

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection Yong Pung How School Of
Law

Yong Pung How School of Law

1-2023

Natural language processing in the legal domain

Daniel Martin KATZ

Dirk HARTUNG

Singapore Management University, dirkhartung@smu.edu.sg

Lauritz GERLACH

Abhik JANA

Michael J. II BOMMARITO

Follow this and additional works at: https://ink.library.smu.edu.sg/sol_research



Part of the [Programming Languages and Compilers Commons](#), and the [Science and Technology Law Commons](#)

Citation

KATZ, Daniel Martin; HARTUNG, Dirk; GERLACH, Lauritz; JANA, Abhik; and BOMMARITO, Michael J. II. Natural language processing in the legal domain. (2023). 1-13.
Available at: https://ink.library.smu.edu.sg/sol_research/4527

This Working Paper is brought to you for free and open access by the Yong Pung How School of Law at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection Yong Pung How School Of Law by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Natural Language Processing in the Legal Domain

Daniel Martin Katz^{1,2,3,4,†,*}, Dirk Hartung^{2,3,†}, Lauritz Gerlach², Abhik Jana⁵, and Michael J. Bommarito^{2,3,4}

[†] These authors have contributed equally to this work and share first authorship.

¹ Illinois Tech - Chicago Kent College of Law, USA

² Bucerius Law School, Germany

³ CodeX, Stanford University, USA

⁴ 273 Ventures, USA

⁵ Universität Hamburg, Germany

* e-mail - dkatz3@kentlaw.iit.edu

ABSTRACT

In this paper, we summarize the current state of the field of NLP & Law with a specific focus on recent technical and substantive developments. To support our analysis, we construct and analyze a nearly complete corpus of more than six hundred NLP & Law related papers published over the past decade. Our analysis highlights several major trends. Namely, we document an increasing number of papers written, tasks undertaken, and languages covered over the course of the past decade. We observe an increase in the sophistication of the methods which researchers deployed in this applied context. Slowly but surely, Legal NLP is beginning to match not only the methodological sophistication of general NLP but also the professional standards of data availability and code reproducibility observed within the broader scientific community. We believe all of these trends bode well for the future of the field, but many questions in both the academic and commercial sphere still remain open.

Introduction

Language is the ‘coin of the realm’ in the legal domain. Not only do legal institutions and actors produce, consume, and interpret massive volumes of text,¹ but virtually every legal process involves either the production or consumption of documents. Careful drafting of documents and the analysis and interpretation of language are among the core activities undertaken by judges, regulators, legislators, and lawyers. Participants in the world’s legal systems “continuously author legal texts such as statutes, regulations, judicial decisions, contracts, patents, briefs, memos, and other related materials.”² Taken together, legal systems output large volumes of documents and these documents are often complex. Indeed, the “language of law” has proven to be so challenging that many laypersons describe legal documents and arguments using terms such as ‘legalese’, ‘legal jargon’ or ‘legal gobbledegook.’

The complexity of the law^{3,4,5} is not just a scientific phenomenon; it has real consequences for many individuals and organizations.^{6,7} In part due to this complexity, legal systems have struggled to assist with “the quantity, quality, and accessibility of legal services demanded by society.”⁸ Yet, despite this underlying and obvious need for improvement in the delivery of justice,^{9,10,11,12} there have been many barriers which have prevented the emergence of scalable solutions to meet various legal needs. These barriers include the culture of law (including lawyers, judges, and legal educators)^{13,14} as well as the regulation of the profession.^{15,16,17} Yet, the primary *technical* challenge limiting transformative technological solutions within the legal sphere is the complex nature of legal language itself.

Simply put, the task of training machines to “understand” legal language has proven to be non-trivial. Notwithstanding the challenge, there has understandably been great interest in exploring the possibility of machines as a force multiplier for helping process complex legal texts. Indeed, both scholars and commercial enterprises have explored the applicability of Natural Language Processing (NLP) technologies for use within the field of law (Legal NLP). In the academic realm, empirical legal studies increasingly rely on a variety of methods from computer science to help support analysis.^{18,19,20} In the commercial sphere, there also have been attempts to embed Legal NLP modules into a number of applications in legal practice,^{21,22,23} from research tools and litigation outcome prediction to drafting support and compliance risk assessment. Overall, despite some laudable attempts, the performance of many commercial applications has at times been undermined by the inability to consistently process legal language in a high-fidelity manner.

Meanwhile in the more general technical literature, the past decade has witnessed major gains in the quality and performance of language models. Building upon foundational advances in neural network research,^{24,25} the broader field of NLP

has been reshaped in this period. While early neural NLP papers were built upon word embeddings,^{26,27,28} the latest wave of LLMs is being built upon the transformer architecture.²⁹ Among other things, the transformer architecture allows for the clever manipulation of the attention mechanism so that training tasks can be scaled through more effective parallelization.³⁰ Notwithstanding some critiques,³¹ successive waves of increasingly large transformer-based large language models (LLMs) have delivered some truly remarkable results.^{32,33,34,35,36}

Despite sometimes being characterized as general models, it is still an open question as to how much uptake or utility such core developments in NLP might offer when directed at complex domain-specific problems. While general models have shown real progress on legal tasks in the zero shot context,^{1,8,37,38,39,40} there are still strong reasons^{41,42,43} to believe that some combination of domain-specific pre-training, prompt engineering, prompt composition or chaining, hyper-parameter optimization, and other model tuning efforts will yield improved results in many substantive use cases. In other words, general NLP models will likely not eclipse the performance of an otherwise equally-sized large language model that has been well trained on the legal domain. That said, general models with enormous scale may very well outperform lower scale domain models.

