

Data-centric and logic-based models for automated legal problem solving

L. Karl Branting¹

Published online: 2 March 2017
© Springer Science+Business Media Dordrecht 2017

Abstract Logic-based approaches to legal problem solving model the rule-governed nature of legal argumentation, justification, and other legal discourse but suffer from two key obstacles: the absence of efficient, scalable techniques for creating authoritative representations of legal texts as logical expressions; and the difficulty of evaluating legal terms and concepts in terms of the language of ordinary discourse. Data-centric techniques can be used to finesse the challenges of formalizing legal rules and matching legal predicates with the language of ordinary parlance by exploiting knowledge latent in legal corpora. However, these techniques typically are opaque and unable to support the rule-governed discourse needed for persuasive argumentation and justification. This paper distinguishes representative legal tasks to which each approach appears to be particularly well suited and proposes a hybrid model that exploits the complementarity of each.

Keywords Logical representation · Empirical methods · Prediction · Hybrid reasoning

1 Introduction

Automation of legal reasoning and problem solving has been a goal of Artificial Intelligence (AI) research from its earliest days. The logical character of legal texts, norms, and argumentation suggests that formal logic can function as a model capable of representing legal rules, case facts, and inference based on the application of rules to facts (Allen 1957). Development of general-purpose techniques for automated inference over logical rules (e.g., Newell et al. 1959), provided a method for implementing this approach in the form of logic-based legal reasoning systems.

✉ L. Karl Branting
lbranting@mitre.org

¹ The MITRE Corporation, McLean, VA, USA

Initially, these systems were purely notional, without any actual implementation, e.g., Mehl (1958) (tax law), McCarty (1977) (tax law), Meldman (1975) (torts). But by the 1980s, a number of researchers had implemented working systems based on manually created logical representations of rules e.g., Sergot et al. (1986) (British Nationality Act), Peterson and Waterman (1985) (product liability). With the optimism characteristic of AI in that era, researchers of the day were confident that the transformation of the law by automated systems was imminent. Broad adoption of legal AI systems never occurred, however. Instead, AI and law remained a niche research area with little practical impact (see, e.g., Oskamp and Lauritsen 2002). Institutional resistance to automation played a role in this phenomenon, but a more fundamental obstacle has been the difficulties of scaling the logic-based approach to the dimensions of complex, dynamic, real-world legal systems. These difficulties arise from two fundamental technical challenges: efficient and verifiable representation of legal texts as logical expressions; and evaluation of legal predicates based on facts expressed in the language of ordinary discourse.

In recent years, a new area of research has emerged that performs legal problem solving using knowledge induced from collections of legal documents or other large data sets. Emergence of this *data-centric* approach coincided with development of techniques for statistical analysis of very large data sets, including large text corpora. By making it possible to rapidly process large data sets, such as all the cases of a given court or all the statutes of a given jurisdiction, these techniques can provide new perspectives on the properties of legal systems, such as the graph topologies of citation networks, the evolution of legal doctrines, and the probability of decisions or other case outcomes. This has led to a proliferation of legal technology companies (631 as of this writing, by one estimate),¹ most focused on lucrative applications in litigation support.

The rapid growth of the data-centric approach raises important issues about its relationship to logic-based approaches: does the data-centric approach render previous approaches obsolete, or are there distinct tasks for which one approach is clearly preferable or tasks for which the two approaches are complementary?

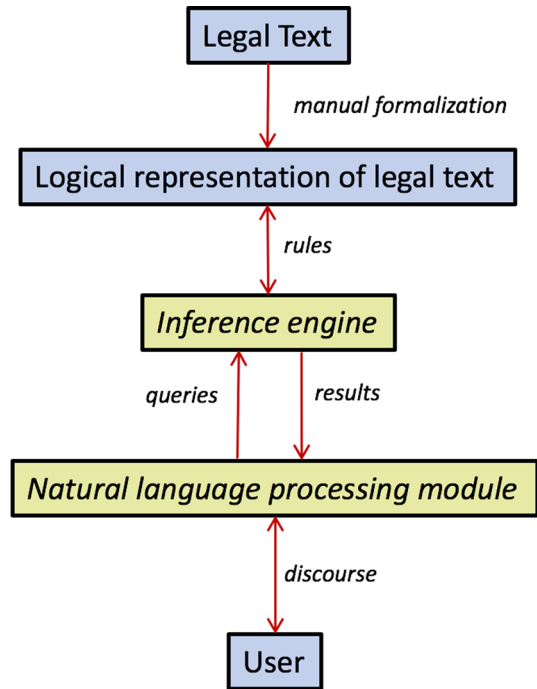
Section 2 reviews the capabilities and challenges of the logic-based approach, and Sect. 3 presents four categories of data-centric approaches. To place the comparison of logic-based and data-centric approaches in context, Sect. 4 identifies which approach appears most appropriate for each of a number of categories of legal tasks. The paper concludes with a proposal for a hybrid approach that exploits the complementarity of the two approaches.

2 The logic-based approach and its discontents

A typical architecture of the pioneers of AI and law is shown in Fig. 1 (Meldman 1975). Starting from the top, legal text is manually formalized in a logical representation. This representation is accessed by an inference engine (e.g., a

¹ <http://tech.law.stanford.edu/> [Accessed: 12 November 2016].

Fig. 1 Jeffrey Meldman's 1975 architecture for a conversational legal reasoning system contained the key elements used in systems for legal analysis in subsequent decades



resolution theorem prover or its equivalent) in response to queries generated by a natural-language processing in response to discourse with the user.

The earliest implemented systems, such as the British Nationality Act as a Logic Program (Sergot et al. 1986) typically used a slightly simpler model, illustrated in Fig. 2. The omission of the natural-language discourse processing module means that queries are posed, and results provided, in the form of Horn clauses, a subset of First-Order Predicate Calculus.

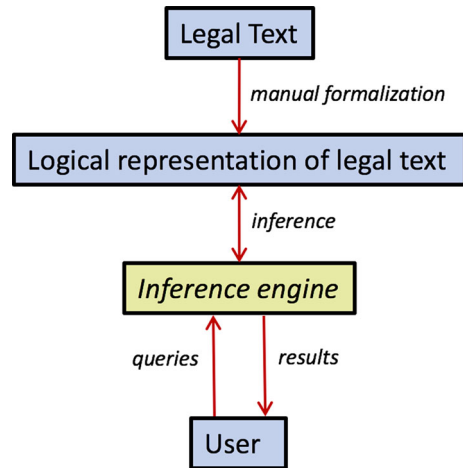
2.1 The challenge of representing texts in logic

Logic-based systems are capable of inferring legal conclusions from appropriately formalized rules and facts. However, as mentioned briefly above, two key obstacles have impeded their acceptance: the challenge of efficiently and verifiably representing legal texts in the form of logical expressions; and the difficulty of evaluating legal predicates from facts expressed in the language of ordinary discourse. Each of these challenges is discussed in turn.

2.1.1 Fidelity and authoritative

A basic assumption of logic-based approaches is that it is possible to create a logical formalization of a given legal text that is a faithful and authoritative expression of the text's meaning.

Fig. 2 An architecture for legal logic programming



Three elements of this assumption can be distinguished: (1) that legal texts have a determinate logical structure; (2) that there is a logical formalism sufficiently expressive to support the full complexity of legal reasoning on these texts, and (3) the logical expression can achieve the same authoritative status as the text from which it is derived.

