v ACTEW Corporation Limited [2009] FCA 943 (24 August 2009)	Heritage Clothing Pty Ltd. trading as Peter Jackson Australia v Mens Suit Warehouse Direct Pty Ltd trading as Walter Withers [2008] FCA 1775 (28 November 2008)	Ashwick (Qld) No 127 Pty Ltd. (ACN 010 577 456) v Commissioner of Taxation [2009] FCA 1388 (26 November 2009)
<i>Query 291: Privileged Communications and Confidentiality</i>				
1	Siam Polyethylene Co Ltd. v Minister of State for Home Affairs (No 3) [2009] FCA 839 (7 August 2009)	Siam Polyethylene Co Ltd. v Minister of State for Home Affairs (No 3) [2009] FCA 839 (7 August 2009)	Siam Polyethylene Co Ltd. v Minister of State for Home Affairs (No 3) [2009] FCA 839 (7 August 2009)	Siam Polyethylene Co Ltd. v Minister of State for Home Affairs (No 3) [2009] FCA 839 (7 August 2009)
2	AWB Limited v Australian Securities and Investments Commission [2008] FCA 1877 (11 December 2008)	AWB Limited v Australian Securities and Investments Commission [2008] FCA 1877 (11 December 2008)	AWB Limited v Australian Securities and Investments Commission [2008] FCA 1877 (11 December 2008)	Krueger Transport Equipment Pty Ltd. v Glen Cameron Storage [2008] FCA 803 (30 May 2008)
3	Brookfield Multiplex Limited v International Litigation Funding Partners Pte Ltd. (No 2) [2009] FCA 449 (6 May 2009)	Brookfield Multiplex Limited v International Litigation Funding Partners Pte Ltd (No 2) [2009] FCA 449 (6 May 2009)	Autodata Limited v Boyce's Automotive Data Pty Limited [2007] FCA 1517 (4 October 2007)	Futuretronics.com.au Pty Limited v Graphix Labels Pty Ltd. [2007] FCA 1621 (29 October 2007)
4	Cadbury Schweppes Pty Ltd. (ACN 004 551 473) v Amcor Limited (ACN 000 017 372) [2008] FCA 88 (19 February 2008)	Barrett Property Group Pty Ltd. v Carlisle Homes Pty Ltd (No 2) [2008] FCA 930 (17 June 2008)	Barrett Property Group Pty Ltd. v Carlisle Homes Pty Ltd (No 2) [2008] FCA 930 (17 June 2008)	Australian Competition & Consumer Commission v Visy Industries [2006] FCA 136 (23 February 2006)
5	Barrett Property Group Pty Ltd v Carlisle Homes Pty Ltd. (No 2) [2008] FCA 930 (17 June 2008)	Cadbury Schweppes Pty Ltd. (ACN 004 551 473) v Amcor Limited (ACN 000 017 372) [2008] FCA 88 (19 February 2008)	Optus Networks Ltd. v Telstra Corporation Ltd (No. 2) (includes Corrigendum dated 7 July 2009) [2009] FCA 422 (9 July 2009)	IO Group Inc v Prestige Club Australasia Pty Ltd. (No 2) [2008] FCA 1237 (11 August 2008)

When $\lambda = 0$, the result set contains the top-five elements of S ranked with the similarity scoring function. The result sets with no diversification contain several almost duplicate elements, defined by terms in the case title. As λ increases, less “duplicates” are found in the result set, and the elements in the result set “cover” many more subjects again, as defined by terms in the case title. We note that the result set with high diversification contains elements that have almost all of the query terms, as well as other terms indicating that the case is related to different subjects among the other cases in the result set.

4. Related Work

In this section, we first present related work on query result diversification, afterwards on diversified ranking on graphs and then on legal text retrieval techniques.

4.1. Query Result Diversification

Result diversification approaches have been proposed as a means to tackle ambiguity and redundancy, in various problems and settings, e.g., diversifying historical archives [19], diversifying user comments on news articles [20], diversifying microblog posts [21,22], diversifying image retrieval results [23], diversifying recommendations [24], utilizing a plethora of algorithms and approaches, e.g., learning algorithms [25], approximation algorithms [26], page rank variants [27] and conditional probabilities [28].