Building a Corpus of Legal NLP Papers

Fueled by the wide scale expansion of digitized legal texts, the prospect for deploying cutting-edge NLP techniques within the legal domain has become increasingly possible. To support those efforts, while also providing a roadmap to various interdisciplinary scholars, we thought it to be useful to summarize emergent trends in the field. Although there have been efforts to characterize some developments in the field, including analysis of Legal NLP alone^{18,44} as well as discussions situated in the broader context of legal informatics,^{45,46} we believe a more holistic treatment of the current state of Legal NLP is justified.

Taking the past decade as our window of analysis, we constructed a corpus of nearly all published Legal NLP papers. To build the corpus, we began by reviewing the proceedings of both general NLP conferences (e.g. ACL, NAACL, EMNLP, EACL, etc.) as well as several speciality Legal NLP gatherings (NLLP, MLLD, Jurix, ICAIL, etc.). This initial screen yielded a large number of results. Next, we performed iterative queries on both major search engines and publication databases. Finally, we undertook a form of ‘snowball sampling’ by manually traversing the citation graph of key publications to scour for additional NLP & Law papers not otherwise identified using the approach above. Absent extraordinary circumstances, we restricted our corpus to include only peer-reviewed scientific journals, technical conference proceedings, and otherwise technically-oriented pre-prints (e.g. SSRN, arXiv). In other words, we generally excluded essays, commentaries, blog posts, social critiques, or otherwise non-technical and/or non-scientific publications.

Law as an intellectual domain is somewhat amorphous and thus some qualitative judgement is required to adjudicate its boundary. Therefore, care was required to determine whether a given paper is or is not ‘Legal NLP.’ As an illustrative example, there are many papers that apply NLP to analyze documents related to financial instruments. While the difference between these financial documents and legal documents often approaches zero, the focus and audience of these papers is typically much different. Therefore, we excluded these sort of efforts due to their limited nexus and connection to the legal domain. Our guiding principle was to include any technically-inclined paper whose target audience or target user was an individual working in the field of law.

Notwithstanding the aforementioned caveats, we believe we have acquired nearly every paper meeting these criteria published over the past decade. Figure 1 provides a time series representation of the volume of papers published over the past decade. We observe a significant increase in the total yearly volume of publications with only a slight diminution in 2020 (which is perhaps attributable to disruptive impacts of the early period of the COVID-19 pandemic).

The Engineering Perspective on Legal NLP Tasks

Law regulates all sectors of society. As such, the topic space covered therein is quite vast. Relevant topics include constitutional law, environmental law, intellectual property, labor law, corporate law, immigration law, criminal law, tax law, family law, and maritime law. Within each of these substantive areas, there are many specific tasks which lawyers must undertake. As reflected in various legal task schemas (e.g. UTBMS, SALI), legal workstreams can be decomposed into their constituent subtasks.^{47,48} For example, within a given legal matter, a legal professional might review documents, draft documents, conduct research, give arguments in court, negotiate with counter parties, etc. This diversity in tasks is reflected across our collection of more than six hundred papers, where scholars generally seek to apply Legal NLP techniques to augment, assist, or replace the lawyers as the sole available resource to complete a given task.




Figure 1. Number of Legal NLP Papers over Time

	Task	Examples / Description
1	Machine Summarization	Abstractive/Extractive Summaries of Legal Documents
2	Pre-Processing	Annotation, Anonymization, Translation
3	Classification	Outcome Prediction, Legal Area Classification, Topic Modeling
4	Information Retrieval	Legal Question Answering, Document Similarity, Document Retrieval
5	Information Extraction	Labeling, Text Extraction, Event Extraction
6	Text Generation	Automated Drafting of Legal Documents
7	Resources	Taxonomies, Ontologies, Datasets, Code Libraries

Table 1. A Taxonomy of Engineering Tasks in Legal NLP

In addition to the more lawyer-centric perspective of legal work, the space of NLP & Law tasks can be considered through the lens of engineering. Relevant engineering tasks could include machine summarization, translation, classification, retrieval, etc. Of course, these perspectives will often intersect. Consider a paper which explores automated extraction of keywords from a patent claim in order to support the search for prior art (i.e. other relevant patents).⁴⁹ This paper can be characterized as a patent paper from a substantive law perspective and as an information retrieval type paper from an engineering perspective. Recognizing this dualism attaches to the vast majority of papers in our corpus, we will—for the purposes of our analysis—privilege the engineering perspective on Legal NLP tasks. While certainly not the only way one could subdivide the space, Table 1 offers an engineering-centric taxonomy of Legal NLP tasks.

Leveraging our taxonomy, we qualitatively reviewed each of the papers in our corpus and determined the relevant engineering category for the work contained therein. While most papers fit squarely within one particular category, there were some papers which combined two or more engineering tasks. Figure 2 reflects the temporal distribution of papers by engineering task.

Subdividing the activity by categories yields interesting results: The general trend—upwards and to the right—is carried through most but not all categories. However, both Summarization and Classification show a change of pace around 2018 or 2019. Information Retrieval, on the other hand, shows fairly stable interest over time. Both Resources and Pre-Processing papers increase over time with the general stable increase of papers.