The assumption that legal texts have a determinate logical structure is at odds with the observation that a single logical proposition may be expressed in multiple distinct alternative natural-language constructions, and that a single natural-language construction may have multiple logical interpretations.² There have been multiple efforts to develop conventions for “normalized” statutory language such that a statute enacted in such a representation would have an unambiguous logical structure, e.g., Allen (1982), but there have been only brief episodes in which legislative bodies were persuaded to adopt such an approach, e.g., Gray (1987). In the absence of such a convention, a logical representation of a statutory rule is simply one interpretation among many possible interpretations. The difficulty of reaching a consensus about the logical interpretation of statutes is exacerbated by dissimilarity between the surface form of statutory texts and the text’s underlying logical structure. Correct interpretation of a given legal provision often requires knowledge of other parts of an overall statutory scheme. Efforts to directly map from even the most carefully structured statutes founder on conventions of statutory writing, such as incorporation by reference and counterfactuals, that have no direct counterpart in the syntax of typical formal logics (Sergot et al. 1986; Bench-Capon and Coenen 1992; Bench-Capon and Gordon 2009).

Even if a determinate logical interpretation of a legal text can be found, it must be expressed in a formalism that supports the inferences needed to achieve the legal task for which the text has been formalized. There is an extensive

² This many-to-many mapping is the result, at least in part, of the many pragmatic functions of natural languages beyond simply expressing propositional content.

literature on logics specialized for the deontic, defeasible, and temporal characteristics of legal reasoning and argumentation. However, there is no consensus as to the most appropriate logic for a given task; instead, the choice is an exercise in judgment.

Finally, the institutional imprimatur of a validly enacted legal text does not, in general, extend to a formal logical representation of that text. Unless a legislature or other law-enacting institution actually enacts a legal rule in logical, as opposed to textual, form or decrees that a logical expression is to be treated as equivalent to a legal text, a logical representation will remain simply one interpretation of the legal text among many. Attorneys could argue that a given logical interpretation is incorrect and that the conclusions of an automated legal reasoning system based on that interpretation are therefore also incorrect.

In summary, the assumption that it is possible to create a logical formalization of a given legal text that is a faithful and authoritative expression of the text's meaning is not justified. Instead, one can create various alternative logical representations corresponding to alternative interpretations of the legal text. The particular logical formalism can be chosen based on a trade-off between simplicity and optimality for a given legal task, but the representation itself, and therefore all inferences drawn from it, has no officially sanctioned status.

2.1.2 Practicality

Representing legal texts in formal logic is a highly specialized skill, requiring insight into the subtleties of both legal writing and logic. The process is both labor-intensive and error-prone. Because of these factors, formalization of legal texts in logic has not yet been shown to be scalable to large or frequently changing legal texts, but instead has typically been limited to boutique applications in which there is a small set of relatively stable rules and the investment required for representation in legal form can be amortized across a great many users.³

2.2 The language gap

A second challenge to logic-based approaches arises from the gap between the abstract legal terms-of-art and ordinary discourse. Even if legal rules are formalized with perfect fidelity, the terms in the rules are typically impossible for a layperson to interpret. Contrast, for example, the description that a layperson might use to pose a legal query from the terms into which an attorney would reformulate this query:

Lay Person

The bookkeeper at my job didn't give me my paycheck last month. Instead, she signed my name on it, cashed it, and left town. I don't know where she went. I asked my boss to give me a new check for my salary, but he said that he had already paid me once and that he didn't have to pay me again. He says

³ For an example of an attempt to scale up formalization of legislation, see (van Engers and Nijssen 2014).

that if anyone owes me the money, it is the bookkeeper. Is he right that he doesn't owe me my wages anymore?

Attorney

Under Article 3 of the Uniform Commercial Code, is a payer's obligation to a payee discharged by a negotiable instrument if the negotiable instrument is paid to a third party over a forged endorsement?

A legal reasoning system for negotiable instruments under Article 3 of the Uniform Commercial Code (UCC) would consist of logical expressions containing terms like "discharge of obligation," "forged endorsement," and "payer's obligation," reflecting the language of the UCC. A system for negotiable-instruments law with the architecture depicted in Fig. 2 could be usable by attorneys but would be incomprehensible to a lay person unfamiliar with the UCC's technical language.

2.2.1 Open texture

The gap between legal terms of art and ordinary discourse goes much deeper than lay persons' unfamiliarity with legal terminology. Legal concepts are characterized by "open texture," e.g., vagueness ("manifestation of assent"), variable standards ("due care," "a reasonable time"), defeasibility (the rule against motor vehicles in the park may not apply to ambulances), and ambiguity ("my son" used in the will of a person with 2 sons). These complexities are not limited to legal concepts (Lakoff 1987), but logic-based legal reasoning systems nevertheless require some approach to assigning truth values to legal terms that function as predicates in logical expressions.

Various approaches to open-texture have been explored. One approach is to apply *canons of construction*, which are heuristics for interpretation of legal provisions dating back to Roman law and elaborated over centuries of legal practice. Applying canons of construction is a skilled legal activity, and attorneys can and do argue about the applicability of competing canons (Brudney and Ditslear 2005). To the best of my knowledge, no AI project has attempted to formalize the process of selecting and applying canons of construction.

A second approach to open texture is to restrict the scope of the system to *prima facie* reasoning, that is, assume that the facts to which the formalized rules are applied constitute a clear case (Branting 2000a). In this approach, the output of the system would come with an implicit proviso that the open-texture of the legal terms might create counterarguments or even negate or invalidate the output. By contrast, a *fully agnostic approach* would generate a set of alternative outputs corresponding to alternative truth-value assignments to each legal predicate. This would produce an unwieldy proliferation of alternative solutions.

Starting in the late 1980s, the dominant approach to open texture has been *Case-Based Reasoning (CBR)*, which involves evaluating whether a legal predicate is satisfied by a given set of facts by comparing those facts to exemplars for which the truth value of the predicate is known, either as a matter of institutional imprimatur, as with legal precedents, or through consensus, as with prototype cases. Fact situations that resemble both positive and negative instances of a legal predicate give rise to competing analyses under this approach, but unlike the fully agnostic

approach, differences in degrees of similarity give rise to a natural ordering of the persuasiveness of case-based arguments.

While CBR is a principled approach to the problem of reasoning about open-textured legal terms and has been demonstrated to lead to accurate legal analyses, current and previous CBR systems take as input not natural language, but formal representations painstakingly engineered to facilitate computational comparison between cases, such as factors (Alevén and Ashley 1996) or semantic nets (Branting 2000b).⁴ As with representing legal rules in logic, representing cases in representations specialized for CBR requires a specialized combination of legal and AI skills, is laborious and error-prone, and has therefore not been scalable.⁵

2.3 Institutional approaches

Despite many successful small-scale demonstrations of legal problem solving by logic-based systems, the challenges of establishing fidelity and authoritativeness, the necessity of large amounts of highly specialized human expertise, and the gap between legal language and ordinary discourse has prevented large-scale development or adoption of logic-based legal systems. However, the first of these three problems could be addressed by institutional changes under which rule-enacting bodies (e.g., legislatures or regulatory bodies) enacted rules in logical, computer-executable form. This could be accomplished either of two ways: (1) through an officially sanctioned text-to-logic compiler together with a rule editor that restricted the syntax and vocabulary of rules to expressions parseable by the rule editor; or (2) enacting rules in logical, computer-executable form together with an officially sanctioned logic-to-text realizer. By producing officially sanctioned rules, either approach could open the door to improved legislative and regulatory consistency (e.g., logical evaluation of consistency with existing rules or automated evaluation with prototype cases intended to be covered by new rules), improved voluntary compliance, since executable rules could be used by members of the public to test alternative actions or situations, and improved administration, since the consistency of procedures, forms, websites, and other entities that mediate between rules and citizens could be subject to formal verification.