Users of (web) search engines typically employ keyword-based queries to express their information needs. These queries are often underspecified or ambiguous to some extent [29]. Different users who pose exactly the same query may have very different query intents. Simultaneously, the documents retrieved by an IR system may reflect superfluous information. Search result diversification aims to solve this problem, by returning diverse results that can fulfill as many different information needs as possible. Published literature on search result diversification is reviewed in [30,31].

The maximal marginal relevance criterion (MMR), presented in [3], is one of the earliest works on diversification and aims at maximizing relevance while minimizing similarity to higher ranked documents. Search results are re-ranked as the combination of two metrics, one measuring the similarity among documents and the other the similarity between documents and the query. In [4], a set of diversification axioms is introduced, and it is proven that it is not possible for a diversification algorithm to satisfy all of them. Additionally, since there is no single objective function suitable for every application domain, the authors propose three diversification objectives, which we adopt in our work. These objectives differ in the level where the diversity is calculated, e.g., whether it is calculated per separate document or on the average of the currently-selected documents.

In another approach, researchers utilized explicit knowledge so as to diversify search results. In [32], the authors proposed a diversification framework, where the different aspects of a given query are represented in terms of sub-queries, and documents are ranked based on their relevance to each sub-query, while in [33] the authors proposed a diversification objective that tries to maximize the likelihood of finding a relevant document in the *top-k* positions given the categorical information of the queries and documents. Finally, the work described in [34] organizes user intents in a hierarchical structure and proposes a diversification framework to explicitly leverage the hierarchical intent.

The key difference between these works and the ones utilized in this paper is that we do not rely on external knowledge, e.g., taxonomy, query logs to generate diverse results. Queries are rarely known in advance, thus probabilistic methods to compute external information are not only expensive to compute, but also have a specialized domain of applicability. Instead, we evaluate methods that rely only on implicit knowledge of the legal corpus utilized and on computed values, using similarity (relevance) and diversity functions in the data domain.

4.2. Diversified Ranking on Graphs

Many network-based ranking approaches have been proposed to rank objects according to different criteria [35], and recently, diversification of the results has attracted attention. Research is currently focused on two directions: a greedy vertex selection procedure and a vertex-reinforced random walk. The greedy vertex selection procedure, at each iteration, selects and removes from the graph the vertex with maximum random walk-based ranking score. One of the earlier algorithms that addresses diversified ranking on graphs by vertex selection with absorbing random walks is Grasshopper [8]. A diversity-focused ranking methodology, based on reinforced random walks, was introduced in [7]. Their proposed model, DivRank, incorporates the rich-gets-richer mechanism to PageRank [11] with reinforcements on transition probabilities between vertices. We utilize these approaches in our diversification framework considering the connectivity matrix of the citation network between documents that are relevant for a given user query.

4.3. Legal Text Retrieval

With respect to legal text retrieval that traditionally relies on external knowledge sources, such as thesauri and classification schemes, various techniques are presented in [36]. Several supervised learning methods have been proposed to classify sources of law according to legal concepts [37–39]. Ontologies and thesauri have been employed to facilitate information retrieval [40–43] or to enable the interchange of knowledge between existing legal knowledge systems [44]. Legal document summarization [45–47] has been used as a way to make the content of the legal documents, notably cases, more easily accessible. We also utilize state of the art summarization algorithms, but under a different objective: we aim to maximize the diversity of the result set for a given query.

Finally, a similar approach with our work is described in [48], where the authors utilize information retrieval approaches to determine which sections within a bill tend to be outliers. However, our work differs in the sense that we maximize the diversify of the result set, rather than detect section outliers within a specific bill.

In another line of work, citation analysis has been used in the field of law to construct case law citation networks [49]. Case documents usually cite previous cases, which in turn may have cited other cases, and thus, a network is formed over time with these citations between cases. Case law citation networks contain valuable information, capable of measuring legal authority [50], identifying authoritative precedent (the legal norm inherited from English common law that encourages judges to follow precedent by letting the past decision stand)[51], evaluating the relevance of court decisions [52] or even assisting with summarizing legal cases [53], thus showing the effectiveness of citation analysis in the case law domain. While the American legal system has been the one that has undergone the widest series of studies in this direction, recently, various researchers applied network analysis in the civil law domain, as well. The authors of [54] propose a network-based approach to model the law. Network analysis techniques were also employed in [55], demonstrating an online toolkit allowing legal scholars to apply network analysis and visual techniques to the entire corpus of EU case law. In this work, we also utilize citation analysis techniques and construct the legislation network, as to cover a wide range of possible aspects of a query.