Some applications, such as Machine Summarization and Text Generation, have seen little activity compared to the more popular fields of Classification and Information Extraction. Given the current interest in text generation through popular exposure to applications such as ChatGPT, we expect a marked increase in the volume of Text Generation papers in the decade to come.




Figure 2. Legal NLP Tasks over Time

The Evolution of Methods in Legal NLP

The ability to work with and process language has long been an interest for scientific researchers. Indeed, arguably the most famous benchmark in the history of Artificial Intelligence, the Turing Test, involves a conversational interaction between a human and a computational agent.⁵⁰ The quest to fulfill the promises and goals of the field has taken scholars in many different directions. As such, NLP has experienced several waves of methodological innovation.

Early work in NLP can be traced to various rules-based systems which were either proposed or implemented.^{51,52,53} The advent of Moore’s Law⁵⁴ and the vast decline in the cost of data storage,⁵⁵ taken together with increasingly clever algorithmic methods, saw the field of NLP transformed from its rules-based AI origins into a data-driven field. This statistical turn, which began in the 1990’s, has more recently given way to the ‘neural era’ within the general field of NLP.⁵⁶

Legal AI and Legal NLP is a long-standing field with papers tracing back many decades.⁵⁷ Over time the field has seen significant methodological innovation. Consider a recent review of the Legal AI papers (authored by several leading scholars in the field) which discusses the increasing use of sophisticated methods. “[L]arge convolutional networks trained on graphics processors achieved breakthrough performance in computer vision [whereby] neural models quickly became the dominant technology . . . This development eventually reached the AI and Law community and led to a surge in the use of such models for the analysis of legal text from around 2017 onwards.”⁵⁷ Although our own qualitative review of our corpus supports this perspective, we thought it might be useful to empirically analyze this proposition.

Working with the full text of the more than six hundred papers in our Legal NLP corpus, we selected a set of key words which we believed could be tracked in order to understand the nature of methods contained within each paper. If the “neural agenda” were indeed ascendant, we would expect to observe an increase in relevant phrases over the most recent years. Having collected a full text copy of nearly every paper within our corpus,¹ we pre-processed each manuscript and extracted the plaintext from each document. While many more sophisticated approaches could be selected, we sought to simply track the emergence of key phrases which were indicative of ‘neural turn’ in Legal NLP.

Figure 3 offers the distribution of key phrases by paper count. While phrases such as ‘neural’ are quite common across the corpus, we recognize this alone could be unrelated to the proposition we seek to evaluate. For example, this could be the byproduct of references to much earlier periods when alternative forms of neural networks were popular.⁵⁸ Therefore, we

¹As of the release of this pre-print, we have collected full text copies of over 99.5% of the papers.




Figure 3. Distribution of Phrases by Paper

included many other key phrases which are associated with the modern neural NLP methods (*e.g.* embedding, word2vec, LSTM, BERT, etc.). Figure 4 provides a normalized time series of these phrases on a yearly basis. While many of these phrases unsurprisingly peak in 2022, we do observe that some first-wave embedding-based approaches (*e.g.* word2vec, doc2vec, etc.) appear to have peaked during the 2018-2021 window.

The Diversity of Languages

Deciding on the language of the data is a foundational decision at the core of each paper. The English language is historically well-positioned in this regard, as it is the *de facto* lingua franca of both computing and international (business) law. For this reason, we expected English to dominate the distribution of languages in our papers, with other world languages such as Chinese, German and French to be the nearest contenders.

This hypothesis was mostly proven correct in our evaluation: The most popular language is—by far—English (56%). The next most important language is Chinese. Other commonly occurring languages include Japanese, French, and German. In the case of Japanese and German, this may be partially explained by data availability: For Japanese, there are widely-used corpora of the Japanese National Pension Act and other statutes that are used in a number of papers surveyed, which might have an influence on the NLP community there. For the German language, there is ample data available, both at the national level² as well as at the supranational level through the European Union.

Figure 5 offers a stack plot of the time series of papers by language. Observing the distribution of the five most common languages over time per year, the proportion of English-language papers remains roughly constant, while papers analyzing Chinese-language corpora increase substantially. Additionally, Figure 5 reveals that the overall diversity of languages has increased, with the sum of less common languages increasing substantially over time (cf. the “other” band in Figure 5). As a word of limitation, it is important to note that we only surveyed papers that were written in the English language. Therefore, given that there is likely significant scholarship published only in languages such as Chinese, German, French, Spanish, etc. the results offered in Figure 5 should be considered as lower bound estimates for linguistic diversity.

²*e.g.*, through offerings such as “Gesetze im Internet” or “Rechtsprechung im Internet” and similar offerings in Austria and the German-speaking Swiss cantons.




Figure 4. Relative Rate of Term Usage over Time. Normalization is per-term relative to the maximum annual rate of mentioning papers.