Unfortunately, the idea of enactment of rules in executable form is, for the most part, incomprehensible to rule-making bodies. Experiments with enactment of rules in executable form are likely to come only when there are well-publicized projects that demonstrate the benefits of automated interpretation of legal rules. Moreover, this institutional reform does not address the language gap; even if a set of logical rules is an authoritative and faithful representation of the intention of the rule-creating body and even if there are positive and negative exemplars of the legal concepts in those rules, there will often remain a gap between problem-fact descriptions and the language of rules or exemplars.

⁴ While there have been experiments in automating the extraction of predefined factors from curated collections, e.g., Brüningshaus and Ashley (2001), this work depended on a pre-existing set of factors manually engineered for each domain.

⁵ For a discussion of other approaches to open texture, see Bench-Capon and Visser (1997).

3 The rise of data-centric legal systems

Recent advances in both Human Language Technology (HLT) and techniques for large-scale data analysis (“Big Data”) have vastly increased capabilities for automated interpretation of legal text. These advances have coincided with a rapid expansion of interest in automated processing and understanding of legal texts on the part of industry, government agencies, court personnel, and the public. The rapid development of large-scale data analytics has attracted many researchers into legal text analysis who have little familiarity with the aims or methods of more traditional logic-based approaches.

As notionally illustrated in Fig. 3, data-centric approaches can, under some circumstances, finesse the two key impediments to logic-based systems: rule formalization and the language gap. However, as described below, data-centric techniques typically address somewhat different legal tasks than logic-based approaches; in general, they are better-suited for tasks that involve exploitation of knowledge latent in legal documents and corpora or that require empirical or statistical characterization of a case than tasks that depend on generation or analysis of highly rule-governed discourse or documents.

Three areas of recent data-centric research can be distinguished: (1) case-oriented, (2) document-oriented, and (3) corpus-oriented. Each will be described in turn, followed by a discussion of opportunities for hybrid logical/data-centric models.

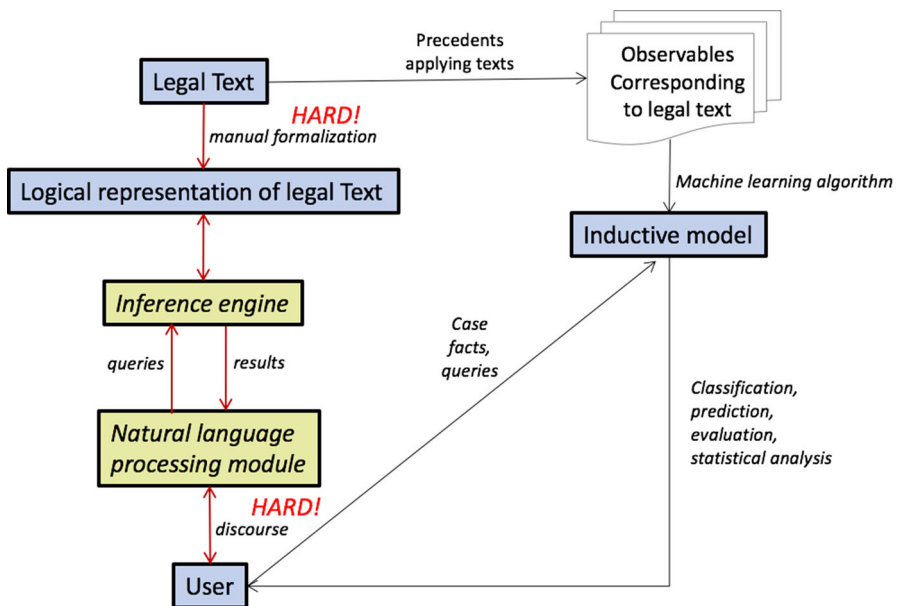


Fig. 3 A data-centric approach can finesse the most challenging impediments to logic-based approaches for some classes of legal tasks

3.1 Case-oriented approaches

One strand of data-centric legal technology focuses on the significant characteristics of cases considered as a whole, such as duration, costs, and potential awards or punishments, and probability of success of claims, motions, or other pleadings. An area of particularly active commercial activity is *litigation assistance*, that is, providing information to improve the probability of success at trial. Often such information takes the form of prediction of case events, such as the probability of success of a given action as a function of particular attorneys, judges, courts, etc. Recent startups marketing such predictive services include Lex Machina,⁶ LexPredict,⁷ and Premonition.⁸ These litigation-assistance tools are in the tradition of a long history of proprietary prediction systems in the insurance industry, e.g., Peterson and Waterman (1985).

Prediction is often an important component of the legal advice provided by attorneys. For example, an attorney's advice to a client about whether a certain tax shelter is legal or about the likelihood that a judge will find an argument persuasive may be based on a combination of weighing the relative strengths of alternative competing legal arguments and predicting the behavior of judges and other participants in the legal disputes based on the attorney's experience. Data-centric prediction can thus be viewed as an automation of the component of an attorney's expertise that is based on experience. The driving force behind adoption of predictive systems is the expense of corporate litigation. Even a small, competitive advantage can be extremely valuable in, e.g., a high-stakes patent-infringement case. Factors shown to be predictive include factors unrelated to the merits of the case, such as the nature of the suit, attorneys, forum, judge, and parties, (Surdeanu et al. 2011) as well factors that may be related to the merits, such as lexical features, events, narratives, and procedural history [see discussion in Aletras et al. (2016)].

A limitation of predictive analytics is that a probability unaccompanied by a legal justification is not useful for actually producing legal arguments, documents, or discourse, although it may guide the decisions about what argument, document, or discourse to create. The opacity of statistical or machine-learning-based decision algorithms has led the EU to enable a "General Data Protection Regulation" that guarantees the right to "an explanation of the decision reached after (algorithmic) assessment" (Goodman and Flaxman 2016a).

On the positive side, predictive systems have the potential to compensate for institutional biases that sometimes favor unnecessary litigation. A study by Loftus and Wagenaar showed that attorneys systematically overestimated their probability of success at trial (Loftus and Wagenaar 1988). More accurate information about the expected costs and rewards of litigation could help clients make more informed decisions about whether to go to trial or to settle.

⁶ <https://lexmachina.com/> [Accessed: 27 November 2016].

⁷ <https://lexpredict.com/> [Accessed: 29 November 2016].

⁸ <https://premonition.ai/> [Accessed: 27 November 2016].

3.2 Document-oriented approaches

A second strand of data-centric legal technology focuses on analysis of individual documents.

3.2.1 Information extraction

Information extraction is the process of identifying *named entities* such as places, persons, organizations, dates, claims, etc., as well as extracting more complex information, such as events and narratives. A relatively mature technology, information extraction has been applied to identify named entities in legal texts (Dozier et al. 2010), legal claims (Surdeanu et al. 2010), and descriptions of events and narratives giving rise to legal claims (Vilain 2016). Most approaches to information extraction require annotated corpora to train the extraction models.

3.2.2 Automated summarization

An important work product in legal research consists of summaries of case facts, decisions, and other legal documents. Case summaries are an important part of the services delivered by legal publishers such as Lexis-Nexis and Thomson Reuters. Traditionally, summarization is performed by junior personnel at law firms or editorial staff at legal publishers.

