



Modeling law search as prediction

Faraz Dadgostari¹ · Mauricio Guim² · Peter A. Beling¹ · Michael A. Livermore¹  · Daniel N. Rockmore³ 

© The Author(s) 2020

Abstract

Law search is fundamental to legal reasoning and its articulation is an important challenge and open problem in the ongoing efforts to investigate legal reasoning as a formal process. This Article formulates a mathematical model that frames the behavioral and cognitive framework of law search as a sequential decision process. The model has two components: first, a model of the legal corpus as a *search space* and second, a model of the search process (or *search strategy*) that is compatible with that environment. The search space has the structure of a “multi-network”—an interleaved structure of distinct networks—developed in earlier work. In this Article, we develop and formally describe three related models of the search process. We then implement these models on a subset of the corpus of U.S. Supreme Court opinions and assess their performance against two benchmark prediction tasks. The first is to predict the citations in a document from its semantic content. The second is to predict the search results generated by human users. For both benchmarks, all search models outperform a null model with the learning-based model outperforming the other approaches. Our results indicate that through additional work and refinement, there may be the potential for machine law search to achieve human or near-human levels of performance.

✉ Michael A. Livermore
mlivermore@virginia.edu

Faraz Dadgostari
fd4cd@virginia.edu

Mauricio Guim
mauricio.guim@itam.mx

Peter A. Beling
pb3a@virginia.edu

Daniel N. Rockmore
daniel.n.rockmore@dartmouth.edu
<https://web.cs.dartmouth.edu/people/daniel-rockmore>

¹ University of Virginia, Charlottesville, USA

² Instituto Tecnológico Autónomo de México (ITAM), Mexico City, Mexico

³ Dartmouth College, Hanover, USA

1 Introduction

In her now classic definition of the program of research into artificial intelligence and law, Edwina Rissland sketched out the goals of the field and identified three related undertakings: first, to know “how to represent several types of knowledge, such as cases, rules, and arguments”; second, to know “how to reason with [those types of knowledge]”; and third, to know “how to use [those types of knowledge] ultimately in a computer program that can perform tasks in legal reasoning and argumentation” (Rissland 1990).

This agenda has helped animate research into artificial intelligence and law for nearly three decades, with some important and notable successes (Bench-Capon et al. 2012). Rissland’s list does the crucial work of articulating much of the conceptual and process framework of legal argumentation, but unarticulated is the crucial and foundational step of *identifying* the legal knowledge relevant to the issue or argument at hand. In any given matter, before legal reasoning can take place, the reasoning agent must first engage in a task of “law search” to identify the legal knowledge—cases, statutes, or regulations—that bear on the questions being addressed. This task may seem straightforward or obvious, but upon inspection, it presents difficult problems of definition and is challenging to represent in a tractable formalization that can be computationally executed. In this Article we take on the problem of articulating and formalizing the law search process and report the results of an initial experiment that provides a framework for comparing machine and human performance on this core legal task.

The issue of “law search” has both conceptual and practical significance. Conceptually, on some accounts, what distinguishes legal reasoning from other forms of reasoning (such as practical or moral reasoning) is precisely its reliance on legal materials (Schauer and Wise 1997, 2000). The search for legal materials, then, implicates the very nature of legal reasoning and some of the most controversial questions in legal philosophy, such as “what is law?” (Green 1996). Although the tools of artificial intelligence research may not be well-suited to answering core questions in jurisprudence, they can be used to explore the implications of answering these philosophical questions in different ways. From a practical perspective, the ability to recognize and search for relevant legal authorities is one of the most basic skills of a legal professional. Were this skill readily accomplished by a machine, it could substantially lower the cost of legal services and democratize access to the law.

While fundamental to law, law search also connects to a wide range of research programs across disciplines. The question of *search*, understood broadly has been analyzed in a number of fields, including economics (Kohn and Shavell 1974; Diamond 1982; Mortensen and Pissarides 1994), medicine (Ledley and Lusted 1959), biology (Ramos-Fernández et al. 2004), and game theory (Alpern and Gal 2006). Calvin Mooers (1950) coined the phrase “information retrieval” to describe “[t]he problem of directing a user to stored information, some of which may be unknown to him [or her]”. Since that time, a substantial research program in computer science has developed to analyze and respond to the information

retrieval problem, especially in the context of unstructured corpora of text documents (Manning et al. 2010). The advent of the World Wide Web and other new collections of digitized documents enlivened the field with both research and commercial opportunities.

Information retrieval intersects with the field of natural language processing and computational text analysis more generally because statistical representations of language are fundamental to effective text-oriented information retrieval systems (Smeaton 1999). Early approaches to the statistical modeling of documents, such as term-frequency/inverse-document frequency (tf-idf) helped demonstrate the usefulness of mathematical representations of documents for information retrieval purposes (Luhn 1957; Sparck Jones 1972). Deep learning techniques such as neural networks have been used to address information retrieval problems (Joachims 1998). Most recently, topic modeling has become a widely used form of statistical representation of document corpora (Blei et al. 2003; Blei and Lafferty 2007; Blei 2012). For citation-linked corpora, various kinds of network-based approaches, sometimes coupled with text representations, are used (Leibon et al. 2018; Clark and Lauderdale 2012). The last of these is most relevant to the work presented herein.

Information retrieval frames the question of search in formal terms. By contrast, researchers in the field of information behavior take a user-centric approach that focuses on how people approach the problem of information search across media (Fisher et al. 2005). Early work emphasized the personal and social factors that give rise to the need for information (Wilson 1981). Scholars also have addressed the iterative nature of search (early search results influence subsequent searches), the complexity of the roles and tasks that motivate search, and the importance of environmental factors that can facilitate or inhibit successful search (Leckie et al. 1996). Researchers have also focused on the cognitive and psychological aspects of search (Meho and Tibbo 2003). In recent years, search behavior has been examined in a range of contexts, including by social and natural scientists (Hemminger et al. 2007), doctors (Davies 2007), engineers (Robinson 2010), and lawyers (Wilkinson 2001). Of late, there has also been a good deal of attention paid to bias in search engine results (Höchstötter and Lewandowski 2009).

The project described in this Article draws from both fields—information behavior and information retrieval—to construct and test a computational model of law search. We have two general goals. The first is to understand the practice of law search, as instantiated in the behavior of actual practitioners. We follow Rissland in the view that formalizing legal tasks in simplified terms and instantiating them in executable code can facilitate understanding of legal decision making. For us, the specific question that drives our research is how human agents engage in law search. This work can also shed light on potential biases in human law search. Our second goal is practical: we develop the foundations for future research into implementable versions of *automated law search* that could be deployed by commercial enterprises or governments. An automated law search tool that approaches human-level proficiency has substantial potential to lower barriers to accessing the law.

The plan for this Article is as follows. We first propose a qualitative description of an optimization function of a human agent engaged in law search. Based on the role of law in legal reasoning and argumentation, we define this optimization function as a

predictive task. Many search tasks can be framed as prediction, especially when implicitly understood as providing a *recommendation* to searchers: the search engine “predicts” what a user wants to find (Page et al. 1999). In the context of law search, the optimization function involves identifying legal materials that would be deemed relevant by an authoritative legal decision maker. Our definition of the optimization problem in law search is thus akin to a model of law championed by American legal realists in the twentieth century and made famous by Justice Holmes: “A legal duty so called is nothing but a *prediction* that if a man does or omits certain things he will be made to suffer in this or that way by judgment of the court” (Holmes 1897, 458). The predictive theory of law has its critics, most famous among them H.L.A. Hart. But we defend our choice of a Holmesian optimization model based on the internal/external divide introduced by Hart himself: our goals involve the scientific study of legal practices as outsiders, rather than the practice of law as internal participants.

We next describe our formal models of law search. Our model comprises a *search space* that can be traversed using a *search strategy*. The search space is based on the *legal landscape* developed in Leibon et al. (2018), and we introduce three search algorithms in this article. The first, which we call the *proximity* algorithm, serves as a baseline that we use to benchmark the other approaches. The *proximity* algorithm matches naive behavior executed over the Leibon et al. (2018) landscape, which itself is constructed using state-of-the art approaches from the information retrieval literature. The second is the *covering* algorithm, which builds in parameters to balance the depth and breadth of exploration during the search process. A third *adaptive* algorithm combines a Markov Decision Process and a reinforcement learning framework to learn the values for the free parameters in the *covering* algorithm based on information in the relevant legal corpus. The adaptive algorithm most closely approximates search as prediction: because it *learns* parameters from the data, its outputs can be interpreted as predictions based on prior evidence.

After describing these models, we discuss an implementation on a selection of U.S. Supreme Court opinions. For purposes of assessment, we identify two tasks. The first is to use the semantic content of a source judicial opinion to predict the other opinions that are cited in the source document. The second task is based on data generated by law student research assistants. The research assistants were given a source document and asked to identify related opinions, and the models are evaluated according to how well they predict the research assistants’ answers. As will be discussed in more detail below, the search approach described in this article—and in particular the training of a predictive model based on the semantic and citation information in judicial opinions—shows substantial promise and, with refinement, could lead to automated law search that approaches human levels of performance, a development that would have substantial practical significance.

2 The goals of a law searcher

2.1 Describing search behavior

What happens when someone sits down to search through the law? Typically, the starting point will be a legal question, and the person engaged in legal search (a “law searcher”) attempts to gain information that is relevant to that legal question. For example, an employer might have a question about her flexibility to change break and shift times for certain employees, which might be governed by contractual provisions, collective bargaining agreements, or state or federal labor or workplace safety laws. The employer (or more likely her representative) then identifies the relevant private agreements, statutes, and case law that might bear on the question of employer/employee relations in this particular context.

Legal education provides some insight into the process of law search, as that training reflects widely shared views and practices about how law search is and should be carried out. Although contemporary lawyers now rely on digital tools to engage in law search, it is not uncommon for an introductory course in legal research to begin with analogue approaches to legal research. By introducing students to these (largely outdated) analog tools, instructors presumably hope to impart some insight into the general task of search and help their students develop a mental model for it, which can be used even with more sophisticated computationally based search tools.