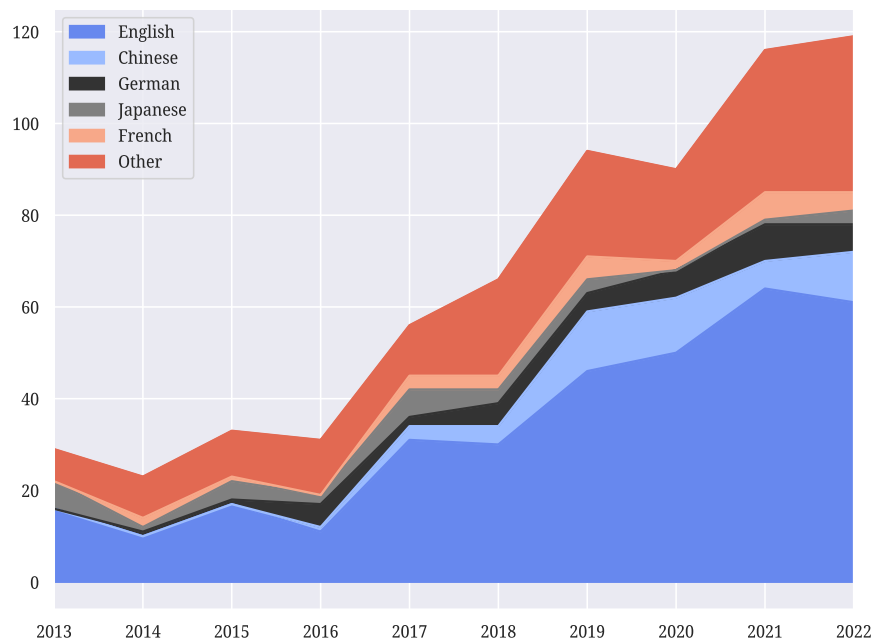


Figure 5. Temporal Distribution of the Most Popular Legal NLP Languages as a Function of Time

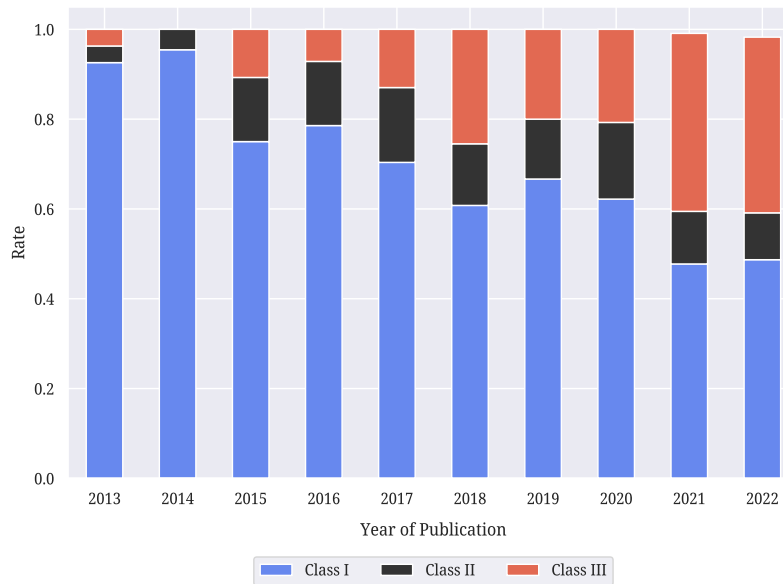


Figure 6. Replication Material Availability as a Function of Time

Reproducibility and Data Availability

Reproducibility is an increasing concern of not only the NLP community but also the broader scientific community.^{59,60} Beyond the important task of verifying existing results, transparent and replicable outputs can help accelerate the pace of additive innovation.

For each of the more than six hundred papers, we comprehensively evaluated the nature of replication resources made available by the respective authors. Of course, in judging the reproducibility of a given paper, many different aspects can be considered including the availability of datasets, corpora, models, proper documentation, etc. For ease, we categorize each the papers in our corpus into one of three classes. *Class I* contains those papers for which either no proper links to the resources are mentioned or the given link does not work, or no links are mentioned at all in the paper. For this class of papers, it would be quite difficult to reproduce the result only based upon the information provided by the authors. *Class II* papers provide partial resources to support replication / implementation. The nature of the incomplete element varies. Some papers provide only data sources, whereas, for some papers, only source codes are available, often with little or no documentation. Therefore, the frameworks or the results presented in this class of papers can not be reproduced straightforwardly, without material assistance from the authors. By contrast, *Class III* papers offer a well-organized repository of resources (both code and data) with proper documentation.

The three classes of papers from a reproducibility standpoint are summarized as follows:

- (1) **Class I** - No Data or Code is Available
- (2) **Class II** - Partial or Incomplete Data and/or Code Availability
- (3) **Class III** - Data and Code are Available in an Organized Repository

We analyze the availability of resources from two different perspectives. To start in Figure 6, we investigate whether with time, Legal NLP authors have become more inclined to make resources publicly available. The statistics for the last decade (2013-2022) are quite promising. We observe that the percentage of *Class I* papers has decreased significantly from 92.59% (2013) to 48.70% (2022), whereas the percentage of *Class III* papers increased noticeably from 3.70% (2013) to 39.13% (2022). These statistics highlight an increased commitment to reproducibility within the Legal NLP community.