While automated summarization has been an active research area for many years, the most popular approach has been *extractive summarization* in which the summary consists of a subset of the text in the original document. While this approach has been applied to patent-claim analysis (Brügmann et al. 2014), it appears in general to work much better for expository texts such as news article, which have an inverted pyramid structure that separates summary text from explanatory and background text, than for legal text genres characterized by complex discourse structures.

More recent work focuses on *abstractive summarization*, an alternative approach that selects content appropriate for a summary and combines it into coherent text. One recent approach to abstractive summarization involves *Deep Learning*, i.e., neural network techniques (Rush et al. 2015). Legal Robot is a Silicon Valley startup that is applying Deep Learning to legal text summarization.⁹ An alternative data-centric approach harvests parenthetical descriptions that judicial opinions routinely place after a citation to another case. The parenthetical descriptions of a given case by citing cases are assembled into a single summary. This approach is under development at the startup CaseText¹⁰ [as described by Pablo Arrendondo and Ryan Walker in Branting and Conrad (2016)].

⁹ <https://www.legalrobot.com/> [Accessed: 2 December 2016].

¹⁰ <https://casetext.com/> [Accessed: 2 December 2016].

3.2.3 Parsing statutory text

The bottleneck of manual representation of legal texts as logical rules, described above, could be ameliorated if legal texts could be automatically parsed into executable logical rules. Many factors make such parsing challenging, including the distinctive vocabulary of law, the unique characteristics of legal texts (such as the extreme length and syntactical complexity of legal sentences and deeply nested bulleted lists), and a lack of consensus about the target representation into which legal texts should be compiled.

However, recent work has suggest that, with the help of resources tailored for this genre, extension of standard tools of the HLT community should make the conversion of statutory text to machine-interpretable rules increasingly feasible (Buabuchachart et al. 2013). Recent IARPA-sponsored work demonstrated how statutory text can be parsed into a standard rule-markup representation, with the biggest single apparent barrier being that there are no tree banks comparable to the Penn Tree Bank¹¹ for statutory text (Morgenstern 2014). For other recent efforts at parsing statutory text, see Wyner and Peters (2011).

3.2.4 Predictive retrieval and form completion

Legal information retrieval is an essential part of modern legal practice. Two recent projects demonstrate how modern text analytics can permit real-time interactive retrieval of legal texts to operate predictively in the form of cognitive assistants.

Schwartz et al. (2015) have developed a Microsoft Word plug-in that uses Latent Dirichlet Allocation (LDA) to identify and suggest prior authorities on the same topic as the paragraph that the user is typing. The system is an intelligent citation assistant that constantly updates its topic model of the document that the user is writing and proactively searches for documents whose topic-model similarity indicates relevance to the user's topic.

Similarly, *Proactive Legal Information Retrieval and Filtering*¹² is a system by Brian Carver and Yi Zhang at the University of California at Berkeley that infers models of the users' interests from the query histories and uses these models to anticipate the user's next queries. Both projects illustrate how user modeling through topic or goal detection can permit a system to anticipate, cache, and suggest potential resources for achieving users' document-oriented goals.

A related idea is to anticipate form-fillings actions to reduce the time and effort necessary to complete forms. There is often redundancy among fields of documents such that the value of the contents of one field may be predictable based on the contents of other fields. Under these circumstances, supervised concept learning can be used to infer values for and pre-populate text fields based on values in other fields and on previous values.

¹¹ <http://www.cis.upenn.edu/~treebank/home.html>. The Penn Treebank Project annotates naturally-occurring text for linguistic structure.

¹² <http://citris-uc.org/ctrs-groups/project/proactive-legal-information-retrieval-and-filtering/>.

An early prototype of this approach was used in a cognitive assistant for routine form completion described in Hermens and Schlimmer's 1994 "Machine-learning apprentice for the completion of repetitive forms" (Hermens and Schlimmer 1994). Hermens and Schlimmer found that the best-performing algorithm reduced the number of keystrokes required for form completions by 87% as compared to the no-learning method.

3.3 Corpus-oriented approaches

The third strand of data-centric legal technology focuses on the properties of entire collections of legal texts, including network structures, temporal and sequential characteristics, and content distribution.

3.3.1 Network analysis

Individual legal rules seldom exist in isolation, but instead typically occur as components of broader statutory, regulatory, and common-law frameworks consisting of numerous interconnected rules, regulations, and rulings. Such systems can be modeled as *networks* in which *nodes* represent rules, cases, or textual divisions of legal authorities, such as sections of regulations or statutes, and *edges* represent relationships between nodes, such as citations, amendment history, or semantic relations, such as topical similarity.

There is an active legal network analysis research community, including workshops on Network Analysis in Law (NAiL 2013; NAiL 2 2014; NAiL 3 2015). A particularly active area of network analysis involves citation networks of judicial decisions, such as U.S. Supreme Court decisions. Network metrics such as degree, betweenness, and centrality can provide insights into the importance or influence of particular decisions, the evolution of legal doctrine, and the behavior of, and alliances among, judges (Chandler 2005).

Statutory citation network visualizations can be useful for multiple stakeholders in the legal systems. For legislative drafters, regulated agencies, and citizens, citation networks can provide a mechanism for visualizing complex structures at multiple scales, including both the global structure and substructures relevant to particular topics or domains (Bommarito and Katz 2009).

For agencies, typical applications include identifying subgraphs of citation networks that correspond to the rules governing a particular scenario or domain, e.g., the authority relationships among agencies and officials responsible for responding to public health emergencies, or the network of regulations that might be applicable to drone operations. For agencies governed by regulations that are frequently amended, a network representation can help visualize the effects of amendments as changes propagate through the citation network.

For legislative drafters, network visualizations can facilitated "copy-and-edit" operations whereby an existing regulatory scheme, e.g., for manned aircraft, can be adapted for a different but related area, e.g., unmanned aircraft, and can help identify provisions that should be modified to maintain consistency in response to a given amendment.

Citation networks are relatively easy to create. In common-law decision networks, citations to prior cases and other legal authorities tend to be expressed in a consistent fashion, with clear semantics (e.g., affirmed, reversed, followed, etc). Similarly, citations between regulatory and statutory provisions are typically syntactically clear, but the semantics of these links are typically not standardized (Sadeghian et al. 2016).

However, the relationships most important for understanding the structure and dynamics of statutory and regulatory systems are not necessarily the citation links. Semantic or topical similarity, or applicability to similar legal entities, for example, may be more important in many applications. For example, if one would like to derive, or verify the correctness of, an administrative process governed by a set of rules, the citation structure of the rules may be a weak proxy for the procedural or administrative relationships created by the rules. Improved methods for deriving such implicit relationships is an important issue for legal network research.

3.3.2 Argumentation mining

Argumentation mining consists of a set of techniques for detecting and extracting discourse elements of individual texts or collections of texts related to justification, disputation, or other aspects of argumentation (Mochales and Moens 2011). Most argumentation-mining research has addressed non-legal genres, but interest in mining legal arguments is growing rapidly (see, e.g., Savelka and Ashley 2016; Walker et al. 2014; Nay 2016; Wyner et al. 2010).