This paper is an extended version of our paper [2], where we originally introduced the concept of diversifying legal information retrieval results with the goal of evaluating the potential of results diversification in the field of legal information retrieval. We adopted various methods from the literature that were introduced for search result diversification (MMR [3], Max-sum [4], Max-min [4] and Mono-objective [4]) and evaluated the performance of the above methods on a legal corpus objectively annotated with relevance judgments. In our work [56], we address result diversification in the legal IR. To this end, we analyze the impact of various features in computing the query-document relevance and document-document similarity scores and introduce legal domain-specific diversification criteria.

In this work, extending our previous work presented in [2], we extended the methods utilized by incorporating methods introduced for text summarization (LexRank [5] and Biased LexRank [6]) and graph-based ranking (DivRank [7] and Grasshopper [8]) in parallel with the search result diversification methods utilized in our previous work (MMR, Max-sum, Max-min and Mono-objective). We adopted the text summarization and graph-based ranking methods in our diversification schema, and additionally, we utilized various features of our legal dataset.

In detail, text summarization methods (LexRank and Biased LexRank) were originally proposed for computing the relative importance of textual units within a document for assisting summarization tasks. They utilize a document as a graph with vertices consisting of its sentences and edges formed among the vertices based on the sentence textual similarity. To utilize such an approach in our diversification scenario, we: (i) introduce a query-dependent prior, since in our approach, we utilize documents that are in the initial retrieval set N for a given query; and (ii) we build a graph using the document similarity on the result set. If we consider documents as nodes, the result set document collection can be modeled as a graph by generating links between documents based on their similarity score. Analogously to adopting graph-based ranking methods (DivRank and Grasshopper) in our setting, we: (i) introduce a query-dependent prior, since these methods were originally proposed in a query independent context, thus not directly applicable to the diversification of search results; and (ii) we use documents that are in the initial retrieval set N for a given query q and create the citation network between those documents based on the features of the legal dataset we use (prior citation analysis).

Furthermore, we complemented the presentation of the diversification methods utilized with an algorithmic presentation, alongside a textual explanation of each diversification objective, evaluated all of the methods using a real dataset from the common law domain and performed an exhaustive accuracy analysis, fine-tuning separately for each method, trade-off values between finding relevant to the user query documents and diverse documents in the result set. Additionally, we identified statistically-significant values in the retrieval performance of the tested algorithms using the paired two-sided t -test.

5. Conclusions

In this paper, we studied the problem of diversifying results in legal documents. We adopted and compared the performance of several state of the art methods from the web search, network analysis and text summarization domains so as to handle the problems challenges. We evaluated all of the methods using a real dataset from the common law domain that we objectively annotated with relevance judgments for this purpose. Our findings reveal that diversification methods offer notable improvements and enrich search results around the legal query space. In parallel, we demonstrated that that web search diversification techniques outperform other approaches, e.g., summarization-based, graph-based methods, in the context of legal diversification. Finally, we provide valuable insights for legislation stakeholders though diversification, as well as by offering balance boundaries between reinforcing relevant documents or information space sampling around legal queries.

A challenge we faced in this work was the lack of ground-truth. We hope for an increase of the size of truth-labeled dataset in the future, which would enable us to draw further conclusions about the diversification techniques. To this end, our complete dataset is publicly available in an open and editable format, along with ground-truth data, queries and relevance assessments.

In future work, we plan to further study the interaction of relevance and redundancy, in historical legal queries. While access to legislation generally retrieves the current legislation on a topic, point-in-time legislation systems address a different problem, namely that lawyers, judges and anyone else considering the legal implications of past events need to know what the legislation stated at some point in the past when a transaction or events occurred that have led to a dispute and perhaps to litigation [57].