In keeping with its common law tradition, in the United States law students will frequently be introduced to law search in the context of judicial opinions, which are collected and organized chronologically by court and region in various law reporter volumes, which are then indexed via the West keynote system. Legal search using West reporters and indexes involves, as a first step, placing a legal question within a West category or subcategory and then finding the related cases. Once some set of opinions is identified using an index, the opinions themselves and the citations they contain serve as an additional resource. Embedded within the relevant language in the decisions will be citations to relevant authorities (or backward citations). An additional reference book, *Shepard’s Citations*, collects forward citations, which are subsequent decisions that cite back to the opinion of interest. Law search involves at least partially tracing forward and backward along this citation network. In addition, other opinions may include different legal categories than those that the searcher initially identified, and the process can begin anew for those categories within the West index.

This type of law search may serve a useful pedagogical function, but, as discussed above, in practice legal search now relies on different tools and encompasses a larger number of sources. Rather than using reporters and paper indexes, searchers typically rely on search tools that are available on commercial legal databases. These search tools rely on either simple keyword approaches or use natural language algorithms to identify relevant areas of law. Mutual citation is also easily traceable through the commercial databases, and digital versions of the keynote system also exist.

A researcher will typically toggle between keyword/natural language searches, a keynote-like annotation system, secondary sources, statutory/regulatory hierarchy, and citation/cross-references to gather the relevant legal texts that he or she is looking for. This model of search leads to a search chain. The first link in the chain is identified through a secondary source, annotation system, or (most likely) search engine. Subsequent links are found via hierarchical structure, citation, or subsequent passes at the index, secondary source, or the search engine. Each step in the process provides some information that can be used to inform the next step, and so on. Eventually, a cutoff point is reached where the researcher decides that he or she has found enough information. At this point, the cost of searching out more authority (which is likely of diminishing relevance) is greater than the benefits of finding additional sources.

2.2 Relevance and prediction

The description of the search process in the preceding subsection explains *how* researchers seek out legal information, but it does not explain *what* they are looking for. At the core of this second question is one of *relevance*: the search process is directed at identifying legal materials that are relevant to a given legal question.

The concept of relevance has been much discussed both within the fields of information retrieval and information science. The information science scholar Tefko Saracevic (1996) has written extensively on the topic of relevance and argues for a multi-faceted understanding of the term that includes several types of subjective relevance judgments, such as topical relevance (aboutness), cognitive relevance (informativeness and novelty), situational relevance (usefulness), and motivational relevance (satisfaction). Often, relevance is used for purposes of evaluating information retrieval systems. When used for this purpose, ground truth is considered to be human agreement on judgements of relevance (Carletta 1996). Researchers in information retrieval have developed a number of metrics for assessing the performance of systems based on subjective human judgments, such as precision (relevant retrieved document over all retrieved documents), recall (relevant retrieved documents over relevant documents), and f-score (a combined metric of precision and recall) (Blair and Maron 1985).

We define a notion of *legal relevance* that borrows from Saracevic and the uses of the concept of relevance in the field of information retrieval. The definition we offer is sociological in nature, and is based on the judgments made by a legal community. In our definition, *a document is legally relevant exactly when it is understood by the dominant legal community as containing information that bears on a legal question of concern*. In other words, the legal document is relevant when it can be used as the basis for argumentation and analysis with regard to that question. This notion of relevance is related to Saracevic's "situational relevance" and is highly purpose-driven. Relevance is determined functionally with respect to norms and practices concerning legal reasoning and argumentation within a legal community.

This sociological notion of relevance carries through the institutional setting of courts and the role that texts play in legal argumentation. In an adversarial

context, lawyers will seek out the authority that best supports their position, but they will also seek out the authority that best supports their opponent's case, so that they can prepare and develop counterarguments. Both sides will seek any authority that they believe a judge would find binding or persuasive. The judge, in carrying out (or having a clerk carry out) additional research, will search for authority that a reviewing panel will find binding or persuasive. Even a final reviewing court searches for authority in this predictive fashion, both to project influence down the judicial hierarchy (if lower courts better conform to precedent that they respect) and forward in time (assuming that appropriately defended decisions will be more resistant to change).

Legal relevance, as we define it here, is related to the predictive theory of the law first introduced by Oliver Wendell Holmes. The predictive theory can be summed up in Holmes's famous declaration: "The prophecies of what the courts will do in fact, and nothing more pretentious, are what I mean by the law" (Holmes 1897, 461). *Mutatis mutandis*, we take "prophecies of [the materials] courts will [examine] in fact, and nothing more" to be the definition of legal relevance. One of the core advantages of the predictive theory of law and law search, is that it "accurately describes what lawyers do-in fact: [U]se legal doctrine to predict the outcomes of future cases" (Hantzis 1987, 342).

Despite its pedigree and relationship to actual law practice, the predictive theory of law has its critics. Most famously, H.L.A. Hart argued that the predictive theory failed to account for the "internal" perspective of legal actors who chose to follow the law out of felt obligation, rather than out of concern for sanctions (Hart 1961). For someone interested in the law's normative force, it is not clear that predictions about what judges would do are the appropriate inquiry. The predictive theory has also been critiqued because it seems to give judges little guidance.

A related claim is that it is incoherent to think of unearthing the law's content absent an exercise in normative reasoning (Dworkin 1978; Fuller 1958; Finnis 2011). Under this view, any statement about what the law says is irreducibly normative because the law simply is a normative enterprise. Note that this jurisprudential view is not merely that people might apply their normative priors or be affected by confirmation bias when engaged in legal reasoning (Jones and Sugden 2001). Rather, the idea is that legal reasoning necessarily relies on moral reasoning, so that merely predicting the behavior of judges is not, as a definition matter, legal reasoning.

Both Hart's and the related critiques could be applied to a predictive theory of law search as well, in that such a theory can never answer the question of what the "truly" relevant legal authority is for a given legal question, with respect to a normative theory of law. The question of normative legal relevance may be a worthwhile line of jurisprudential inquiry, given the centrality of law search for legal reasoning. To date, legal philosophers have not taken up the question of relevance, or have done so in fairly limited ways. For example, Dworkin (1978) assumes that *all* law is in some way relevant to *every* legal question, and then side-steps the obvious practical challenges that would pose by imagining a theoretical super-judge with unlimited cognitive capacity.

However, one need not adopt Hart's internal perspective or the theory of law as inherently moral. The alternative is the *external* perspective of a non-participant. The external perspective might be adopted by an economist or a historian who seeks to understand legal decisions not in terms of the legal reasons given by the participants, but on the basis of (for example) ideological preferences or historical circumstances. This external perspective is, in a sense, the scientist's perspective—indeed, Holmes's predictive theory has been characterized by some legal philosophers as a “scientific” approach to law (Wells 1993). For researchers in artificial intelligence and law, this external scientific perspective seems entirely appropriate, even if there are certain types of (normative) questions that necessarily remain outside the field.

2.3 Convergence

According to the discussion in the prior subsection, relevance is understood as a social fact that amounts to a prediction about what others will deem relevant. This prediction is both substantive and procedural. It is substantive in the sense that a searcher reviewing a document must judge (or predict how others will judge) whether that document bears in a substantive way on the legal question at hand. The prediction is procedural in the sense that there are norms, conventions, and practices concerning how search is conducted. A document that is difficult to locate using the law search methods of the dominant legal culture is less relevant precisely because it is unlikely to be found, and as a consequence is unlikely to be deemed relevant by other legal searchers. And so the process of search as actually practiced has consequences for what is or is not relevant law.

When there is agreement about the relevant law for a given legal question or dispute, the parties will converge on some set of sources. “Convergence” can be understood as a characteristic of the legal system that arises from several interacting features, which can all facilitate or reduce convergence. The law itself is one feature: legal systems that have a larger number of documents can expect, other things being equal, less convergence. The information systems used to organize the law and facilitate document retrieval are another factor. Some information systems may facilitate convergence by channeling searchers in particular ways, while others may facilitate idiosyncratic search patterns. The community of legal searchers, and the norms, practices, and conventions that structure their behavior, is another feature that might influence convergence. If informal practices are widely shared or well known, that may enhance convergence. Finally, the incentives of the parties to conform to community expectations concerning relevant legal authority can matter. The higher the penalty for failing to identify law that is deemed relevant, the more effort will be undertaken to predict what others will think the law is.

Although the notion of convergence has procedural elements, convergence of results does not imply convergence in the search chains that generate those results. In practice, the search chains generated by legal searchers could be idiosyncratic. Even the same researcher, approaching a problem at different times, may follow different paths through a corpus, depending on a wide variety of factors that could be

as ephemeral as the loading time of a particular document in a browser window. The differences between individuals is likely to be even greater, as legal searchers, by habit, inclination, or experience, come to rely on different sources or approaches. Nevertheless, the materials that are discovered through these different chains can converge.

A final related feature of the search process is that it is an individual behavior that occurs with little centralized coordination. Although at a micro level it is possible that search tasks are divided within small hierarchies (for example, a law firm team), there is no large-scale centralizing body that coordinates law search. Entities like Westlaw or LexisNexis and texts such as the American Law Institute's Restatements play a moderate coordinating role by standardizing search practices and providing focal points. But, at least for more sophisticated questions, legal searches are carried out in a decentralized fashion. To the extent that there are system-wide properties associated with law search (including the level of convergence in the system) those properties emerge from uncoordinated behavior of individuals.

3 The search space

The rough qualitative description of the optimization problem of law search can be translated to a more formal model. The starting place for the models that will be discussed in this paper is the notion of *navigating* through a corpus of documents. Navigation focuses on the portion of search in which a searcher moves from one document to another within a corpus. It is important to note that navigation does not include the initial query that led to some starting place within the corpus. We refer to this starting place as a source document. The process of identifying the source document is undoubtedly an important part of law search, and may be amenable to study based on insights from artificial intelligence and law. The commercial databases now deploy sophisticated natural language algorithms capable of generating meaningful search returns based on unstructured user queries. Navigation, however, focuses on the subsequent step of the search process, after a source document is identified, in which information in that document is used to inform future steps in a search path. This navigation process is our focus here.

The space of documents that is being navigated (i.e., the search space) is represented as a network in which documents are nodes with edges between them. The search space is a “multinetwork” because it is based on two network structures: the directed network derived from citation information and the network based on textual similarity (see generally Tseng 2010). Our multinetwork representation comes from Leibon et al. (2018) and generalizes the well-known PageRank algorithm (Page et al. 1999; Langville and Meyer 2011) to produce a symmetric matrix that incorporates textual similarity and citation networks.