We next consider to what extent the language of legal data correlates with the underlying availability of resources. The observations for some of the most researched languages (English, Chinese, German, Japanese & French) are presented in Figure 7. Even though the number of research papers analyzing English legal data is the highest by a considerable margin,

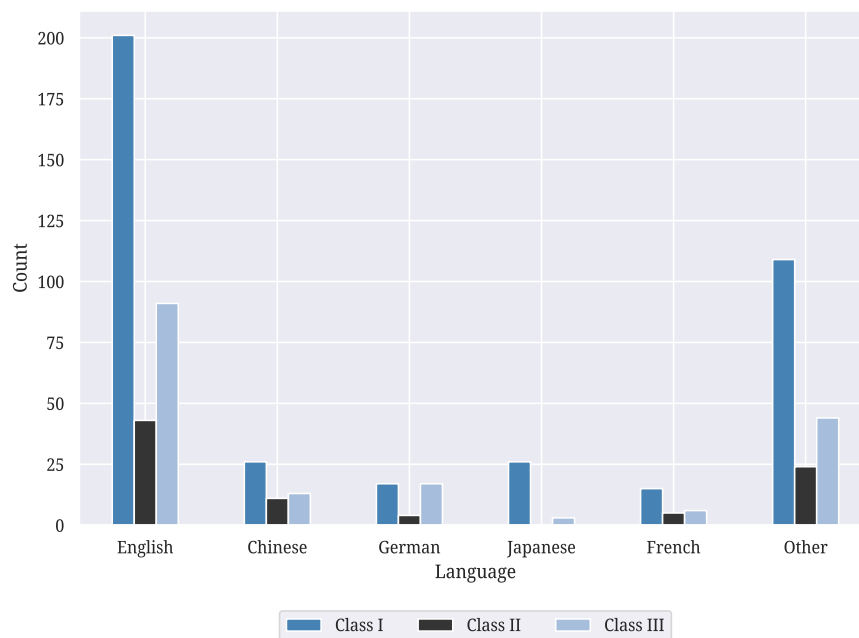


Figure 7. Replication Material Availability by Language

the percentage of *Class I* papers for English is also quite high (59.59%), and the percentage of *Class III* papers is quite low (26.84%). Chinese, the second most researched language, has 52% *Class I* papers and only 26% *Class III* papers. On the other hand, for not-so-resource-rich languages like Deutsch, the percentages of *Class III* papers for Deutsch are quite high (44.74%) compared to English and Chinese.

The detailed statistics are shared in our GitHub repository. Overall, researchers are getting more interested in making their resources (code, data, documentation, etc.) publicly available over time, irrespective of the underlying language in question.

Legal NLP Citation Analysis

As highlighted in herein, the field of Legal NLP is growing in many interesting ways. To help round out our analysis and understand which papers are commonly referenced by scholars, we collected basic citation data. Citations can provide a useful perspective to help identify key papers within the field. While citations are a noisy measure of ‘quality,’ they do reflect a crowd-sourced measure of prominence within the scientific community.

Leveraging Google Scholar as a source for citation information, we automated the collection of citation data while qualitatively reviewing the outputs of this automation for quality control purposes. Citations counts reflect the current status as of early January 2023. Across all of the papers, citation counts widely varied but something akin the classic bibliometric skew⁶¹ towards the top papers is present within the field of Legal NLP.

In reviewing the raw citations patterns, it was clear that the age of a publication was a contributing factor to citation count. This is, of course, not surprising as the longer a paper exists, the more opportunities a paper has to be cited. To help correct for this issue, we normalized citation counts. Although more sophisticated normalization methods are certainly possible,⁶² we conducted linear normalization of the citations counts for each paper in our corpus. All citation data will be made available in our online repository.

Table 2 offers a normalized list of the most cited papers in field authored in the past decade. Even a cursory review reflects the wide variety of substantive topics and jurisdictions. Collectively, these papers represent a good initial starting point for those interested in learning more about the emerging field of Legal NLP.

Rank	Title	Authors	Year
1	Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective	Aletras, Tsarapatsanis, Preoțiu-Pietro, Lampos	2016
2	LEGAL-BERT: The Muppets straight out of Law School	Chalkidis, Fergadiotis Malakasiotis, Aletras, Androutsopoulos	2020
3	Using machine learning to predict decisions of the European Court of Human Rights	Medvedeva, Vols, Wieling	2020
4	How Does NLP Benefit Legal System: A Summary of Legal Artificial Intelligence	Zhong, Xiao Tu, Zhang, Liu, Sun	2020
5	Legal Judgment Prediction via Topological Learning	Zhong, Haoxi, Guo, Tu, Xiao, Liu, Su	2019
6	LexGLUE: A Benchmark Dataset for Legal Language Understanding in English	Chalkidis, Jana, Hartung, Bommarito, Androutsopoulos, Katz, Aletras	2022
7	Neural Legal Judgment Prediction in English	Chalkidis, Androutsopoulos, Aletras	2019
8	Learning to Predict Charges for Criminal Cases with Legal Basis	Luo, Feng, Xu, Zhang, Zhao	2017
9	Few-Shot Charge Prediction with Discriminative Legal Attributes	Hu Li Tu, Liu, Sun	2018
10	Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measure	Arts, Hou, Gomez	2021
11	TechNet: Technology semantic network based on patent data	Sarica, Luo, Wood	2020
12	Text summarization from legal documents: a survey	Kanapala, Pal, Pamula	2019
13	Patent claim generation by fine-tuning OpenAI GPT-2	Lee, Hsiang	2020
14	CLAUDETTE: an automated detector of potentially unfair clauses in online terms of service	Lippi, Palka, Contissa, Lagioia, Micklitz, Sartor, Torroni	2019
15	A comparative study of automated legal text classification using random forests and deep learning	Chen, Wu, Chen, Lu, Ding	2022
16	Exploring the Use of Text Classification in the Legal Domain	Sulea, Zampieri, Malmasi, Vela, Dinu, van Genabith	2017
17	When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset of 53,000+ Legal Holdings	Zheng, Guha, Anderson, Henderson, Ho	2021
18	Deep learning in law: early adaptation and legal word embeddings trained on large corpora	Chalkidis, Kampas	2018
19	Legal Judgment Prediction via Multi-Perspective Bi-Feedback Network	Yang, Jia, Zhou, Luo	2019
20	Distinguish Confusing Law Articles for Legal Judgment Prediction	Xu, Wang, Chen, Pan, Wang, Zhao	2020
21	CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review	Hendrycks, Burns, Chen, Ball	2021
22	BERT-PLI: Modeling Paragraph-Level Interactions for Legal Case Retrieval	Shao, Mao, Liu, Ma, Satoh, Zhang, Ma	2020
23	Lawformer: A pre-trained language model for chinese legal long documents	Xiao, Hu, Liu, Tu, Sun	2021
24	EC-QA: a legal-domain question answering dataset	Zhong, Xiao, Tu, Zhang, Liu, Sun	2020
25	Extracting Contract Elements	Chalkidis, Androutsopoulos, Michos	2017
26	Automated requirements identification from construction contract documents using natural language processing	Hassan, Le	2020
27	Fine-Grained Named Entity Recognition in Legal Documents	Leitner, Rehm, Moreno-Schneider	2019
28	Iteratively Questioning and Answering for Interpretable Legal Judgment Prediction	Zhong, Wang, Tu, Zhang, Liu, Sun	2020
29	Legal Area Classification: A Comparative Study of Text Classifiers on Singapore Supreme Court Judgments	Soh, Lim, Chai	2019
30	VICTOR: a dataset for Brazilian legal documents classification	Araujo, de Campos, Braz, DaSilva	2020