Author Contributions: Marios Koniaris conceived of the idea, designed and performed the experiments, analyzed the results, drafted the initial manuscript and revised the manuscript. Ioannis Anagnostopoulos analyzed the results, helped to draft the initial manuscript and revised the final version. Yannis Vassiliou provided feedback and revised the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Alces, K.A. Legal diversification. *Columbia Law Rev.* **2013**, *113*, 1977–2038.
2. Koniaris, M.; Anagnostopoulos, I.; Vassiliou, Y. Diversifying the Legal Order. In Proceedings of the IFIP International Conference on Artificial Intelligence Applications and Innovations, Thessaloniki, Greece, 16–18 September 2016; pp. 499–509.
3. Carbonell, J.; Goldstein, J. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval, Melbourne, Australia, 24–28 August 1988; pp. 335–336.
4. Gollapudi, S.; Sharma, A. An Axiomatic Approach for Result Diversification. In Proceedings of the 18th International Conference on World Wide Web, Madrid, Spain, 20–24 April 2009; pp. 381–390.
5. Erkan, G.; Radev, D.R. LexRank: Graph-based lexical centrality as salience in text summarization. *J. Artif. Intell. Res.* **2004**, *22*, 457–479.
6. Otterbacher, J.; Erkan, G.; Radev, D.R. Biased LexRank: Passage retrieval using random walks with question-based priors. *Inf. Process. Manag.* **2009**, *45*, 42–54.
7. Mei, Q.; Guo, J.; Radev, D. Divrank: The interplay of prestige and diversity in information networks. In Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, 25–28 July 2010; pp. 1009–1018.
8. Zhu, X.; Goldberg, A.B.; Van Gael, J.; Andrzejewski, D. Improving Diversity in Ranking using Absorbing Random Walks. In Proceedings of the Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT), 2007, Rochester, NY, USA, 22–27 April 2007; pp. 97–104.
9. Hand, D.J.; Mannila, H.; Smyth, P. *Principles of Data Mining*; MIT Press: Cambridge, MA, USA, 2001.
10. Wong, S.M.; Raghavan, V.V. Vector space model of information retrieval: A reevaluation. In Proceedings of the 7th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Cambridge, UK, 2–6 July 1984; pp. 167–185.
11. Page, L.; Brin, S.; Motwani, R.; Winograd, T. The PageRank Citation Ranking: Bringing Order to the Web; Technical Report for Stanford InfoLab: Stanford University, Stanford, CA, USA, 1999.
12. Galgani, F.; Compton, P.; Hoffmann, A. Combining different summarization techniques for legal text. In Proceedings of the Workshop on Innovative Hybrid Approaches to the Processing of Textual Data, Avignon, France, 22 April 2012; pp. 115–123.
13. Sanderson, M. Test Collection Based Evaluation of Information Retrieval Systems. *Found. Trends® Inf. Retr.* **2010**, *4*, 247–375.
14. Radlinski, F.; Bennett, P.N.; Carterette, B.; Joachims, T. Redundancy, diversity and interdependent document relevance. In *ACM SIGIR Forum*; Association for Computing Machinery (ACM): New York, NY, USA, 2009; Volume 43, pp. 46–52.
15. Clarke, C.L.A.; Kolla, M.; Cormack, G.V.; Vechtomova, O.; Ashkan, A.; Büttcher, S.; MacKinnon, I. Novelty and diversity in information retrieval evaluation. In Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Singapore, 20–24 July 2008; pp. 659–666.
16. Chapelle, O.; Metlzer, D.; Zhang, Y.; Grinspan, P. Expected reciprocal rank for graded relevance. In Proceedings of the 18th ACM conference on Information and Knowledge Management—CIKM '09, Hong Kong, China, 2–6 November 2009; pp. 621–630.
17. Zhai, C.X.; Cohen, W.W.; Lafferty, J. Beyond independent relevance. In Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval, Toronto, ON, Canada, 28 July–1 August 2003; pp. 10–17.
18. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent dirichlet allocation. *J. Mach. Learn. Res.* **2003**, *3*, 993–1022.