The multinetwork recognizes two types of edges between documents. One type of edge is based on cross-reference information. For judicial opinions, that means that a directed edge is created when one document cites to another. This structure is grounded in the qualitative observation discussed above that searchers often use citations as one way to identify documents of interest. Leibon et al. (2018) use citation

data to produce two matrices, “*cited-by*” and “*cited*”, which together represent the citation network of documents. Citations can be taken as proxy for another kind of relatedness between pairs of documents that is not captured by textual similarity.

A second set of edges between documents is constructed based on their semantic content (i.e., the words contained in those documents). Semantic content can be understood as a proxy for several different actual search mechanisms used by law searchers, including keyword searches, curated categorizations (i.e., Westlaw headnotes), and other sources (such as treatises). The operating assumption is that documents with similar words will show up together in keyword searches, be grouped under similar headnotes, and appear in the same treatises.

There are many ways to represent semantic content, and there is a balance that needs to be struck between total information, dimensionality, coarse versus fine-graining, and computational costs. We opt for a *topic model* representation (Blei et al. 2003; Blei and Lafferty 2007; Roberts et al. 2014). Topic models are becoming an increasingly common tool for analytic work that is textually based (e.g., political science and literary studies) (Quinn et al. 2010; Livermore et al. 2017; Ridell 2014). They are based on term frequency vector representations of documents (i.e., “bag-of-words” representations), which are effectively lists of word frequencies or proportions, indexed by a vocabulary of words. Bag-of-words representations achieve (by some measure) massive dimension reduction of the corpus, because word order is ignored. The highest possible dimensionality would effectively assign a dimension to the word type and position of each word in the document, leading to an explosion of dimensions that would prove computationally intractable. Topic models reduce dimensions even further from the vocabulary size to a relative handful of dimensions equal to the number of topics. Very roughly, this is achieved through the use of word co-occurrence to construct subject matter categories represented as distributions over the vocabulary (these are the “topics”). Documents are then represented as distributions over topics. Prior research has shown that topic model representations of judicial opinions retain a considerable amount of the original data found in a full-term frequency vector (Livermore et al. 2017). Thus, topic models achieve some level of coarse-graining, which can reduce the influence of highly idiosyncratic language.

To retrieve a latent structure of the legal corpora from the texts and network structure of existing legal documents, we follow the approach in Leibon et al. (2018). In addition to the directed citation network, we use a latent Dirichlet allocation (LDA) topic model (see e.g., Blei and Lafferty (2007)) to build a low-dimensional mathematical representation of the legal corpora. The topic model is used to construct the foundation of a measure of textual similarity between two documents in the legal corpus.

The citation information and topic model data enables the construction of a metric space of legal documents, $(O, \text{PageDist})$ where $O = \bigcup_{i=1}^N \{o_i\}$ is the set of opinions in the legal corpus of size N and PageDist is a metric defined on O , representing a distance function of opinions $\text{PageDist} : O \times O \rightarrow \mathbb{R}^+$, derived from the stochastic matrix of the transition probabilities of a random walk model over the legal corpus. PageDist is derived from a weighted aggregation of citation matrices and the textual similarity matrix. Different weighting settings result in different PageDist functions.

Accordingly, different distance functions can be defined over the same opinion set. Here for simplicity, *PageDist* is derived assigning equal weights to each of three matrices:

$$w_{sim} = w_{cited} = w_{cited-by} = 1/3; w_{sim} + w_{cited} + w_{cited-by} = 1$$

As described in Leibon et al. (2018), all three matrices are normalized; accordingly, a weighted aggregation of the matrices after normalization produces a stochastic matrix. This stochastic matrix can be considered a transition probability matrix for a random walk over the legal corpus which is equivalent to a Markov chain with a state space given by the documents. *PageDist* as defined by Leibon et al. (2018) holds the properties of a metric (distance) which can further be used to build a geometry that models the legal corpus as a normed space with a relevant notion of curvature. As a normed space of legal documents, this *legal landscape* embodies a meaningful mathematical structure for legal documents. Accordingly, we use it to define the space over which law search occurs (i.e., the *search space*).

4 Search strategies

The search strategies that are described below all operate over the same search space. However, they allow users to navigate through the space differently, resulting in different search outcomes. By way of analogy, imagine a robot that is programmed to navigate through some physical space for purposes of cleaning the floors. The robot could move around using different strategies: one strategy might be to cover all of the ground that is close to its starting position and move outward; another approach might be to just randomly set off in a given direction and periodically make random turns; another strategy might be to clean in a single location for a while and then periodically move to a distant portion of the space. One can think of the search strategies and search space in a similar fashion, with a shared space (the Leibon et al. (2018) legal landscape) that is navigated using different approaches (the *proximity*, *covering*, or *adaptive* algorithms).

The first search strategy that we introduce is the *proximity algorithm*, which serves as the baseline model that can be used to test the performance of the *covering* and *adaptive* approaches. The proximity algorithm baseline represents the state of the art drawn from the information retrieval literature. The PageDist measure proposed in Leibon et al. (2018) is similar in spirit to leading information retrieval model that combine language modeling and link-analysis into a unified setting, such as RankLDA, HITSLDA (Zhang et al. 2018), LIMTopic (Duan et al. 2014), LC-LDA (Liu and Xu 2017) and RankTopic (Duan et al. 2012).

Under the proximity strategy, the searcher navigates through the space based purely on proximity in the Leibon et al. (2018) legal landscape. Beginning with a source document as a starting place, the proximity algorithm simply picks up all of the documents that are closest to the source document within the space. An important feature of the proximity algorithm is that its results are particularly strongly

related to the search space itself, as the results that are generated using this strategy directly reflect proximity as represented in the search space. Because the distance measure is based on a random walk of the network, it can be thought of as encoding, in a probabilistic fashion, how a “random law surfer” will move from one legal document to the next. Similar Markov machinery underlies the earliest form of Google’s search algorithm, *PageRank* (Brin and Page 1998).

The second search strategy that we introduce is the *covering algorithm* (see Caprara et al. 2000). The covering algorithm is meant to capture the fact that there are often multiple legal issues within a single document. The covering model begins with a source document and then identifies the most proximate document. Then, based on some fixed parameters, it determines whether to continue navigating along that line of documents, or to return to the source document and begin the search again along a different line. The idea is that there can be multiple legal issues present in a document, and once a searcher is satisfied with the results on one issue, it may make sense to explore a second or third legal issue, rather than continuing collecting documents on the first. Embedded within the covering algorithm is an assumption about how to make the tradeoff between depth (i.e., exploring one issue in more detail) and breadth (i.e., exploring a larger number of issues). This tradeoff is expressed in terms of a set of parameter values concerning the number of issues to explore and how deeply to explore them.

The final search strategy that we introduce is the *adaptive algorithm*. As far as we know, our adaptive algorithm is an innovative approach to representing the search problem. This strategy is akin to the covering algorithm, but rather than using defined values for the breadth/depth tradeoff, the parameters are learned from the corpus, using a reinforcement learning approach. The learner uses the documents in the corpus and the citations included in those documents as data for a training procedure in which parameter values that correctly predict citation are reinforced. One can think of the adaptive algorithm as akin to a law student who learns the types of cases to identify (and cite) by studying the cases that have been identified (and cited) by the experts who produced the existing stock of documents in the corpus.

Each of these strategies takes different approaches to capturing features of how law search is carried out. Because they must be formalized and converted into an executable program, they are by nature simplified representations of the complex, idiosyncratic, and stochastic human search process. However, they nevertheless capture many important features of law search: relevance of semantic content; guidance via citations; and tradeoffs between depth and breadth. In practice, human legal researchers rely on a variety of tools not explicitly represented in the models. These may include their background understanding of the relevant law or secondary sources, such as treatises or the ALI’s Restatements. But a good deal of this “out-of-model” information may be proxied in features that can be extracted from the documents and therefore can be, at least loosely, captured by the models.

4.1 Problem formulation

The purpose of law search is to identify legal authorities that are relevant to a legal question. To formalize this problem, we define a “citation free legal text” (CFLT), as a formal semantic description of a legal matter that lacks citations to any authority. A CFLT can be constructed by removing the citation information from a legal opinion. We then define a legal corpus (which can roughly be understood as corresponding to a jurisdiction) and then define a metric (called *PageDist*) that gives rise to a geometric structure over the corpus. We then use the resulting definition of the legal corpus to define the problem of search as finding a subset of the legal corpus that can be recommended as an appropriate citation set for a CFLT.

4.1.1 Definition—citation free legal text (CFLT)

A CFLT is a text string that discusses a legal matter and lacks citation information. A CFLT and an opinion are both textual data addressing a legal matter—the primary relevant difference is that observed opinions contain citation data, whereas a CFLT does not. In other words, a CFLT is an opinion that cites no other document and is cited by no other document.

Accordingly, if we add a CFLT to a set of opinions comprising a legal corpus, the resulting set is a new body of legal documents, $D = O \cup \{CFLT\}$, which can be modeled as a geometrized metric space (D , *PageDist*) in the manner described earlier. This is helpful because having the given CFLT and all opinions of the legal corpus in the same metric space simplifies the development of the citation recommendation algorithm. Specifically, consider an opinion o_i and a CFLT formed by deleting the citation information from o_i . Then, if the CFLT is fed to a recommendation algorithm, the result can be evaluated in comparison to the citations in o_i .

Abstractly, we can formalize the legal search problem as a citation recommendation problem, where legal search aims to find the most suitable set of opinions $C \subset O$ to be cited for a given CFLT. The target set of opinions, $C(CFLT)$ can be called *search results* or *citation recommendations*. These results are assumed to have an exogenously set size, $|C(CFLT)| = n$, which represents the limited capacity of the search process. A means of evaluating the suitability of $C(CFLT)$ is to compare them to the actual citation information in o_i .

4.2 Proximity strategy

An intuitive citation recommendation algorithm selects as the recommended citation set those opinions in O that, in the landscape of legal documents, (D , *PageDist*), are closest to the CFLT. This is precisely the proximity method defined above.