3.3.3 Legal landscapes

A *legal landscape* is a global characterization of the state of the law relevant to a given set of tasks. Understanding the legal landscape can be important for legal planning and decision making. For example, a *patent landscape* is the collection of existing patents relevant to (e.g., representing prior art similar to) a given topic. *Automated Patent Landscaping* (Abood and Feltenberger 2016) addresses the task of identifying all patents related to a given topic using a semi-supervised machine learning model to prune an initial high-recall set of candidates down to a high-precision subset.

A related problem is detection concept drift in legal concepts. This problem was addressed by applying machine learning to cases manually represented as sets of features in Rissland and Friedman (1995), but current language processing technology is making unsupervised approaches to legal landscape analysis increasingly feasible. For example, vendors of text-analytic-based litigation services mentioned above, such as Lex Machina,¹³ market the ability to predict the likely outcome of an action in a given litigation context based on a global analysis of the prior outcomes of similar actions. While the academic research literature seems to contain few examples of unsupervised legal landscape analysis, work in this area may become increasingly common.

¹³ <https://lexmachina.com/> [Accessed: 27 November 2016].

3.3.4 *Judicial database analysis*

The transition from paper to electronic filing, which began in many jurisdictions in the late 1990s, has transformed how courts operate and how judges, court staff personnel, attorneys, and the public create, submit, and access court filings. However, despite many advances in judicial access and administration brought about by electronic filing, courts are typically unable to interpret the contents of court filings automatically. Instead, court filings are interpreted only when they are read by an attorney, judge, or court staff member.

A court docket is a register of document-triggered litigation events, where a litigation event consists of either (1) a pleading, motion, or letter from a litigant, (2) an order, judgment, or other action by a judge, or (3) a record of an administrative action (such as notifying an attorney of a filing error) by a member of the court staff. Docket events can be viewed as illocutionary speech acts: if well-formed and in the right context, they change the state of the litigation and, potentially, the legal relationships between the litigants.

Because documents are typically accompanied by metadata generated at the time of filing (e.g., parties, court, judge, intended docket event, etc.) docket databases can be viewed as large annotated corpora to which machine learning algorithms and text analytics can be applied. The knowledge latent in these datasets may be as useful for courts and the public as for attorneys.

Many of the predictive litigation services mentioned above make use of docket databases purchased from the U.S. federal judiciary. However, courts themselves are only in the early stages of exploiting this latent information to improve court administration.

One application is filing-error detection, a significant issue for courts in which electronic filing has replaced the docket clerks who formerly checked paper submissions for accuracy and completeness. Text classification and information extraction techniques have been applied to detect filing errors, such as filing a document to a wrong case or as a wrong type (e.g., filing a motion as an order or an exhibit as a separate filing event from the document to which it is attached) or omitting critical information (Branting 2016).

Another application is improving court management by predicting resource allocation factors, such as case duration, probability of settlement prior to trial, probable trial duration, etc., from case-specific information such as cause of action, amount in controversy, attorneys, parties, and, ultimately, the description of the facts and issues in controversy.

3.4 **The opacity of data-centric approaches**

Data-centric approaches promise the benefits of scalability and obviation of laborious human encoding of legal texts in logical formalisms. However, the benefits come at the cost of decreased transparency and explanatory capability. As mentioned above, the EU's "General Data Protection Regulation" is intended to guarantee an explanation for algorithmic assessment of rights (Goodman and Flaxman 2016a, b). In the United States there has been controversy over use of

recidivism risk assessment (Angwin et al. 2016) and sentencing models (Sidhu 2015) that are biased but whose opacity makes challenging their conclusions impossible. There is active research both in understanding and reducing bias in empirical models (Chouldechova 2016) and in making neural network models of text analysis more comprehensible, e.g., Li et al. (2016). However, application of data-centric models for adjudication or assessment of rights or penalties remains quite problematical. Data-centric approaches may be better used as tools for increasing the comprehensibility of legal texts and corpora and the predictability of legal systems than as black-box components of those systems. Data-centric approaches have the potential to increase access to justice and the predictability and efficiency of legal systems, but are capable of misuse.

4 Legal tasks amenable to each technique

The relative merits of logic-based and data-centric approaches depend on the particular legal problem-solving tasks to which they are applied. It is therefore important to distinguish among the variety of different legal tasks.

The legal systems of democratic nations comprise a multiplicity of legal policies, statutes, norms, and regulations. The complexity of the interactions among these authorities and between these authorities and the behavior of citizens defies concise summary. However, it is possible to enumerate some basic categories of tasks relevant to the main stakeholders in the legal system, who include citizens, judges and magistrates, attorneys, legislators and regulatory rule makers, government agencies, scholars and teachers, and juries. A tentative assessment is proposed below for each such category as to the degree to which one approach or the other is likely to be more appropriate.

1. **Analysis** Among the first tasks proposed for automation was *analysis*, determining and justifying the legal consequences of a given set of facts. Many more complex legal tasks depend on the fundamental ability to draw legal conclusions from facts based on a given set of rules. In the earliest systems, analysis consisted of deduction over logical representation of legal rules and facts. Later work, particularly in common-law domains, incorporated *case-based reasoning* into the analytic process as an approach to grounding abstract legal predicates in the specific case facts. In general, these case facts were themselves represented in some logical formalism rather than in free text. Since legal conclusions and justifications are ordinarily expressed in terms of legal rules, legal analysis seems the ideal application for logic-based reasoning. Data-centric approaches seem ill-suited except in the case of evaluating a single legal predicate, in which case an explanation in terms of relevant similarities and differences to exemplars (e.g., precedents) can sometimes be acceptable.
2. **Retrieval** Information retrieval (IR) in the legal context consists of finding legal texts useful for one of the other tasks in this enumeration.

IR is probably the earliest application of data-centric techniques to legal problem solving.

3. **Prediction** Predictability is a fundamental characteristic of well-functioning legal systems (Lindquist and Cross 2017). The input to a typical prediction task consists of case facts, claims, a forum, and other contextual information, such as attorneys, judges, etc. Typical desired outputs include the likelihood that a claim or motion will prevail, the expected judgment or award, or the case duration. Prediction is a very natural and increasingly popular data-centric application, in part because it can be framed as a classification or regression task, for which there are very extensive resources. Since this requires nothing more than a corpus of fact/outcome pairs, it is highly scalable but, depending on the inductive method, may have negligible explanatory capability. However, predictions may be acceptable simply on the basis of past accuracy of the predictive model. Logic-based approaches to prediction typically require comparing the persuasiveness of alternative justifications for a given legal conclusion. Human attorneys may provide both types of explanation: empirical, based on past experience; and logical, based on the relative persuasiveness of competing arguments.
4. **Argument generation** Legal discourse typically consists of arguments for or against a given legal conclusion, because unsupported legal assertions have little persuasive weight. Legal adjudication typically consists of the choice of the most persuasive of two conflicting arguments. Many case-based reasoning systems were designed to produce alternative arguments on a given legal issue. The logical character of legal arguments makes this a task for which logical approaches are currently much better suited. However, there is growing interest in discourse generation using neural networks (see, e.g., Serban et al. 2016), so eventually these techniques may be applicable to argumentation, but currently this seems to be solely the ambit of logic-based approaches.
5. **Dialectical argumentation** Some legal discourse, such as pleadings in court cases, have a dialectical character in which each party asserts arguments in response to prior assertions or arguments by the other party. This requires the ability not just to generate a well-structured argument, but to recognize what argument is appropriate at a given stage of discourse.¹⁴

Since dialectical argumentation seems to require argument generation as a subtask, one might think that it is even less amenable to data-centric approaches. However, there are some cases in which empirical analysis is informative. For example, sequences of court pleadings can be analyzed using sequence models, such as Conditional Random Fields (Lafferty et al. 2001), to learn the most likely next event or the event most likely to lead to a given outcome. So, both logic-based and data-centric approaches may be appropriate to dialectical argumentation, the former addressing the precise logical character of the individual argumentation steps, the latter on the regularities that can be induced from collections of dialectics.