Table 2. Most Cited Papers in Legal NLP Published Since 2013 (Normalized)

An Interactive Living Survey

A survey spanning a decade of global research is an enormous task, even in an emerging field such as Legal NLP. The sheer amount of information creates a significant risk of overburdening the community, resulting in scientific findings being overlooked or getting lost. Even at the beginning of the period under review, the growth rate of publications in our field (see *infra* 1) made it difficult to identify all relevant publications. As the number of new publications over time grew and methodical and linguistic diversity increased, the task has become constantly more challenging. Similar developments in science in general⁶³ have inspired a series of innovative attempts to solve the problem of scientific knowledge discovery and education³⁶ by using of large language models to support both research⁶⁴ and writing of scientific publications⁶⁵ albeit with limited success and apparent upper bounds despite large numbers of parameters at present.⁶⁶ Given current technical limitations and employing an established design mechanism in machine learning,^{67,68} we suggest a hybrid, human-in-the-loop approach to information management, combining state-of-the-art natural language processing tools to find and curate publications for review with subject matter expert human control. This idea of a *living* survey has been successfully implemented in other domains of machine learning research such as deep neural networks⁶⁹ and natural language processing research such as explainable artificial intelligence.⁷⁰

While we are confident that our extensive review of the existing literature has yielded a comprehensive collection of papers, the fast publication pace in the field as well as the linguistic diversity of legal research makes it probable that we may have missed relevant, individual contributions. For this survey to be as helpful for the community as possible we have therefore built web infrastructure³ to accept additional contributions from the public and continuously update our results accordingly. Researchers can provide a link to or upload their contribution, fill in a small number of fields with meta-data so that the task of maintaining the collection can be distributed across the community building on successful models of operation from the open source movement.

This digital infrastructure provides us with an opportunity to not only create a living but also an *interactive* survey. Since publications in the field of computational legal studies and natural legal language processing should be more than mere papers and ideally contain both code and data for reproduction,⁷¹ we have provided an easy-to-navigate graphical interface to explore the collections of papers contained in this review. Users can search publications and filter the collection according to their needs and based on the taxonomy laid out in Table 1. In addition, they can select individual and sets of publications to display temporal dynamics in methods, topics and languages.

Conclusion and Future Perspectives

As highlighted above, the field of Legal NLP is growing in volume, diversity of languages and sophistication of methods. As a result, research is conducted on an increasing variety of tasks as scientists push the limits of technical feasibility and demonstrate successful resolutions of increasingly difficult real-world challenges.⁸ Over our period of observation, an ever-growing amount of legal data and computational resources have become publicly available, certainly contributing to major advances in the field. However, the effects of training data, model architectures, and modeling techniques compared to the continuous increase in scale of general models requires extensive further research.

Large language models have recently fully entered the public's perception, resulting in viable commercial interest from players historically disinterested in Legal NLP research, such as publishing houses, law firms and courts. They often possess vast collections of data and might be increasingly willing to share them with researchers. This, in turn, is likely to lead to an uptick in research using real-world, commercial, or administrative datasets. While commercial players will drive the interest in Legal NLP research their very own challenge will lie in the connection of the resulting, ever increasing technical capabilities into products. Among the broader world of legal technology, it is therefore language-centric technologies that are likely to play an ever-increasing role. Public players will likely focus on deployment in the context of digital justice, leveraging models of the neural era to reduce case backlogs, improve access, and develop the ability to deliver justice at scale.

Our analysis has revealed that legal text generation has played a rather limited role over the last decade. With generative models capturing much of the attention recently, Legal NLP research will likely reflect this interest. As we demonstrated the usefulness of NLP methods for literature surveys themselves, one potential avenue for further research is the role of generative models for science itself and the division of labor between human researchers and large language models in our very own domain.