More formally, consider the ball of radius $r \in \mathbb{R}^+$ centered at $CFLT \in D$ defined as:

$$B_r(CFLT) = \{o_i \in O \mid \text{PageDist}_D(CFLT, o_i) \leq r\}.$$

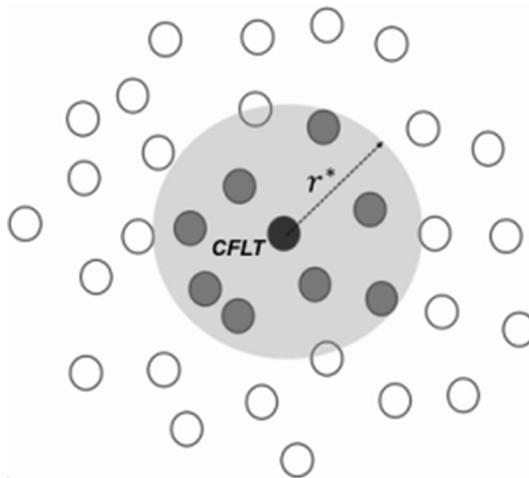


Fig. 1 Proximity ball

Then by solving the problem:

$$\begin{aligned} & \underset{r \in \mathbb{R}^+}{\text{maximize}} && r \\ & \text{subject to} && |B_r(CFLT)| \leq n \end{aligned} \quad (1)$$

we can obtain the radius r^* for the largest ball centered at $CFLT$ with at most n opinions contained in it. Therefore, $B_{r^*}(CFLT)$ represents the citation recommendation set determined by the proximity algorithm, or formally $C(CFLT) = B_{r^*}(CFLT)$ (Fig. 1).

4.3 Covering strategy

To model behavioral aspects of the search process, we assume that the searcher not only tries to find opinions similar to the $CFLT$, but also considers that there may be several different subjects or issues addressed in the $CFLT$. Accordingly, the searcher attempts to find opinions that are relevant to these different issues in the limited number of citations, n .

Let us assume that m is a predetermined number of issues within the $CFLT$, and n_k is the predetermined depth of the required search for each issue where, $k \in \{1, 2, \dots, m\}$ and $\sum_{k=1}^m n_k = n$. Then the result of the search will be the set of citations denoted as $C(CFLT) = \bigcup_{k=1}^m C_k(CFLT)$, where $C_k(CFLT)$ represents the suggested issue set k . We develop an iterative search process that begins initially with empty sets for the set of citations and the issue sets, or formally $C(CFLT) = \emptyset$, and $C_k(CFLT) = \emptyset, \forall k \in \{1, 2, \dots, m\}$. Then, the iterative procedure starts with the first issue set, $C_{k=1}(CFLT)$, and at each step selects an opinion addressing the associated issue and adds it to $C_{k=1}(CFLT)$ until the issue set is filled with respect to its predetermined size ($|C_1(CFLT)| \leq n_1$). Then the algorithm moves to the next issue and follows the same procedure until all issues have been addressed and the associated opinions have been identified and added to the appropriate sets (Fig. 2).

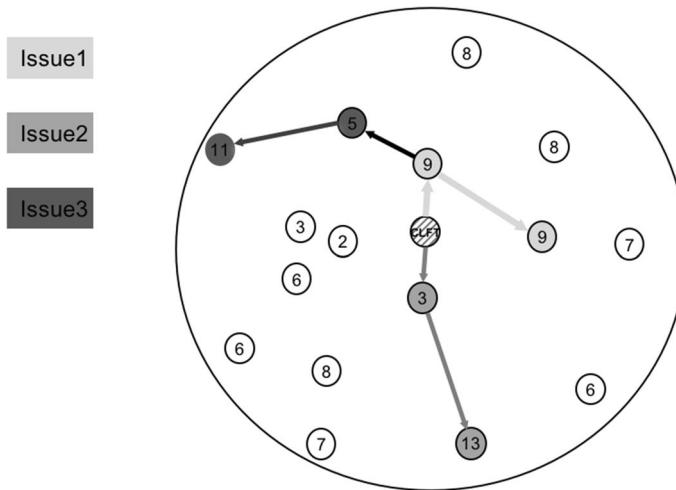


Fig. 2 Covering strategy

At each step of the covering method, we represent the opinions that have not yet been added to a prior set (and therefore are still candidate opinions to be added to the list of search results) as $O/C(CFLT)$. Also, if we refer to the issue that is being addressed at the current step of the procedure as \bar{k} , then $C(CFLT)/C_{\bar{k}}$ represents the set of opinions that have been added to the search results at the beginning of \bar{k} .

Here we define three criteria for the selection of an opinion to recommend at each step:

First, $b^{CFLT}(o_i) = \text{PageDist}(CFLT, o_i); \quad \forall o_i \in O/C(CFLT)$.

This represents the distance of a candidate opinion, o_i from $CFLT$.

Second, $b^c(o_i) = \sum_{o_j \in C(CFLT)/C_{\bar{k}}} \text{PageDist}(o_i, o_j); \quad \forall o_i \in O/C(CFLT)$.

$C(CFLT)/C_{\bar{k}}$ is the set of opinions that are already added to the search result. The term $b^c(o_i)$ represents the sum of distances of the candidate opinion, o_i from the opinions that were added to the search results prior to beginning of the current step in the search process (represented by \bar{k}).

Third, $b^s(o_i) = \sum_{o_j \in C_{\bar{k}}} \text{PageDist}(o_i, o_j); \quad \forall o_i \in O/C(CFLT)$.

This represents the sum of distances of the candidate opinion, o_i from the opinions that are already in the current issue set, $C_{\bar{k}}$.

If $C_{\bar{k}}$ is empty, the search process seeks the first opinion to recommend that addresses issue \bar{k} . In this case, the covering strategy will seek an opinion that is the closest to the original $CFLT$ but also different from the other opinions that have already been added to search results. It will do this by minimizing b^{CFLT} and maximizing b^c .

If $C_{\bar{k}}$ is not empty, then some opinions have been selected that address issue \bar{k} . The search process now seeks to add another opinion to $C_{\bar{k}}$. Here, the covering strategy seeks to suggest an opinion that is the closest to the opinions in $C_{\bar{k}}$ (which address the subject \bar{k}) while trying to be different from the opinions suggested for

other subjects and the *CFLT*. Therefore it should seek to minimize b^s and maximize b^{CFLT} and b^c .

The simplest way to combine these criteria is to use the weighted sum function defined as:

$$b(o_i) = w_1 b^c(o_i) + (-1)^{\mathbf{1}_{\{C_k=\emptyset\}}} w_2 b^{CFLT}(o_i) - \mathbf{1}_{\{C_k=\emptyset\}} w_3 b^s(o_i); \quad \forall o_i \in O/C(CFLT)$$

where $w_1 + w_2 + w_3 = 1$, $w_1 + w_3 \neq w_2$, $w_1 \neq w_2$ and $w_1, w_2, w_3 \geq 0$.

For the covering algorithm, we set these weights as $w_1 = 1/2$, $w_3 = w_2 = 1/4$. The general setting of these parameters should satisfy the previous four constraints and otherwise can be tuned. For purposes of our analysis, we chose to assign half of the weight to textual similarity ($w_1 = 1/2$) and then divided the remaining weight equally between the quality measure (discussed below) ($w_2 = 1/4$) and exploration ($w_3 = 1/4$). We did not undertake an effort to tune these parameters, rather, we developed the *adaptive* algorithm specifically to learn these parameters from the data.

4.3.1 Quality measure

In implementing the covering algorithm, we also introduce a quality measure for opinions, based on the number of citations that it has received, in the selection procedure. Let us define $\lambda(o_i)$, $\forall o_i \in O$, to represent the number of times opinion $o_i \in O$ has been cited by other opinions in O . Then to implement the covering algorithm with the quality measure, we modify the formulation to find:

$$o^* = \underset{o_i \in O/C(CFLT)}{\operatorname{argmin}} B \times b(o_i) + (1 - B) \times \lambda(o_i)) \quad (2)$$

where B represents the weight associated with b as the similarity measure and $(1 - B)$ represents the weight associated with the quality measure λ .¹

4.3.2 Algorithm

Initialize $C(CFLT) = \emptyset$ and $C_k(CFLT) = \emptyset$, $\forall k \in \{1, 2, \dots, m\}$

Set values for n , m and n_k where $k \in \{1, 2, \dots, m\}$ and $\sum_{k=1}^m n_k = n$

Set values for w_1, w_2, w_3 where $w_1 + w_2 + w_3 = 1$; $w_1 + w_3 \neq w_2$; $w_1 \neq w_2$ and $w_1, w_2, w_3 \geq 0$

Set value for B

For $\bar{k} = 1$ To m Repeat:

While $|C_{\bar{k}}(CFLT)| < n_{\bar{k}}$:

$$o^* \leftarrow \underset{o_i \in O/C(CFLT)}{\operatorname{argmin}} (B \times b(o_i) + (1 - B) \times \lambda(o_i))$$

$$C_{\bar{k}}(CFLT) \leftarrow C_{\bar{k}}(CFLT) \cup \{o^*\}$$

$$C(CFLT) \leftarrow C(CFLT) \cup C_{\bar{k}}(CFLT)$$

Return $C(CFLT)$

¹ $\lambda(o_i)$ only considers citations of the o_i by other opinions in the relevant corpus. This could be extended to other legal and non-legal documents. In the implementation of the covering algorithm, we set $B = 1/5$.

4.4 A learning approach

The covering algorithm is based on a search model in which all of the parameters that shape the behavioral structure of the model are given. Accordingly, the covering method has a behavioral structure, but the behavioral parameters are assumed and are not learned directly through the data.

The *adaptive algorithm* takes the search procedure as described in the covering algorithm section and learns the model parameters through reinforcement learning. Like law practitioners, the agent aims to learn an optimal search behavior by attempting to replicate the practice of others. In our case, the data for the learning process is generated through the citations in judicial opinions, and the goal of the algorithm is to predict the citations in a document based on its text alone. These predicted citations serve as the search results. Accordingly, in the adaptive algorithm, legal search is modeled as a Markovian sequential decision-making process in which the agent (searcher) tries to learn a set of policies (the parameter values) that generate results that mimic the citation behaviors of the judges. Note that there is a close link between search and citation, because before an opinion can be cited, it must be identified through a search process.

4.4.1 Markov decision processes

A Markov decision process (MDP) is a convenient and powerful mathematical framework to model decision contexts where outcomes are determined via a mix of random factors and agent-directed choices. MDPs have proven to be particularly fruitful when combined with reinforcement learning techniques.