19. Singh, J.; Nejd, W.; Anand, A. History by Diversity: Helping Historians Search News Archives. In Proceedings of the 2016 ACM on Conference on Human Information Interaction and Retrieval, Carrboro, NC, USA, 13–17 March 2016; pp. 183–192.
20. Giannopoulos, G.; Koniaris, M.; Weber, I.; Jaimes, A.; Sellis, T. Algorithms and criteria for diversification of news article comments. *J. Intell. Inf. Syst.* **2015**, *44*, 1–47.
21. Cheng, S.; Arvanitis, A.; Chrobak, M.; Hristidis, V. Multi-Query Diversification in Microblogging Posts. In Proceedings of the 17th International Conference on Extending Database Technology (EDBT), Athens, Greece, 24–28 March 2014; pp. 133–144.
22. Koniaris, M.; Giannopoulos, G.; Sellis, T.; Vasileiou, Y. Diversifying microblog posts. In Proceedings of the International Conference on Web Information Systems Engineering, Thessaloniki, Greece, 12–14 October 2014; pp. 189–198.
23. Song, K.; Tian, Y.; Gao, W.; Huang, T. Diversifying the image retrieval results. In Proceedings of the 14th ACM International Conference on Multimedia, Santa Barbara, CA, USA, 23–27 October 2006; pp. 707–710.
24. Ziegler, C.N.; McNee, S.M.; Konstan, J.A.; Lausen, G. Improving recommendation lists through topic diversification. In Proceedings of the 14th International Conference on World Wide Web, Chiba, Japan, 10–14 May 2005; pp. 22–32.
25. Raman, K.; Shivaswamy, P.; Joachims, T. Online learning to diversify from implicit feedback. In Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Beijing, China, 12–16 August 2012; pp. 705–713.
26. Makris, C.; Plegas, Y.; Stamatiou, Y.C.; Stavropoulos, E.C.; Tsakalidis, A.K. Reducing Redundant Information in Search Results Employing Approximation Algorithms. In Proceedings of the International Conference on Database and Expert Systems Applications, Munich, Germany, 1–4 September 2014; pp. 240–247.
27. Zhang, B.; Li, H.; Liu, Y.; Ji, L.; Xi, W.; Fan, W.; Chen, Z.; Ma, W.Y. Improving web search results using affinity graph. In Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Salvador, Brazil, 15–19 August 2005; pp. 504–511.
28. Chen, H.; Karger, D.R. Less is more: probabilistic models for retrieving fewer relevant documents. In Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Seattle, WA, USA, 6–10 August 2006; pp. 429–436.
29. Cronen-Townsend, S.; Croft, W.B. Quantifying Query Ambiguity. In Proceedings of the Second International Conference on Human Language Technology Research, San Diego, CA, USA, 24–27 March 2002; pp. 104–109.
30. Santos, R.L.T.; Macdonald, C.; Ounis, I. Search Result Diversification. *Found. Trends® Inf. Retr.* **2015**, *9*, 1–90.
31. Drosou, M.; Pitoura, E. Search result diversification. *ACM SIGMOD Rec.* **2010**, *39*, 41.
32. Santos, R.L.; Macdonald, C.; Ounis, I. Exploiting query reformulations for web search result diversification. In Proceedings of the 19th International Conference on World Wide Web, Hong Kong, China, 1–5 May 2010; pp. 881–890.
33. Agrawal, R.; Gollapudi, S.; Halverson, A.; Ieong, S. Diversifying search results. In Proceedings of the second ACM international conference on web search and data mining, Barcelona, Spain, 9–11 February 2009; pp. 5–14.
34. Hu, S.; Dou, Z.; Wang, X.; Sakai, T.; Wen, J.R. Search Result Diversification Based on Hierarchical Intent. In Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, Melbourne, Australia, 19–23 October 2015; pp. 63–72.
35. Langville, A.N.; Meyer, C.D. A survey of eigenvector methods for web information retrieval. *SIAM Rev.* **2005**, *47*, 135–161.
36. Moens, M. Innovative techniques for legal text retrieval. *Artif. Intell. Law* **2001**, *9*, 29–57.
37. Biagioli, C.; Francesconi, E.; Passerini, A.; Montemagni, S.; Soria, C. Automatic semantics extraction in law documents. In Proceedings of the 10th international conference on Artificial intelligence and law, Bologna, Italy, 6–11 June 2005.
38. Mencia, E.L.; Fürnkranz, J. Efficient pairwise multilabel classification for large-scale problems in the legal domain. In *Machine Learning and Knowledge Discovery in Databases*; Springer: Berlin, Germany, 2008; pp. 50–65.