Concerning the engineering part of our survey, future work consists of two tasks: updating the *living* survey with state-of-the-art publications and building and releasing new *interactive* features for exploration. While we hope that the community

³Available at URL: %%% (Coming Soon)

can assist in the former we plan to focus on the latter, starting with a functionality to extract and visualize the reference graph of our collection and corresponding network metrics.

All of the above developments finally point to a growing awareness and better understanding of computational legal studies among traditional legal as well as other empirically-minded scholars. This might well result in an increased intra-disciplinary collaboration and a greater openness toward quantitative methods, thereby fostering the relevance of Legal NLP in the academic community.

Data Availability

All data collected in this survey will be made available on GitHub upon paper publication.

Acknowledgements

This work was partially funded by the German Federal Ministry of Education and Research (BMBF) through the [HILANO](#) project (funding code 01IS18085).

References

1. Chalkidis, I. *et al.* Lexglue: A benchmark dataset for legal language understanding in english. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 4310–4330 (2022).
2. Coupette, C., Beckedorf, J., Hartung, D., Bommarito, M. & Katz, D. M. Measuring law over time: A network analytical framework with an application to statutes and regulations in the united states and germany. *Front. Phys.* **9**, 658463 (2021).
3. Ruhl, J., Katz, D. & Bommarito, M. Harnessing legal complexity. *Science* **355**, 1377–1378 (2017).
4. Bommarito II, M. & Katz, D. Measuring and modeling the us regulatory ecosystem. *J. Stat. Phys.* **168**, 1125–1135 (2017).
5. Katz, D. M., Coupette, C., Beckedorf, J. & Hartung, D. Complex societies and the growth of the law. *Sci. reports* **10**, 1–14 (2020).
6. Ruhl, J. B. & Katz, D. M. Measuring, monitoring, and managing legal complexity. *Iowa L. Rev.* **101**, 191 (2015).
7. Staudt, R. W. All the wild possibilities: Technology that attacks barriers to access to justice. *Loy. LAL Rev.* **42**, 1117 (2008).
8. Bommarito II, M. & Katz, D. M. Gpt takes the bar exam. *arXiv preprint arXiv:2212.14402* (2022).
9. Rhode, D. L. *Access to justice* (Oxford University Press, 2004).
10. Susskind, R. E. *Online courts and the future of justice* (Oxford University Press, 2019).
11. Sandefur, R. L. & Teufel, J. Assessing america’s access to civil justice crisis. *UC Irvine L. Rev.* **11**, 753 (2020).
12. Prescott, J. J. Improving access to justice in state courts with platform technology. *Vand. L. Rev.* **70**, 1993 (2017).
13. Susskind, R. E. *Tomorrow’s lawyers: An introduction to your future* (Oxford University Press, 2017).
14. Barton, B. H. & Bibas, S. *Rebooting justice: More technology, fewer lawyers, and the future of law* (Encounter Books, 2017).
15. Kobayashi, B. H. & Ribstein, L. E. Law’s information revolution. *Ariz. L. Rev.* **53**, 1169 (2011).
16. Hadfield, G. K. The cost of law: Promoting access to justice through the (un) corporate practice of law. *Int. Rev. Law Econ.* **38**, 43–63 (2014).
17. Barton, B. H. & Rhode, D. L. Access to justice and routine legal services: New technologies meet bar regulators. *Hast. LJ* **70**, 955 (2018).
18. Fagan, F. Natural language processing for lawyers and judges. *Mich. L. Rev.* **119**, 1399 (2020).
19. Livermore, M. A. & Rockmore, D. N. *Law as Data: Computation, Text, & the Future of Legal Analysis* (Santa Fe Institute Press, 2019).
20. Kolt, N. Predicting consumer contracts. *Berkeley Technol. Law J.* **37** (2022).
21. Bommarito II, M. J., Katz, D. M. & Detterman, E. M. Lexnlp: Natural language processing and information extraction for legal and regulatory texts. In *Research Handbook on Big Data Law*, 216–227 (Edward Elgar Publishing, 2021).
22. Dale, R. Law and word order: Nlp in legal tech. *Nat. Lang. Eng.* **25**, 211–217 (2019).