Under a MDP, a decisionmaker in a given state s chooses an action a that is available in that state and then transitions to a new state s' based on a transition function, which is influenced by the choice a . A reward function determines the value, for the decision maker, of the move to s' . Importantly, given s and a , the state s' is conditionally independent of all previous states and actions. This process is then repeated from state s' , from which a new action a' is selected.

The MDP model can be applied to law search as follows. Imagine a legal actor in the United States who seeks to identify judicial opinions that bear on the following question: “What (if any) limitation on the sentencing discretion of federal judges are permissible under the U.S. constitution?” Using a keyword search tool, a researcher might identify *Sumner v. Shuman*, 483 U.S. 66 (1987), a case in which the Court addressed the question of “whether a statute that mandates the death penalty for a prison inmate who is convicted of murder while serving a life sentence without possibility of parole comports with the Eighth and Fourteenth Amendments”. This case can be understood as a state s that serves as the starting point for the search process.

Based on the information found in *Sumner v. Shuman*, including citations to other cases as well as specific legal terminology, the researcher may then take action a that transitions to a new case, *Lockett v. Ohio*, 438 U.S. 586 (1987), which deals with the question of whether an Ohio statute “that narrowly limits the sentencer’s discretion to consider the circumstances of the crime and the record and character of the offender as mitigating factors” offends the Constitution. After processing the

information in *Lockett*, the searcher may then move on to *Eddings v. Oklahoma*, 455 U.S. 104 (1982)—another death penalty case—and then to *Harmelin v. Michigan*, 501 U.S. 957 (1991), which upholds a statute that imposes a life sentence without the possibility for parole for a drug crime without considering mitigating factors. At each stage in this search chain, the researcher considers the information contained in the relevant document and then transitions to a new document. The reward function can be understood as arising from the usefulness of the information in a given document with respect to the legal question that motivated the search. The policies that are learned are a set of search practices that concern how the searcher transition from one document to the next.

In more formal terms, a finite-state MDP is defined as a tuple $M = (S, A, T, \gamma, r)$ where $S = \{s_1, s_2, \dots, s_n\}$ is a set of n states; $A = \{a_0, a_1, \dots, a_{m-1}\}$ is a set of m actions; $T : S \times A \rightarrow S$ is the transition function,² denoted by $T(s, a)$ which we define as the effects of an action $a \in A$ taken in a state $s \in S$; γ is a discount factor; and $r : S \times A \rightarrow R$ is the reward function, denoted by $r(s, a)$ which we define as depending on state $s \in S$ and action $a \in A$.

Consider a decisionmaker who selects actions according to a policy $\pi : S \rightarrow A$ that maps states to actions. Define the value function at state s with respect to policy π to be $V^\pi(s) = \sum_{t=0}^{\infty} \gamma^t r(s^t, \pi(s^t))$ where the sum is over the state sequence $\{s^0, s^1, \dots\}$, given policy π , where superscripts index time. A decisionmaker who aims to maximize reward will at every state s choose the action that maximizes $V^\pi(s)$. Similarly, define the *Q-factor* for state s and action a under policy π as $Q^\pi(s, a)$, to be the reward from state s , taking action a and thereafter following policy π . Given a policy π , $\forall s \in S, a \in A$, $V^\pi(s)$ and $Q^\pi(s, a)$ satisfy:

$$V^\pi(s) = r(s, \pi(s)) + \gamma V^\pi(s')$$

$$Q^\pi(s, a) = r(s, a) + \gamma V^\pi(s') = r(s, a) + \gamma Q^\pi(s', a')$$

where $s' = T(s, \pi(s))$ and $a' = \pi(s)$.

The well-known Bellman optimality conditions state that π is optimal if and only if, $\forall s \in S$ we have $\pi(s) = \operatorname{argmax}_{a \in A} Q(s, a)$ (Qiao and Beling 2011).

4.4.2 Q-learning algorithm

Reinforcement learning (RL) is a class of solution methods for MDP problems. To develop the adaptive algorithm, we make use of a RL method known as Q-Learning to solve the MDP formulation of the search problem. Q-Learning is more efficient than conventional methods such as dynamic programming. Also, Q-Learning is an appropriate approach for solving model-free problems, where it iteratively updates *Q-factors* based on sampling state-action rewards.

² A more general format for T is a $n \times n$ transition probability matrix, but here we restrict the model to be representative of a deterministic Markov process.

Q-Learning, similar to other RL paradigms, assumes that an agent can learn how to maximize the reward it gets from its environment by keeping track of the effects of actions in terms of both the immediate reward and positioning for future rewards, with the latter being expressed in terms of a state. Given a MDP problem, $M = (S, A, T, \gamma, r)$, the Q-Learning algorithm consists of state-action values, denoted by $\bar{Q}(s, a), \forall s \in S; \forall a \in A$, which represents an approximation of $Q^\pi(s, a)$ where

$$\pi(s) = \operatorname{argmax}_{a \in A} \bar{Q}(s, a)$$

determines the action taken in a given state $s \in S$, based on policy π .

Before the learning process starts, the agent assigns arbitrary values (like zeros) to $\bar{Q}(s, a)$ pairs, but each time it interacts with the environment by choosing an action $a \in A$ it receives an immediate reward r and accordingly updates the corresponding state-action value $\bar{Q}(s, a)$ as follows:

$$\bar{Q}(s, a) \leftarrow (1 - \alpha)\bar{Q}(s, a) + \alpha \left[r + \gamma \max_{a' \in A} \bar{Q}(s', a') \right] \quad (3)$$

where $s' = T(s, a)$ and α represents a learning rate ($0 \leq \alpha \leq 1$) which determines how much weight new information is given when updating the state-action values and γ is the discount rate ($0 \leq \gamma \leq 1$) for future rewards.

This procedure automatically updates policy π since it is dependent on the state-action values $\bar{Q}(s, a), \forall s \in S; \forall a \in A$. However it also uses an ϵ -greedy action selection procedure to determine the next action to take in the learning episode. This means that most of the time the action chosen has maximal action value $a^* = \arg \max_{a \in A} \bar{Q}(s, a)$ in the current state s , but with a small probability ϵ , the action is selected randomly.³

4.4.3 The search problem as a decision process

Legal search as modeled in previous sections is a sequential decision process of selecting opinions, which can be modeled as a MDP and solved by the Q-Learning method as follows.⁴

Under our MDP model of law search, the agent intends to identify cases that address different issues that are present in the source opinion. There are also a limited number of search results that are permissible. Similar to the covering algorithm, the searcher decomposes the source document, *CFLT*, into an ordered set of issues and then selects a set of opinions to recommend for each issue. These sets collectively form the set of search results. At each step of the search, depending on if the current issue set is empty or not, the searcher has a different set of actions to choose from.

³ In the implementation of the adaptive algorithm we set $\epsilon = 0.2$ for the first 2000 learning episodes and then it is reduced to $\epsilon = 0$ linearly in the next 3000 learning episodes.

⁴ In the developed adaptive algorithm as a Q-learning procedure, we set $\gamma = 0.4$, $\alpha = 1$.

One possible action is to start search on a new issue. This is referred to here as a *new-issue-action*. A *new-issue-action* step does not select a new opinion to add to the search results—it only starts a new issue and then proceeds to the next step. Other actions will all lead to the selection of a new opinion that will be added to the search results for the current issue. We call these *selection-actions*. If the current issue set is not empty, the searcher can choose the new-issue-action or one of the selection-actions.

If she chooses the new-issue-action, no new opinion has been selected, but the search results for the last issue are finalized and a new issue starts with an empty set. On the other hand, if the current issue set is empty, it means that the issue set has been initiated in the previous step. If that is the case, the only available actions consist of selection-actions, any of which will lead to the selection of the first opinion to be added to the current issue set. In both situations, if a selection-action is chosen, the searcher uses measures of similarity and quality to choose the next opinion to add to the search results. Here we use the same measures as defined in the covering algorithm where $b(o_i)$ and $\lambda(o_i)$ are measures of similarity and quality and the next opinion to add to the search results is selected based on a weighted sum of the similarity measure and the quality measure as:

$$o^* = \underset{o_i \in O/C(CFLT)}{\operatorname{argmax}} (B \times b(o_i)) + (1 - B) \times \lambda(o_i).$$

However the weight distribution denoted by B is determined based on the searcher's policy in the current state of the search. In fact, each of the selection-actions corresponds to a weighted distribution of the similarity and quality measures. Here, we assume that the searcher has only two settings of the weighted distribution.⁵ Accordingly, the selection-actions for any given search state consists of two actions: choosing the first weighting setting, B_1 ; or choosing the second weighting setting, B_2 .

Let us define a state for the search process as $s = (t, m, l, v)$, where t is the number of opinions already selected in the search process, m is the number of issues addressed and l is the number of the opinions selected for the current issue, based on the first weighting setting, and v is the number of the opinions selected for the current issue based on the second weighting setting. Accordingly, state space S is constructed as all possible values for the states.

Recalling the notation from the covering algorithm section, we keep track of the set of search results as $C(CFLT) = \bigcup_{k=1}^m C_k(CFLT)$, where $C_k(CFLT)$ represents the suggested issue set k . Also, recall $C_{\bar{k}}(CFLT)$ refers to the issue \bar{k} , which is being filled at the current step of the algorithm.

If the new-issue-action is denoted by a_0 and the selection-actions are denoted by a_1 and a_2 (indexed according to the weight settings they use to select the next opinion), we can define the action space for any given state $s \in S$ as: (1) if the current issue set is not empty, the corresponding action space is $A(s) = \{a_0, a_1, a_2\}$ and (2) if the current issue set is empty, then the corresponding action space is $A(s) = \{a_1, a_2\}$.

⁵ In the implementation of the adaptive algorithm, we use two weight distributions: $B_1 = 1/5$ and $B_2 = 4/5$.