¹⁴ See Prakken (2005) for a review of the history of argumentation schemes in automated legal reasoning.

6. **Document drafting** Much of the work product of attorneys typically consists of creating individual documents to achieve legal goals, e.g., pleadings, trust instruments, wills, contracts, orders, or decisions (Lauritsen and Gordon 2009).
Past AI models of document drafting have been logic-based. As mentioned above, discourse generation through neural networks is in its infancy. Recent work on inducing document plans from corpora for use in generating new documents is a promising development that may enable data-centric methods to contribute to the document drafting task (see, e.g., Konstas and Lapata 2013).
7. **Planning** An important activity of many attorneys is structuring transactions or enterprises to achieve particular legal goals, such as reducing taxes or conveying ownership.
In general, legal planning appears to be a highly rule-governed activity, requiring logical representations of both legal rules, states, and actions (Sanders 1994). Recent work has demonstrated how plans for tax evasion and for detecting tax evasion can be developed by genetic algorithms (Hemberg et al. 2015). However, the chromosomal representations in this project were manually engineered, so this approach is a kind of hybrid of symbolic and inductive techniques.
8. **Legislative drafting** Legislative drafting often involves changes to large and complex statutory or regulatory texts that comprise many interacting provisions. Crafting an amendment to a legislative text requires understanding how the changes will effect other provisions so that the goals of the changes can be achieved while maintaining consistency with the rest of the text.
As discussed above, network models of legislation, e.g., Koniaris et al. (2015), are an area of active research. The full potential of these models is still being explored, but network analysis promises improved analysis and consistency-checking for legislators.
9. **Trend analysis** Identifying past and predicting future changes in legal doctrine or outcomes.
Trend analysis is a paradigmatic application of data-centric techniques, since it involves temporal analysis over many examples rather than logical inference. Most current work in this area appears to be in the area of litigation support, discussed above.
10. **Adjudication** Resolving conflicting claims based on factual evidence and legal arguments constitutes the task of adjudication. There are human-rights and jurisprudential reasons that adjudication should remain a matter of human discretion. However, routine disputes can often be settled inexpensively through automated Alternative Dispute Resolution (ADR) techniques. If the decision in an adjudication can be anticipated, proposed work products such as proposed orders can be pre-drafted, increasing administrative efficiency.
One approach to decision support for adjudication makes use of logical models of the rules that the adjudicator will apply. For example, a decision support system can back-chain from a top-level *adjudicate* goal, querying the adjudicator at each decision point and either instantiating a document

template with values gathered during the inference process (Pethe et al. 1989) or constructing the document through unification as a side effect of the inference (Branting et al. 1998; Branting 1998). An alternative, data-centric, approach organizes and analyzes data relevant to a decision, highlighting the similarities and differences between the instance case and prior cases (Schild 1998; Tata 1998). More modern data-centric approaches have the potential to inform adjudicators about relevant aspects of a case's record and context, in much the same way that data-centric litigation support tools can greatly enhance the insight of litigators into a case.

11. **Legal document auditing and quality control** The widespread adoption of electronic case management systems has greatly improved the efficiency of document handling. However, by eliminating docket clerks who formerly checked filings for completeness and correctness, electronic filing introduces a small but steady stream of erroneous documents into court dockets. Manually auditing these documents is tedious for staff and expensive for court administrators. Automating the detection of document filing errors would significantly improve court administration.

Data-centric approaches are a promising approach to automated error detection. As discussed above, court dockets can be treated as annotated corpora. Models induced from prior audit actions can be used to detect new errors, significantly reducing courts' audit loads. Error detection at submission time can reduce the number of errors reaching the audit staff and can reduce the frustration and delays experienced by attorneys as a result of inadvertent filing errors.

5 Conclusion

This paper has contrasted the capabilities and limitations of logic-based and data-centric approaches to legal problem solving. The previous section distinguished representative tasks for which the logic-based approach is more suitable from those for which the data-centric approach is preferable. The larger theme of this discussion, however, is that these two approaches are complementary. Many forms of legal argumentation and discourse are structured around precise rules that are modeled well by logic. Other aspects of legal problem solving, such as grounding the semantics of legal terms in the language of ordinary discourse, identifying global characteristics of systems of laws, or predicting legal events, have no natural fit within the logical framework but are better suited to empirical analysis.

However, many tasks could benefit from both approaches. Figure 4 sets forth a notional model in which the AI and law model for legal analysis proposed by Meldman in the 1970s is updated in two ways. First, logical representations of rules are created by a "logic-aware" rule editor that gives rise to an authenticated logical representation of the rules intended by the drafters. Subgoals arising from inference over this representation, consisting of legal concepts or predicates, are evaluated on the basis of an inductive model consisting either of exemplar-based reasoning over

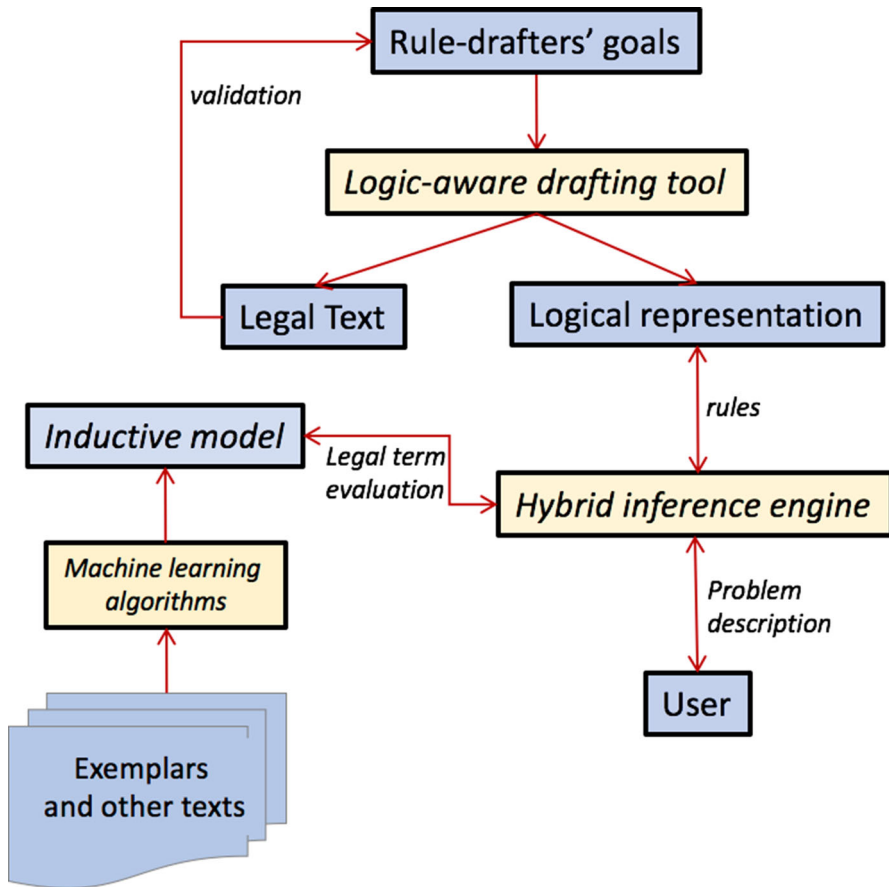


Fig. 4 Logic-based and data-centric approaches play complementary roles in a hybrid model of legal reasoning

legal precedents expressed in text or of models induced from such precedents. In this approach, data-centric reasoning would bridge the language gap between legal precedents and problem descriptions, while understandable and persuasive justifications were produced from the legal rules.