23. Engstrom, D. F. & Gelbach, J. B. Legal tech, civil procedure, and the future of adversarialism. *U. Pa. L. Rev.* **169**, 1001 (2020).
24. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
25. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
26. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S. & Dean, J. Distributed representations of words and phrases and their compositionality. *Adv. neural information processing systems* **26** (2013).
27. Pennington, J., Socher, R. & Manning, C. D. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543 (2014).
28. Peters, M. E. *et al.* Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2227–2237 (2018).
29. Vaswani, A. *et al.* Attention is all you need. *Adv. neural information processing systems* **30** (2017).
30. Tay, Y., Dehghani, M., Bahri, D. & Metzler, D. Efficient transformers: A survey. *ACM Comput. Surv.* **55**, 1–28 (2022).
31. Bender, E. M., Gebru, T., McMillan-Major, A. & Mitchell, M. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623 (2021).
32. Kenton, J. D. M.-W. C. & Toutanova, L. K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171–4186 (2019).
33. Brown, T. *et al.* Language models are few-shot learners. *Adv. neural information processing systems* **33**, 1877–1901 (2020).
34. Zaheer, M. *et al.* Big bird: Transformers for longer sequences. *Adv. Neural Inf. Process. Syst.* **33**, 17283–17297 (2020).
35. Scao, T. L. *et al.* Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100* (2022).
36. Thoppilan, R. *et al.* Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239* (2022).
37. Zheng, L., Guha, N., Anderson, B. R., Henderson, P. & Ho, D. E. When does pretraining help? assessing self-supervised learning for law and the casehold dataset of 53,000+ legal holdings. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law*, 159–168 (2021).
38. Chalkidis, I., Fergadiotis, M. & Androutopoulos, I. Multieurlex-a multi-lingual and multi-label legal document classification dataset for zero-shot cross-lingual transfer. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 6974–6996 (2021).
39. Nay, J. J. Large language models as corporate lobbyists. *arXiv preprint arXiv:2301.01181* (2023).
40. Bommarito, J., Bommarito, M., Katz, D. M. & Katz, J. Gpt as knowledge worker: A zero-shot evaluation of (ai) cpa capabilities. *arXiv preprint arXiv:2301.04408* (2023).
41. Huang, J. *et al.* Large language models can self-improve. *arXiv preprint arXiv:2210.11610* (2022).
42. Wu, T. *et al.* Promptchainer: Chaining large language model prompts through visual programming. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–10 (2022).
43. Zhou, Y. *et al.* Large language models are human-level prompt engineers. *arXiv preprint arXiv:2211.01910* (2022).
44. Zhong, H. *et al.* How does nlp benefit legal system: A summary of legal artificial intelligence. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5218–5230 (2020).
45. Katz, D. M., Dolin, R. & Bommarito, M. J. *Legal informatics* (Cambridge University Press, 2021).
46. Ashley, K. D. *Artificial intelligence and legal analytics: new tools for law practice in the digital age* (Cambridge University Press, 2017).
47. Bartolo, M., Tylinski, K. & Moore, A. Pre-trained contextual embeddings for litigation code classification. In *LegalAIIA@ ICAIL*, 38–45 (2019).
48. Constantinou, V. & Kabiri, M. Detecting anomalous invoice line items in the legal case lifecycle. *arXiv preprint arXiv:2012.14511* (2020).

49. Rossi, J. & Kanoulas, E. Query generation for patent retrieval with keyword extraction based on syntactic features. In *JURIX*, 210–214 (2018).
50. Turing, A. M. Computing machinery and intelligence. *Mind* **59**, 433–460 (1950).
51. Chomsky, N. *Syntactic structures* (De Gruyter Mouton, 1957).
52. Schank, R. C., Goldman, N. M., Rieger III, C. J. & Riesbeck, C. Margie: Memory analysis response generation, and inference on english. In *IJCAI*, 255–261 (1973).
53. Lehnert, W. G. A conceptual theory of question answering. In *Proceedings of the 5th international joint conference on Artificial intelligence-Volume 1*, 158–164 (1977).
54. Shalf, J. The future of computing beyond moore’s law. *Philos. Transactions Royal Soc. A* **378**, 20190061 (2020).
55. Gupta, P. *et al.* An economic perspective of disk vs. flash media in archival storage. In *2014 IEEE 22nd International Symposium on Modelling, Analysis & Simulation of Computer and Telecommunication Systems*, 249–254 (IEEE, 2014).
56. Zhou, M., Duan, N., Liu, S. & Shum, H.-Y. Progress in neural nlp: modeling, learning, and reasoning. *Engineering* **6**, 275–290 (2020).
57. Governatori, G. *et al.* Thirty years of artificial intelligence and law: the first decade. *Artif. Intell. Law* **30**, 481–519 (2022).
58. Anderson, J. A. & Rosenfeld, E. *Talking nets: An oral history of neural networks* (MIT Press, 2000).
59. Munafò, M. R. *et al.* A manifesto for reproducible science. *Nat. human behaviour* **1**, 1–9 (2017).
60. Ivie, P. & Thain, D. Reproducibility in scientific computing. *ACM Comput. Surv. (CSUR)* **51**, 1–36 (2018).
61. Brzezinski, M. Power laws in citation distributions: evidence from scopus. *Scientometrics* **103**, 213–228 (2015).
62. Owlia, P., Vasei, M., Goliaei, B. & Nassiri, I. Normalized impact factor (nif): an adjusted method for calculating the citation rate of biomedical journals. *J. biomedical informatics* **44**, 216–220 (2011).
63. Bornmann, L. & Mutz, R. Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references. *J. Assoc. for Inf. Sci. Technol.* **66**, 2215–2222 (2015).
64. Beltagy, I., Lo, K. & Cohan, A. Scibert: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676* (2019).
65. Taylor, R. *et al.* Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085* (2022).
66. Hong, Z. *et al.* Scholarbert: Bigger is not always better. *arXiv preprint arXiv:2205.11342* (2022).
67. Wu, X. *et al.* A survey of human-in-the-loop for machine learning. *Futur. Gener. Comput. Syst.* (2022).
68. Wang, J., Guo, B. & Chen, L. Human-in-the-loop machine learning: A macro-micro perspective. *arXiv preprint arXiv:2202.10564* (2022).
69. Wang, Q. *et al.* Visual genealogy of deep neural networks. *IEEE transactions on visualization computer graphics* **26**, 3340–3352 (2019).
70. Qian, K. *et al.* Xnlp: A living survey for xai research in natural language processing. In *26th International Conference on Intelligent User Interfaces-Companion*, 78–80 (2021).
71. Coupette, C. & Hartung, D. Sharing and caring: Creating a culture of constructive criticism in computational legal studies. *arXiv preprint arXiv:2205.01071* (2022).