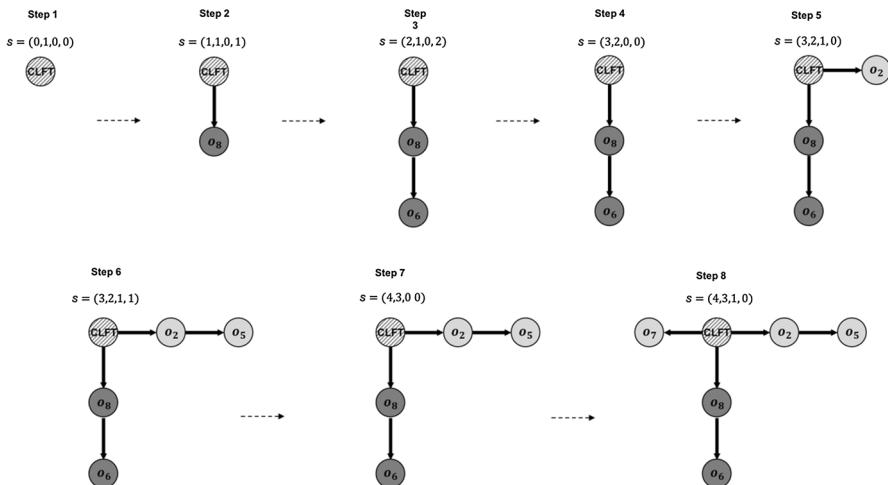


Fig. 3 Adaptive method

Imagine a short search episode with total capacity of 5 search results. The starting state, similar to any other search episode, is $s = (0, 1, 0, 0)$, representing no search results at the beginning. Also, an empty set, representing the current issue set is initiated to keep track of the search results for the first issue. Then the searcher chooses an action based on a policy among the first and the second weighting setting. Say, based on the policy, the searcher chooses to use the second weighting setting and accordingly ends up with the selection of o_8 (see Fig. 3), which is added to the first issue set. Then, since the first issue set is not empty anymore, the searcher needs to choose an action among three choices: (1) to start a new issue; (2) to continue the current issue and use the first weighting setting; (2) or to continue the current issue and use the the second weighting setting. As illustrated in Fig. 3, the second step may end up with the selection of o_6 to be added to the first issue set. Say in the next step, the policy is to start a new issue set. Accordingly, no new opinion is selected, but the first issue set is finalized and again an empty set, representing the current issue set, is initiated. This process continues to the point that the search results reaches the predetermined capacity, which in this case is 5.

At each step $s \in S$, if the policy selects an opinion using the first weighting setting, the next state is determined as $s^{new} = s + (1, 0, 1, 0)$. In the same way, if the policy is to choose the second weighting setting, we have $s^{new} = s + (1, 0, 0, 1)$. On the other hand, if the policy is to start a new issue, finalizing the last issue set and starting over with an empty set, we need to update the state as $s^{new} = (t, m + 1, 0, 0)$.

Accordingly, no matter the state we are in and what the policy is in that state, the next state is determined based on the current state and does not depend on the states before that. This means that the corresponding decision process holds the Markov

property and accordingly we can optimize the decision process, learning the optimal policy at each step, using a reinforcement learning procedure (here Q-learning).

To define an appropriate reward function, recall that the *CFLT* was formed by deleting the citation information from an opinion both in the training and testing data. Accordingly, an appropriate reward function would be based on whether, at each step, the recommended opinion o^* is in the list of the actual citations that were included in the opinion that was used to generate the *CFLT*. Then, if in a search state, a new-issue-action is selected by the policy, since no opinion is selected, the reward is zero. Otherwise a selection-action represented by the weight setting B is implemented based on the policy and it leads to a recommendation of:

$$o^* = \underset{o_i \in O/C(CFLT)}{\operatorname{argmax}} (B \times b(o_i) + (1 - B) \times \lambda(o_i))$$

to be cited, with the reward determined as:

$$R_{C(CFLT)}^\lambda = \begin{cases} 1, & \text{if } o^* \in \bar{C}(CFLT) \\ -0.1, & \text{if } o^* \notin \bar{C}(CFLT) \end{cases}$$

where $\bar{C}(CFLT)$ is the set of the opinions cited by the original opinion from which the *CFLT* was generated. After implementation of the adaptive algorithm, the resulting *Q-factor* is used as an approximation for the optimal search policy for a given *CFLT*, which means that the selection of opinions in the search results for the *CFLT* under this policy is an approximation of the citations in the underlying document.

4.4.4 Algorithm

Initialize $\bar{Q}(s, a)$, $\forall s \in S, a \in A(s)$, arbitrarily
 Set values for n, γ and α where $0 \leq \alpha \leq 1$ and $0 \leq \gamma \leq 1$
 Repeat (for each learning episode)
 Initialize $s = (0, 1, 0, 0)$
 Choose a from $A(s)$ using policy derived from \bar{Q} (ϵ -greedy)
 Repeat until $s[0] < n$
 Take action a , observe R , determine s'
 choose a' from $A(s')$ using ϵ -greedy policy derived from \bar{Q}
 $\bar{Q}(s, a) \leftarrow (1 - \alpha)\bar{Q}(s, a)$
 $+ \alpha[R + \gamma \max_{a' \in A(s')} \bar{Q}(s', a')]$
 $s \leftarrow s', a \leftarrow a'$
 Return $\bar{Q}(s, a)$

5 Implementation

5.1 Citation Prediction Task

Perhaps the best source of data available to the public concerning law search is the law itself, and specifically the corpus of judicial opinions. An observer of a judicial opinion can safely draw two conclusions concerning the sources of law that are cited in the opinion. One is that those documents were deemed relevant by the court to the legal questions discussed in the opinion. The second is that they were found through a process of law search by some actor (either a judge, a clerk, or a litigating attorney). In light of the predictive-strategic characteristic of law search, the parties' goal when conducting search can be thought of, at least partially, as predicting the citations in the final issued opinion.

Our first test for our algorithms is the ability of each algorithm to predict the citations within a judicial opinion based exclusively on the topic proportions within that document. For this task, we construct a CFLT from a randomly chosen opinion from the corpus. We then reduce the text to a distribution over the topics generated by the topic model. This opinion is the source document for the search. We then construct the search space without that opinion and, using the topic proportions in the source document, attempt to reconstruct the citations in that document (which have been held out) based on the search algorithms.⁶ The more accurate the reconstruction, the better the algorithm.⁷

We note two interesting features of this test. The first is that a null model in which citations are unrelated to document content can be rejected on intuitive grounds. It would be surprising indeed if documents with certain words (e.g., “interrogation”) are not more likely to include citations to certain documents (e.g., *Miranda v. Arizona*, 384 U.S. 436 (1966)) than documents without those words. Given the extremely strong theoretical reasons to believe that there is a relationship between the predictor and outcome variables, the question is not really whether substance is actually related to citation. The goal is to develop models capable of predicting relationships that are known to exist.⁸

⁶ This treatment of citation ignores information concerning the frequency or intensity of citation within an opinion. A single citation in a footnote that is part of a string of citations for a minor point is treated the same as an extensive discussion in the main body of the opinion. This loss of information could be corrected for in future models, which would likely increase predictive power.

⁷ Note that, for this test, all three algorithms are programmed to exclude from their results all opinions that are published after the date of the source document. Generally, the adaptive algorithm is always trained on a sub-set of prior in time opinions because those are the only ones that can be cited by a source document.

⁸ This approach is different than many standard inferential approaches take by social scientists, in which the relationship itself is being investigated, rather than the model.

5.2 Data

For data, we relied on the Free Law Project’s website CourtListener, which provides free public access to legal materials, including Supreme Court decisions. The CourtListener data includes 63,864 text files for decisions issued from the late 18th century until 2016. Each file includes the full text of all opinions along with metadata that includes the date issued, citation for the decision, and a list of citations within the opinions. We restricted our analysis to full decisions issued in the period 1946–2005 that contained at least one citation, for a total of 9575 files. Function words (such as “is” and “at”) were removed, as is common for many natural language processing applications. In addition, we removed all in-text citations so that they would not be erroneously picked up by the topic model.⁹ The citation network was generated using only Supreme Court opinions issued during the study period—any edges created by citation to appellate court opinions or secondary materials were not represented. The mean number of cited cases is 12, with a standard deviation of 11 and a maximum value of 121. A naive impact factor was constructed for the reduced corpus based on a weighted sum of the number of times a case was cited in another Supreme Court opinion and the number of times it was cited by the entire U.S. legal corpus (both estimates that can be derived via CourtListener data). We have verified our data against the supplemental information provided by Fowler et al. (2007) and Fowler and Jeon (2008) in their studies of the Supreme Court’s citation network.¹⁰

We use two measures to evaluate the performance of the algorithms. The first is recall, which is the size of the intersection of actual citation and predicted citations, divided by the number of actual citations. The second is precision, which is the size of the intersection of actual and predicted citations, divided by the number of predicted citations. Both of these measures can be estimated against any fixed number of predicted citations in the sense that for a given X we can compute recall@ X and precision@ X which are the recall and precision respectively in the top X returned predictions. For example, recall@10 is the recall estimate based on the top ten predictions; precision@20 is the precision estimate based on the top twenty predictions, and so on.

For the adaptive algorithm, we divided the opinions into ten clusters based on a partitioning of the search space.¹¹ The justification for the clustering is that search behavior may differ based on the substantive area of law being searched. Accordingly, the learning process is implemented on each cluster separately to capture potential differences in the depth/breadth tradeoff. For the training process, the source documents are first identified with its cluster, and then the training procedure is implemented on that cluster. The training and test sets are constructed at a 9:1 ratio. The reinforcement learner was implemented on the training set of each

⁹ For this procedure, all of the identifying information concerning a citation was removed, including the caption, date, and the Reporter information.

¹⁰ See <http://fowler.ucsd.edu/judicial.htm> for the supplemental information.

¹¹ The choice of ten clusters was arbitrary. We implemented the clustering by use of the common k-means algorithm (Huang 1998).

Table 1 Results of three search approaches

Method	Precision@10 (%)	Precision@20 (%)	Recall@10 (%)	Recall@20 (%)
Proximity	13	10	3	7
Covering	16	18	5	13
Adaptive	19	18	7	15

of the clusters separately, and the resulting search policy was used to predict the citation set of the associated test set. The proximity and covering algorithms were implemented on the same test datasets to facilitate inter-model comparisons. For the recall and precision estimates, “actual citation” is the set of cases that were cited in the opinion. The predicted citations are considered as a set with no order or ranking.

Table 1 compares the three models according to four performance measures, with estimates averaged over the ten clusters.