Use of data-centric methods to evaluate predicates arising in legal rules seems currently to be a much more achievable goal than a logic-aware rule-drafting tool carrying an official imprimatur. In the near term, hybrid systems may be developed that are similar to previous hybrid rule/CBR systems except that they will not require the manual case representation of, e.g., Ashley (1990), Rissland and Skalak (1991), and Branting (2000b). This means that they will be scalable with respect to new precedents even though conversion of rule texts into a logical representation would continue to be manual.

In the meantime, the financial rewards to developers of data-centric systems are such that development in this area is likely to continue apace. It will be the

challenge of the AI and law community to integrate these new technologies with older logic-based techniques to fully exploit the complementarity of the two technologies.

Acknowledgements The MITRE Corporation is a not-for-profit Federally Funded Research and Development Center chartered in the public interest. This document is approved for Public Release, Distribution Unlimited, Case Number 16-4565 ©2016 The MITRE Corporation. All rights reserved.

References

- Abood A, Feltenberger D (2016) Automated patent landscaping. In: Proceedings of the workshop on legal text, document, and corpus analytics (LTDC A-2016), pp 1–8. http://law-and-big-data.org/LTDC A_2016_Workshop_Report.pdf
- Aletras N, Tsarapatsanis D, Preotiu-Pietro D, Lampos V (2016) Predicting judicial decisions of the European Court of Human Rights: a natural language processing perspective. *PeerJ CompSci*. <https://peerj.com/articles/cs-93/>
- Aleven V, Ashley K (1996) Doing things with factors. In: Proceedings of the 3rd European workshop on case-based reasoning (EWCR-96). Lausanne, Switzerland, pp 76–90
- Allen L (1957) Symbolic logic: a razor-edged tool for drafting and interpreting legal documents. *Yale Law Journal* 66:833–879
- Allen L (1982) Towards a normalized language to clarify the structure of legal discourse. In: Edited versions of selected papers from the international conference on logic, informatics, law, vol 2, pp 349–407. North-Holland Publishing Company, Amsterdam
- Angwin J, Larson J, Mattu S, Kirchner L (2016) Machine bias: there’s software used across the country to predict future criminals. and it’s biased against blacks. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Ashley K (1990) Modelling legal argument: reasoning with cases and hypotheticals. MIT Press, Cambridge
- Bench-Capon T, Coenen F (1992) Isomorphism and legal knowledge based systems. *Artif Intell Law* 1(1):65–86
- Bench-Capon T, Gordon TF (2009) Isomorphism and argumentation. In: Proceedings of the 12th international conference on artificial intelligence and law, ICAIL ’09, ACM, New York, NY, USA, pp 11–20
- Bench-Capon TJM, Visser PRS (1997) Open texture and ontologies in legal information systems. In: Proceedings of the eighth international workshop on database and expert systems applications, 1997. pp 192–197
- Bommarito MJ, Katz D (2009) Properties of the United States code citation network. <https://ssrn.com/abstract=1502927>
- Branting K (1998) Techniques for automated drafting of judicial documents. *Int J Law Inf Technol* 6(2):214–229
- Branting LK (2000a) An advisory system for pro se protection order applicants. *Int Rev Law Comput Technol* 14(3): 357
- Branting LK (2000b) Reasoning with rules and precedents: a computational model of legal analysis. Kluwer, Dordrecht
- Branting LK (2016) Vocabulary reduction, text excision, and procedural-context features in judicial document analytics. In Proceedings of the workshop on legal text, document, and corpus analytics (LTDC A-2016), pp 30–36. http://law-and-big-data.org/LTDC A_2016_Workshop_Report.pdf
- Branting LK, Conrad J, editors (2016). Legal text, document, and corpus analytics (LTDC A 2016) workshop report. http://law-and-big-data.org/LTDC A_2016_Workshop_Report.pdf
- Branting LK, Lester J, Callaway C (1998) Automating judicial document drafting: a unification-based approach. *Artif Intell Law* 6(2–4):111–149
- Brudney J, Ditslear C (2005) Canons of construction and the elusive quest for neutral reasoning. *Vanderbilt Law Rev* 58:1
- Brügmann S, Bouayad-Agha N, Burga A, Carrascosa S, Ciaramella A, Ciaramella M, Codina-Filba J, Escorsa E, Judea A, Mille S, Miller A, Saggion H, Ziering P, Schtze H, Wanner L (2014) Towards

- content-oriented patent document processing: intelligent patent analysis and summarization. *World Patent Inf* 40:30–42
- Brüninghaus S, Ashley K (2001) The role of information extraction for textual CBR. In: Aha D, Watson I (eds) *Case-based reasoning research and development*. Springer, Berlin, pp 74–89
- Buabuchachart A, Metcalf K, Charness N, Morgenstern L (2013) Classification of regulatory paragraphs by discourse structure, reference structure, and regulation type. In: *Proceedings of the 26th international conference on legal knowledge-based systems JURIX*, University of Bologna, Bologna, Italy
- Chandler SJ (2005) The network structure of supreme court jurisprudence. Technical report tech report no. 2005-W-01, University of Houston Law Center
- Chouldechova A (2016) Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. In: *Proceedings of fairness, accountability, and transparency in machine learning*, New York University, New York, NY. [arXiv:1610.07524](https://arxiv.org/abs/1610.07524)
- Dozier C, Kondadadi R, Light M, Vachher A, Veeramachaneni S, Wudali R (2010) Named entity recognition and resolution in legal text. Springer, Berlin
- Goodman B, Flaxman S (2016) EU regulations on algorithmic decision-making and a “right to explanation”. In *ICML workshop on human interpretability in machine learning*, New York, NY
- Goodman B, Flaxman S (2016) European union regulations on algorithmic decision-making and a “right to explanation”. *ArXiv e-prints*
- Gray G (1987) Reducing unintended ambiguity in statutes: an introduction to normalization of statutory drafting. *Tenn Law Rev* 54:433
- Hemberg E, Rosen J, Warner G, Wijesinghe S, O’Reilly U.-M. (2015) Tax non-compliance detection using co-evolution of tax evasion risk and audit likelihood. In: Atkinson K, Sichelman T (eds) *Proceedings of the 15th international conference on artificial intelligence and law, ICAIL-2015*, San Diego, USA. ACM, pp 79–88
- Hermens LA, Schlimmer JC (1994) A machine-learning apprentice for the completion of repetitive forms. *IEEE Expert* 9:28–33
- Koniaris M, Anagnostopoulos I, Vassiliou Y (2015) Network analysis in the legal domain: a complex model for european union legal sources. *CoRR*, [arXiv:abs/1501.05237](https://arxiv.org/abs/1501.05237)
- Konstas I, Lapata M (2013) Inducing document plans for concept-to-text generation. In: *Conference on empirical methods in natural language processing (EMNLP 2013)*, Seattle, USA
- Lafferty J, McCallum A, Pereira F (2001) Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: *Proceedings of the 18th international conference on machine learning (ICML 2001)*, pages 282–?289. Morgan Kaufmann
- Lakoff G (1987) *Women, fire, and dangerous things*. University of Chicago Press, Chicago
- Lauritsen M, Gordon T (2009) Toward a general theory of document modeling. In: *Proceedings of the 12th international conference on artificial intelligence and law (ICAIL 2009)*. ACM, pp 202–211
- Li J, Chen X, Hovy EH, Jurafsky D (2016) Visualizing and understanding neural models in NLP. In: *NAACL HLT 2016*, the 2016 conference of the north American chapter of the association for computational linguistics: human language technologies, San Diego California, USA, June 12–17, 2016, pp 681–691
- Lindquist S, Cross F (2017) Stalility, predictability and the rule of law: stare decisis as reciprocity. <https://law.utexas.edu/conferences/measuring/The%20Papers/Rule%20of%20Law%20Conference.crosslindquist.pdf>
- Loftus EF, Wagenaar WA (1988) Lawyers’ predictions of success. *Jurimetr J* 29:437–453
- McCarty L (1977) Reflections on “taxman”: an experiment in artificial intelligence and legal reasoning. *Harv Law Rev* 90:837–893
- Mehl L (1958) Automation in the legal world: from the machine processing of legal information to the “law machine”. In: *Symposium on the mechanisation of the thought process*
- Meldman JA (1975) A preliminary study in computer-aided legal analysis. PhD thesis, Massachusetts Institute of Technology
- Mochales R, Moens M-F (2011) Argumentation mining. *Artif Intell Law* 19(1):1–22
- Morgenstern L (2014) Toward automated international law compliance monitoring (tailcm). Technical report, LEIDOS, INC. AFRL-RI-RS-TR-2014-206
- NAiL (2013). Workshop on network analysis in law. In: *The 14th international conference on ai and law, held in conjunction with ICAIL 2013*. <http://www.leibnizcenter.org/~winkels/NAiL2013.html>. Accessed 2 Dec 2016