39. Grabmair, M.; Ashley, K.D.; Chen, R.; Sureshkumar, P.; Wang, C.; Nyberg, E.; Walker, V.R. Introducing LUIMA: an experiment in legal conceptual retrieval of vaccine injury decisions using a UIMA type system and tools. In Proceedings of the 15th International Conference on Artificial Intelligence and Law, San Diego, CA, USA, 8–12 June 2015; pp. 69–78.
40. Saravanan, M.; Ravindran, B.; Raman, S. Improving legal information retrieval using an ontological framework. *Artif. Intell. Law* **2009**, *17*, 101–124.
41. Schweighofer, E.; Liebwald, D. Advanced lexical ontologies and hybrid knowledge based systems: First steps to a dynamic legal electronic commentary. *Artif. Intell. Law* **2007**, *15*, 103–115.
42. Sagri, M.T.; Tiscornia, D. Metadata for content description in legal information. In Proceedings of the 14th International Workshop on Database and Expert Systems Applications, Prague, Czech Republic, 1–5 September 2003; pp. 745–749.
43. Klein, M.C.; Van Steenberghe, W.; Uijttenbroek, E.M.; Lodder, A.R.; van Harmelen, F. Thesaurus-based Retrieval of Case Law. In Proceedings of the 2006 conference on Legal Knowledge and Information Systems: JURIX 2006: The Nineteenth Annual Conference, Paris, France, 8 December 2006; Volume 152, p. 61.
44. Hoekstra, R.; Breuker, J.; di Bello, M.; Boer, A. The LKIF Core ontology of basic legal concepts. In Proceedings of the 2nd Workshop on Legal Ontologies and Artificial Intelligence Techniques (LOAIT 2007), Stanford, CA, USA, 4 June 2007.
45. Farzindar, A.; Lapalme, G. Legal text summarization by exploration of the thematic structures and argumentative roles. In *Text Summarization Branches Out Workshop Held in Conjunction with ACL*; Association for Computational Linguistics (ACL): Barcelona, Spain, 25–26 July 2004; pp. 27–34.
46. Farzindar, A.; Lapalme, G. Letsum, an automatic legal text summarizing system. In Proceedings of the Legal Knowledge and Information Systems. JURIX 2004: The Seventeenth Annual Conference, Berlin, Germany, 8–10 December 2004; pp. 11–18.
47. Moens, M.F. Summarizing court decisions. *Inf. Process. Manag.* **2007**, *43*, 1748–1764.
48. Aktolga, E.; Ros, I.; Assogba, Y. Detecting outlier sections in us congressional legislation. In Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval—SIGIR '11, Beijing, China, 24–28 July 2011; pp. 235–244.
49. Marx, S.M. Citation networks in the law. *Jurimetr. J.* **1970**, *10*, 121–137.
50. Van Opijnen, M. Citation Analysis and Beyond: In Search of Indicators Measuring Case Law Importance. In Proceedings of the Legal Knowledge and Information Systems-JURIX 2012: The Twenty-Fifth Annual Conference, Amsterdam, The Netherlands, 17–19 December 2012; pp. 95–104.
51. Fowler, J.H.; Jeon, S. The Authority of Supreme Court precedent. *Soc. Netw.* **2008**, *30*, 16–30.
52. Fowler, J.H.; Johnson, T.R.; Spriggs, J.F.; Jeon, S.; Wahlbeck, P.J. Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court. *Political Anal.* **2006**, *15*, 324–346.
53. Galgani, F.; Compton, P.; Hoffmann, A. Citation based summarisation of legal texts. In Proceedings of the PRICAI 2012: Trends in Artificial Intelligence, Kuching, Malaysia, 3–7 September 2012; pp. 40–52.
54. Koniaris, M.; Anagnostopoulos, I.; Vassiliou, Y. Network Analysis in the Legal Domain: A complex model for European Union legal sources. *arXiv* **2015**, arXiv:1501.05237.
55. Lettieri, N.; Altamura, A.; Faggiano, A.; Malandrino, D. A computational approach for the experimental study of EU case law: analysis and implementation. *Soc. Netw. Anal. Min.* **2016**, *6*, doi:10.1007/s13278-016-0365-6.
56. Koniaris, M.; Anagnostopoulos, I.; Vassiliou, Y. Multi-dimension Diversification in Legal Information Retrieval. In Proceedings of the International Conference on Web Information Systems Engineering, Shanghai, China, 8–10 November 2016; pp. 174–189.
57. Wittfoth, A.; Chung, P.; Greenleaf, G.; Mowbray, A. AustLII's Point-in-Time Legislation System: A generic PiT system for presenting legislation. Available online: http://portsea.austlii.edu.au/pit/papers/PiT_background_2005.rtf (accessed on 25 January 2017).