The adaptive algorithm generally performs better than the other two models. For purposes of comparison, the precision and recall estimates based on ten results are illustrative. In a group of ten results, the expected number of matches from the proximity algorithm to actual citations is roughly 1.3, the covering algorithm will generate 1.6, and the adaptive algorithm will generate roughly 2. The baseline probability drawing randomly from 10,000 cases is an order of magnitude lower, and so all three algorithms represent a fairly substantial improvement over random choice.¹² In general, precision declines as the number of results increase because the denominator (number of results) increases faster than the numerator (the number of correct results), in part because the lowest-hanging fruit can be expected to be picked up first. These precision and recall estimates compared the adaptive model post-training to the other two models.¹³

In interpreting these results, the proximity model, which draws from state of the art techniques in the information retrieval literature, is taken as the baseline. The covering and adaptive algorithms perform substantially better than this baseline. Note also that oftentimes there will be multiple plausible citations for an opinion author to choose from for any given legal proposition, meaning that there is a degree

¹² The baseline probability of drawing one of (say, generously) 100 cited cases from a pool of 10,000 is 1/100. Over ten draws, the likelihood of drawing at least one correct case is (a bit under) 1/10. The recall estimates take the number of cases cited as the denominator, and so for the recall@10 measure, the maximum possible performance (i.e., perfect precision) with thirty cases cited in a document would be 33% (ten out of the thirty cases). Accurately selecting one of thirty cases would lead to a 3.33% recall estimate, which is roughly that associated with the random walk model. The adaptive model performs a bit over twice as well (picking up 2.2 out of thirty citations, for example) with the covering model in between. It is worth noting that nearly a quarter of cited cases fall within the top fifty recommendations of the adaptive model.

¹³ The adaptive algorithm is initially seeded with random parameters to balance the breadth and depth of treatment for different legal issues, but very quickly (within twenty or so training sessions) determines values that allow it to substantially outperform the covering model.

of random choice in the selected citation. This implies a limit on precision, even if a sufficiently large recommendation set could in theory approach perfect recall.

5.3 Clustering

The takeaway from this analysis is that the three algorithms perform substantially better than random, and that the adaptive algorithm is the best performer. Structurally, the covering algorithm represents how search is carried out better than the proximity algorithm, with the adaptive algorithm having the additional advantage of learning how to balance depth and breadth of search from the data. This fact tells us something interesting about law search and citation in the Supreme Court: there is a tradeoff between exploration and the utility of adding additional citation. Because there are often many distinct legal issues within a Supreme Court opinion, drawing from distinct areas of law, it is not sufficient for a searcher to simply identify all things considered similar documents (as in the proximity model). Rather, search approaches that (essentially) decompose documents into smaller units that represent the many potential issues within an opinion accord better with the actual search-and-citation practices on the Court. The parameter values learned by the adaptive algorithm indicate how the balance between the need to explore multiple issues, and the need to delve into some issues in greater depth, is struck.

The clustering approach allows us to compare the parameter values that were learned in each cluster against the other. These parameters essentially estimate how the substantive area of law affects the breadth/depth balance, with some areas of law tilted toward an exploration of more legal issues, while others tend toward more deep analysis of a smaller number of legal issues.

The average number of predicted citations per issue are one way of expressing the breadth/depth tradeoff, with higher numbers indicating an emphasis on exploring issues in greater depth, and smaller numbers indicating an emphasis on exploring a larger number of issues. Table 2 reports the learned strategy in each of the clusters. We provide an intuitive label of the substantive area generated based on the opinions in the clusters.¹⁴

As can be seen here, there is some variation between substantive areas in this measure. Of course, at this early stage, these statistics should be understood as purely descriptive, and we do not know whether there is any out-of-sample validity—for example, whether the same correlations between subject area of average predicted citations per issue would persist over time or would occur in another related corpus, such as federal appellate opinions. Nevertheless, as a descriptive matter, it is the case that the reinforcement learner was able to improve its performance by differentiating how it made the breadth/depth tradeoff based on the region of the search space where a search was occurring.

¹⁴ To construct these intuitive labels, we used a topic-model based cluster labeling technique and then hand-labeled the clusters based on the top words in each associated topic. These labels are provided for the sake of convenience (see generally Tseng (2010); Tseng et al. (2006); Eikvil et al. (2015)).

Table 2 Comparison of breadth/depth tradeoff

Cluster label	Ave. # predicted citations
0 regulatory/evidence	2
1 private/securities/corporate	1.4
2 immigration/IP/int'l	2.5
3 admin/ratemaking	1.7
4 constitutional	2
5 water/enviro/energ	2.5
6 criminal	2
7 procedure	2
8 bankruptcy/property/credit	1.4
9 sentencing/6th Amend.	3.3

This measure can help formalize and test intuitions about how the substantive area of law affects the content of Supreme Court opinions. For example, one might speculate that opinions involving an area of law that is more specialized would tend to have a larger number of citations per issue than opinions concerning areas of law that tend to intersect with many different types of legal questions. This is a hypothesis that could, in theory, be tested using our measure. The formalization could also, in future work, facilitate comparison between different legal corpora. It may be the case, for example, that in a more specialized court (such as the U.S. Court of Appeals for the Federal Circuit) fewer issues are explored in greater depth. The methods discussed in this section provide a means of testing intuitions of this sort.

5.4 Human coded data

Despite the usefulness of the prior analysis, however, it is certainly the case that it is difficult to calibrate expectations concerning the performance of the algorithms. As mentioned above, outperforming random citation is a basic baseline, as is the proximity algorithm, which is drawn from state-of-the-art methods from the information retrieval literature. To provide an additional test that compares model performance to the underlying behavioral phenomenon of interest, we compare the model-derived results to those generated by research assistance.

The second test used to evaluate the models was based on data generated by law student research assistants. The goal of this analysis was to determine whether the models captured an intuitive notion of “legal similarity” as understood within the existing legal culture. The task requested of the students is akin to the hand coding of judicial opinions familiar from empirical legal studies Hall and Wright (2008). What is being coded here can be understood as a representation of the issues that are present in a case as a set of related cases. This data can be used in the same way that citation data was used in the prior analysis, to determine how well the algorithms successfully replicate the searches generated by the student researchers.

The research assistants participating in this project were second- and third-year students and so had completed at least the core first-year curriculum in which they were exposed to the system of common-law adjudication and the accompanying conceptual apparatus. Although not qualified to practice law, they can and do serve as law clerks during summer employment, participate in law school clinics (which includes representing clients), and take upper-level law classes.

For the coding procedure, seven research assistants each received ten randomly selected U.S. Supreme Court cases during the relevant time period (1946–2005) that served as the source documents. The students were then instructed to retrieve ten “similar” cases for each source case within that same time frame. The student instructions stated, “The definition of ‘similar’ is up to you” and students were told to “use whatever search tools you feel are appropriate”. The research assistants were free to use whatever search tools they were familiar with.

The goal of these instructions was to elicit relatively natural search behavior from the research assistants based on their own subjective understanding of legal similarity, and the typical tools of law search. By limiting the instructions on search and reducing the amount of description of legal similarity, we increased the diversity of the coded data generated by the students—quite the opposite of a typical coding exercise, where “inter-coder reliability” is prized. Instead, our procedure was meant to generate the amount of inter-coder overlap that arose purely from the shared understanding that the research assistants had, based on their common legal education and other experiences, free from coaching. To maintain the distinctness of the research assistants’ definitions and search approaches, they were requested not to discuss these matters with each other.

The seven students returned ten results for each of the ten source cases, resulting in 700 identified cases. Despite having different definitions of similarity, and using different resources to find their results, students’ cases overlapped substantially.

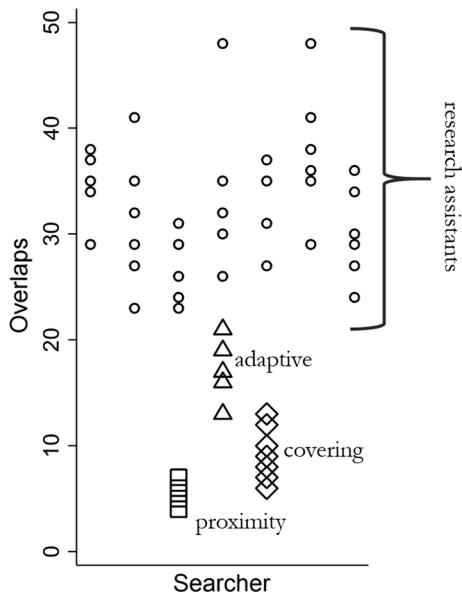
Each algorithm then executed the same task: given the source documents, they generated a list of other opinions based on their search methodology. An important difference between this task and the citation task is that the students looked at the entire universe of opinions, and not just opinions that were decided prior to the source document. To account for this different context, we applied the citation recovery approach discussed above for each algorithm, but treated all opinions—those issued both before and after the source document—as potential citations. Essentially, the algorithms were “asked” what citations would be found in the source document if it were issued today.¹⁵ Figure 4 reports how well the algorithm performed in anticipating the search results from the research assistants.

The search results for each research assistant were compared to the other researchers to determine the degree of overlaps.¹⁶ Each two research assistant pairs generates

¹⁵ More exactly, the hypothesized date of issue was the end of the study period, because more recent documents were not found in the corpus.

¹⁶ The number of overlaps is akin to the Jaccard coefficient—a common measure of similarity between sets—given that the students all generated the same number of returned results for each source document. (Romesburg 2004, see)

Fig. 4 Performance of search algorithms



a number of overlaps, generally between twenty and fifty (out of a total of one hundred possible overlaps). These are represented by each of the circles in Fig. 4. The search results of each of the algorithms were also compared to the research assistants, with an overlap generated for each research assistant-algorithm pair. The adaptive algorithm performed the best, with an average of seventeen overlaps. The covering algorithm averaged nine overlaps and the proximity algorithm averaged six. The average overlaps for the research assistants ranged from twenty-seven to thirty-eight.

There are two noteworthy observations from these results. The first is that there is a substantial spread in the performance of the three algorithms. The proximity model performs considerably above what would be expected from random guess, but would likely not be mistaken for a human searchers. The more sophisticated adaptive algorithm, on the other hand, performs surprisingly well, and, at least in some cases, its degree of overlap with the research assistants is nearly as great as the overlap of the research assistants with each other. It is worth remembering that these algorithms operate purely on term-frequency representations of all of the documents in the corpus, and that the citation information is stripped from the source document. For both of those reasons, the research assistants had useful information that was held back from the algorithms. With improvements, it is not unreasonable to expect that future iterations may begin to approach the amount of overlap found among human researchers.