- NAiL 2 (2014). Second international workshop on network analysis in law. In: The 27th international conference on legal knowledge and information systems. Held in conjunction with JURIX 2014, 10–12 Dec 2014, Krakow, Poland. <http://www.leibnizcenter.org/~winkels/NAiL2014.html>. Accessed 2 Dec 2016
- NAiL 3 (2015). Third international workshop on network analysis in law. In: The 28th international conference on legal knowledge and information systems, Held in conjunction with JURIX 2015, 9–11 December 2015, Braga, Portugal. <http://www.leibnizcenter.org/~winkels/NAiL2015.html>. Accessed 2 Dec 2016
- Nay JJ (2016) Gov2vec: Learning distributed representations of institutions and their legal text. In: Proceedings of the first workshop on NLP and computational social science. Association for Computational Linguistics, pp 49–54, Austin, Texas
- Newell A, Shaw J, Simon H (1959) Report on a general problem-solving program. In: Proceedings of the international conference on information processing, pp 256–264
- Oskamp A, Lauritsen M (2002) AI in law practice? So far, not much. *Artif Intell Law* 10:227–236
- Peterson M, Waterman D (1985) Rule-based models of legal expertise. In: Walters C (ed) *Computing power and legal reasoning*. West Publishing Company, Minneapolis, pp 627–659
- Pethe VP, Rippey CP, Kale LV (1989) A specialized expert system for judicial decision support. In: Proceedings of the 2nd international conference on artificial intelligence and law. Vancouver, BC, pp 190–194
- Prakken H (2005) AI & law, logic and argument schemes. *Argumentation* 19(3):303–320
- Rissland E, Friedman MT (1995) Detecting change in legal concepts. In: Proceedings of the 5th international conference on artificial intelligence and law (ICAIL 1995), College Park, MD, USA. ACM Press, pp 127–136
- Rissland E, Skalak D (1991) Cabaret: rule interpretation in a hybrid architecture. *Int J Man-Mach Stud* 34(6):839–887
- Rush AM, Chopra S, Weston J (2015) A neural attention model for abstractive sentence summarization. CoRR, [arXiv:abs/1509.00685](https://arxiv.org/abs/1509.00685)
- Sadeghian A, Sundaram L, Wang D, Hamilton W, Branting K, Pfeifer C (2016). Semantic edge labeling over legal citation graphs. In: Proceedings of the workshop on legal text, document, and corpus analytics (LTDCA-2016), pp 70–75. http://law-and-big-data.org/LTDCA_2016_Workshop_Report.pdf
- Sanders K (1994) CHIRON: Planning in an Open-Textured Domain. PhD thesis, Brown University
- Savelka J, Ashley K (2016) Extracting case law sentences for argumentation about the meaning of statutory terms. In: Proceedings of the 3rd workshop on argument mining (ArgMining2016), Berlin, Germany. Association for Computational Linguistics
- Schild UJ (1998) Criminal sentencing and intelligent decision support. *Artif Intell Law* 6(2):151–202
- Schwartz T, Berger M, Hernandez J (2015) A legal citation recommendation engine using topic modeling and semantic similarity. In: Law and big data workshop, 15th international conference on artificial intelligence and law
- Serban IV, Klinger T, Tesauro G, Talamadupula K, Zhou B, Bengio Y, Courville AC (2016) Multiresolution recurrent neural networks: an application to dialogue response generation. CoRR, [arXiv:abs/1606.00776](https://arxiv.org/abs/1606.00776)
- Sergot MJ, Sadri F, Kowalski RA, Kriwaczek F, Hammond P, Cory HT (1986) The British Nationality Act as a logic program. *Commun ACM* 29(5):370–386
- Sidhu D (2015) Moneyball sentencing. *Boston Coll Law Rev* 56(2):672–731
- Surdeanu M, Nallapati R, Gregory G, Walker J, Manning C (2011) Risk analysis for intellectual property litigation. In: Proceedings of the 13th international conference on artificial intelligence and law, Pittsburgh, PA. ACM
- Surdeanu M, Nallapati R, Manning C (2010) Legal claim identification: information extraction with hierarchically labeled data. In: Proceedings of the 7th international conference on language resources and evaluation. LREC
- Tata C (1998) ‘Neutrality’, ‘choice’, and ‘ownership’ in the construction, use, and adaptation of judicial decision support systems. *Int J Law Inf Technol* 6(2):143–166
- van Engers T, Nijssen S (2014) From legislation towards the provision of services - an approach to agile implementation of legislation. *Lecture notes in computer science*, pp 163–172
- Vilain M (2016) Language-processing methods for us court filings. In: Proceedings of the workshop on legal text, document, and corpus analytics (LTDCA-2016), pp 81–90. http://law-and-big-data.org/LTDCA_2016_Workshop_Report.pdf

- Walker V, Vazirova K, Sanford C (2014) Annotating patterns of reasoning about medical theories of causation in vaccine cases: toward a type system for arguments. In: Proceedings of the 1st workshop on argumentation mining, Baltimore, Maryland. Association for Computational Linguistics, pp 1–10
- Wyner A, Mochales Palau R, Moens M-F, Milward D (2010) Approaches to text mining arguments from legal cases. In: Francesconi E, Montemagni S, Peters W, Tiscornia D (eds) Semantic processing of legal texts. LNCS (LNAI), vol 6036, pp 60–79. Springer, Heidelberg
- Wyner A, Peters W (2011) On rule extraction from regulations. *Front Artif Intell Appl* 235:113–122