6 Summary

In this article, we define and model law search as a strategic-predictive process and implement two related predictive tasks: a citation prediction task and a human-searcher prediction task. We use both of these tasks to test the proximity,

covering, and adaptive algorithms against each other and find that on both tasks, the adaptive algorithm performs best, and the covering outperforms the proximity algorithm. For these results, we learn that Supreme Court opinions generally tend to explore multiple issues in several different areas of law, rather than focus on a single issue. The better performance of the adaptive algorithm also provides some evidence that the breadth/depth tradeoff is different depending on the substantive area of the law being searched, given the greater ability of the adaptive model to calibrate itself based on the cluster in which a source document occurred. In addition, the overall performance of the adaptive algorithm on both tasks, and especially the human-searcher prediction task indicates that the search space and algorithms capture important features of human search. The characteristics of opinions that are embedded in this model of search—textual similarity and the citation network—are sufficiently related to the optimization function of natural searches that they achieve fairly impressive approximate prediction.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Alpern S, Gal S (2006) The theory of search games and rendezvous, vol 55. Springer, Berlin

Bench-Capon T, Araszkiewicz M, Ashley K, Atkinson K, Bex F, Borges F, Bourcier D, Bourgine P, Conrad J G, Francesconi E, Gordon T F, Governatori G, Leidner J L, David D, Lewis RPL, McCarty LT, Prakken H, Schilder F, Schweighofer E, Paul Thompson AT, Verheij B, Walton DN, Wyner AZ (2012) A history of AI and law in 50 papers: 25 years of the international conference on AI and law. *Artif Intell Law* 20:215–319

Blair DC, Maron ME (1985) An evaluation of retrieval effectiveness for a full-text document-retrieval system. *Commun ACM* 28(3):289–299

Blei DM (2012) Probabilistic topic models. *Commun ACM* 55(4):77–84

Blei DM, Lafferty JD (2007) A correlated topic model of science. *Ann Appl Stat* 1(1):17–35

Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet allocation. *J Mach Learn Res* 3:993–1022

Brin S, Page L (1998) The anatomy of a large scale hypertextual web search engine published version. *Comput Netw ISDN Syst* 30:107–117

Caprara A, Fischetti M, Toth P (2000) Algorithms for the set covering problem. *Ann Oper Res* 98:353–371

Carletta J (1996) Assessing agreement on classification tasks: the kappa statistic. *Comput Linguist* 22(2):249–254

Clark TS, Lauderdale BE (2012) The genealogy of law. *Political Anal* 20:329–350

Davies K (2007) The information-seeking behaviour of doctors: a review of the evidence. *Health Inf Libr J* 24(2):78–94

Diamond PA (1982) Aggregate demand management in search equilibrium. *J Political Econ* 90(5):881–894

Duan D, Li Y, Li R, Zhang R, Wen A (2012) Ranktopic: ranking based topic modeling. In: 2012 IEEE 12th international conference on data mining. IEEE, pp 211–220

Duan D, Li Y, Li R, Zhang R, Gu X, Wen K (2014) Limtopic: a framework of incorporating link based importance into topic modeling. *IEEE Trans Knowl Data Eng* 26(10):2493–2506

Dworkin R (1978) Taking rights seriously. Harvard University Press, Cambridge

Eikvil L, Jenssen TK, Holden M (2015) Multi-focus cluster labeling. *J Biomed Inform* 55:116–123

Finnis J (2011) Natural law and natural rights. Oxford University Press, Oxford

Fisher KE, Erdelez S, McKechnie L (2005) Theories of information behavior. *Information Today Inc, Medford*

Fowler JH, Jeon S (2008) The authority of Supreme Court precedent. *Soc Netw* 30:16–30

Fowler JH, Johnson TR, Spriggs JF, Jeon S, Wahlbeck PJ (2007) Network analysis and the law: measuring the legal importance of precedents at the U.S. Supreme Court. *Political Anal* 15:324–346

Fuller L (1958) Positivism and fidelity to law—a reply to Professor Hart. *Harv Law Rev* 71:630–663

Green L (1996) The concept of law revisited. *Mich Law Rev* 94(6):1687–1717

Hall MA, Wright RF (2008) Systematic content analysis of judicial opinions. *Calif Law Rev* 96(1):63–122

Hantzis CW (1987) Legal innovation within the wider intellectual tradition: the pragmatism of Oliver Wendell Holmes Jr. *Northwest Univ Law Rev* 82(3):541–595

Hart H (1961) The concept of law. Clarendon Press, Oxford

Hemminger BM, Lu D, Vaughan KTL, Adams SJ (2007) Information seeking behavior of academic scientists. *J Am Soc Inf Sci Technol* 58(14):2205–2225

Höchstötter N, Lewandowski D (2009) What users see—structures in search engine results pages. *Inf Sci* 179(12):1796–1812

Holmes OW (1897) The path of the law. *Harv Law Rev* 10:457–478

Huang Z (1998) Extensions to the k-means algorithm for clustering large data sets with categorical values. *Data Min Knowl Discov* 2(3):283–304

Joachims T (1998) Text categorization with support vector machines: learning with many relevant features. In: European conference on machine learning. Springer, pp 137–142

Jones M, Sugden R (2001) Positive confirmation bias in the acquisition of information. *Theory Decis* 50(1):59–99

Kohn MG, Shavell S (1974) The theory of search. *J Econ Theory* 9(2):93–123

Langville AN, Meyer CD (2011) Google's pagerank and beyond: the science of search engine rankings. Princeton University Press, Princeton

Leckie GJ, Pettigrew KE, Sylvain C (1996) Modeling the information seeking of professionals: a general model derived from research on engineers, health care professionals, and lawyers. *Libr Q* 66(2):161–193

Ledley RS, Lusted LB (1959) Reasoning foundations of medical diagnosis. *Science* 130(3366):9–21

Leibon G, Livermore MA, Harder R, Riddell AB, Rockmore DN (2018) Bending the law: geometric tools for quantifying influence in the multinetwor of legal opinions. *Artif Intell Law* 26(2):145–167

Liu Y, Xu S (2017) A local context-aware lda model for topic modeling in a document network. *J Assoc Inf Sci Technol* 68(6):1429–1448

Livermore MA, Riddell AB, Rockmore DN (2017) The Supreme Court and the judicial genre. *Ariz Law Rev* 59(4):837–901

Luhn HP (1957) A statistical approach to mechanized encoding and searching of literary information. *IBM J Res Dev* 1(4):309–317

Manning C, Raghavan P, Schütze H (2010) Introduction to information retrieval. *Nat Lang Eng* 16(1):100–103

Meho LI, Tibbo HR (2003) Modeling the information-seeking behavior of social scientists: Ellis's study revisited. *J Am Soc Inf Sci Technol* 54(6):570–587

Mooers CN (1950) The theory of digital handling of non-numerical information and its implications to machine economics. Zator Co., Boston

Mortensen DT, Pissarides CA (1994) Job creation and job destruction in the theory of unemployment. *Rev Econ Stud* 61(3):397–415

Page L, Brin S, Motwani R, Winograd T (1999) The PageRank citation ranking: bringing order to the web. Technical report. Stanford InfoLab, Stanford. <http://ilpubs.stanford.edu:8090/422/>

Qiao Q, Beling PA (2011) Inverse reinforcement learning with Gaussian process. In: American control conference (ACC), pp 113–118

Quinn KM, Monroe BL, Colaresi M, Crespin MH, Radev DR (2010) How to analyze political attention with minimal assumptions and costs. *Am J Political Sci* 54(1):209–228

Ramos-Fernández G, Mateos JL, Miramontes O, Cocho G, Larralde H, Ayala-Orozco B (2004) Lévy walk patterns in the foraging movements of spider monkeys (*Ateles geoffroyi*). *Behav Ecol Sociobiol* 55(3):223–230

Riddell AB (2014) How to read 22,198 journal articles: studying the history of German studies with topic models. In: Erlin M, Tatlock L (eds) *Distant readings: topologies of German culture in the long nineteenth century*. Camden House, Rochester, pp 91–114

Rissland EL (1990) Artificial intelligence and law: stepping stones to a model of legal reasoning. *Yale Law J* 99(8):1957–1981

Roberts ME, Stewart BM, Tingley D, Lucas C, Leder-Luis J, Gadarian SK, Albertson B, Rand DG (2014) Structural topic models for open-ended survey responses. *Am J Political Sci* 58(4):1064–1082

Robinson MA (2010) An empirical analysis of engineers' information behaviors. *J Am Soc Inf Sci Technol* 61(4):640–658

Romesburg C (2004) *Cluster analysis for researchers*. Lulu Press, Morrisville

Saracevic T (1996) Relevance reconsidered. In: *Proceedings of the second conference on conceptions of library and information science (CoLIS 2)*. ACM, New York, pp 201–218

Schauer F, Wise VJ (1997) Legal positivism as legal information. *Cornell Law Rev* 82(5):1080–1110

Schauer F, Wise VJ (2000) Nonlegal information and the delegalization of law. *J Legal Stud* 29(S1):495–515

Smeaton AF (1999) Using NLP or NLP resources for information retrieval tasks. In: Strzalkowski T (ed) *Natural language information retrieval*. Springer, Dordrecht, pp 99–111

Sparck Jones K (1972) A statistical interpretation of term specificity and its application in retrieval. *J Doc* 28(1):11–21

Tseng YH (2010) Generic title labeling for clustered documents. *Expert Syst Appl* 37:2247–2254

Tseng Y, Lin C, Chen H, Lin Y (2006) Toward generic title generation for clustered documents. In: *Information retrieval technology*

Wells CP (1993) Holmes on legal method: the predictive theory of law as an instance of scientific method. *South Ill Univ Law J* 18:329–345

Wilkinson MA (2001) Information sources used by lawyers in problem solving: an empirical exploration. *Libr Inf Sci Res* 23(3):257–276

Wilson TD (1981) On user studies and information needs. *J Doc* 37(1):3–15

Zhang Z, Li R, Li Y, Gu X (2018) Rank-integrated topic modeling: a general framework. In: *Asia-Pacific Web (APWeb) and web-age information management (WAIM) joint international conference on web and big data*. Springer, pp 16–31

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.