

Making a Computational Attorney

Dell Zhang* Frank Schilder† Jack G. Conrad Masoud Makrehchi
David von Rickenbach Isabelle Moulinier

Abstract

This “blue sky idea” paper outlines the opportunities and challenges in data mining and machine learning involving making a *computational attorney* — an intelligent software agent capable of helping human lawyers with a wide range of complex high-level legal tasks such as drafting legal briefs for the prosecution or defense in court. In particular, we discuss what a ChatGPT-like Large Legal Language Model (L³M) can and cannot do today, which will inspire researchers with promising short-term and long-term research objectives.

1 Introduction

The legal domain has always been an important application area of cutting-edge *data mining* and *machine learning* techniques. For the last three decades, research on legal AI has also been pushing the frontier of data mining and machine learning [8, 14, 19, 2, 17]. In comparison with other application areas, work with legal data is characterized by the following unique features.

- *Massive scale of complex text data.* For example, Thomson Reuters has accumulated over 60,000 TBs worth of data, a substantial portion of which is legal text data. There are also publicly available large legal corpora such as the 256GB Pile-of-Law dataset [9]. The majority of such data are in the form of long documents written in formal and professional language, such as legal judgments, legal opinions, and legal contracts.
- *High labeling cost.* For example, attorney fees in the US usually range from about \$100 to \$1,000 or more per hour, which makes the acquisition of gold-standard annotations from legal *subject matter experts* (SMEs) for training and testing legal AI models very expensive.
- *Emphasis on thoroughness as well as precision.* For example, legal research platforms like Westlaw and e-discovery tools usually put a high weight on *recall* but also demand high *precision* (see latest Westlaw Precision release).
- *Requirement of specialist knowledge.* For example, to qualify as an attorney in the US, one must typically complete 7+ years of post-secondary educa-

tion including 3 years at an accredited law school and then pass a difficult professional-license exam commonly known as the “bar exam” [4].

Apparently, the rise of *pre-trained large language models* such as BERT and GPT is causing a *paradigm shift* in data mining and machine learning. The legal domain is no exception. Hence, we propose to reexamine the current research agenda on legal AI by rethinking the question “what does it take to make a *computational attorney*?” The reason we coin a new term “computational attorney” instead of using the existing term “computational law” is to stress our anticipation for such capabilities to go beyond the automation of legal compliance management (e.g., based on *computable contracts* [16]) or other mundane legal information processing tasks (typically carried out by *paralegals*) and help human lawyers with complex high-level legal tasks (like drafting legal briefs for the prosecution or defense in court). Such computational attorneys are expected not to replace human lawyers but to work as their competent and reliable partners¹.

The success of this vision will change the legal industry (of \$300+ billion annual revenue in the US) in at least two aspects: (i) drastically improving the efficiency of millions of attorneys as well as law firms; and (ii) democratizing the legal services [3] in this law-dependent world where as many as 86% of low-income Americans with civil legal problems report inadequate or no legal assistance due to prohibitively expensive legal fees [5]. Conversely, the opportunities and challenges posed by the pursuit of a computational attorney can also motivate basic research objectives for data mining and machine learning which we will sketch out below.

2 Legal AI Research: The Trajectory

2.1 Past The usage of data mining and machine learning for legal tasks such as computer-assisted classification of legal abstracts dates back to 1990s [22]. In the past 30 years [8, 14, 19, 2, 17], a variety of data mining and machine learning techniques have been applied to automate relatively minor and repetitive low-level legal tasks, e.g., legal text classification and summarization, information extraction from legal documents, similar case matching, and litigation analytics. Major

*Thomson Reuters Labs, dell.z@ieee.org

†Thomson Reuters Labs, frank.schilder@thomsonreuters.com


¹<https://tinyurl.com/y9zgg6n>


data mining conferences such as ICDM have recently included workshops dedicated to this research field². Usually, a specific model will be developed to address a specific legal task.


2.2 Present The research on legal AI is undergoing a fundamental change at the moment: instead of *many small models* each for one specific task, researchers have started to build and utilize *one big model* for many different tasks. Such a big model, aka *foundation model* [5], for the legal domain, is a large language model either pre-trained on legal corpora from scratch or adapted from a general model with further pre-training on legal corpora [6, 25], which we call a Large Legal Language Model (L³M).


The characteristics of the legal domain mentioned in Section 1 suggest that L³M has huge advantages over traditional technical approaches to legal AI problems. On one hand, the massive scale of complex text data enables or facilitates the (self-supervised) pre-training of L³M. On the other hand, the *few-shot* prompting (i.e., *in-context learning*) or *zero-shot* prompting capability of L³M for downstream tasks can greatly alleviate or even avoid the high labeling cost, while the flexibility of L³M to accommodate ambiguity and idiosyncrasies can help to meet the challenges of thoroughness and specialized knowledge. It is not surprising that with L³Ms such as LEGAL-BERT [6] and Lawformer [21], we are seeing new heights achieved in legal text classification and other tasks [25, 15].

More importantly, when the scale of a L³M goes above a certain phase-change threshold, it will start to show some *emergent abilities* [20] that are present in only very large models but not smaller ones. It seems that *legal reasoning* (which is largely beyond the reach of traditional techniques) is one of these emergent abilities. L³M based *legal prompt engineering* (aka *legal prompting*) [23, 18] has demonstrated impressive performance and great potential on COLIEE entailment tasks where the questions come from Japan’s rigorous bar exam (which has less than a 40% passing rate in 2020). Furthermore, OpenAI’s recently released GPT-3.5 and ChatGPT³, though not adapted to the legal domain, has shown the preliminary ability of *legal reasoning* which is lacking in its previous version GPT-3 (see Figure 1). A comprehensive study of GPT-3.5 (text-davinci-003) shows that it achieved an average accuracy of above 50% on the complete MBE section of the National Conference of Bar Examiners (NCBE) practice exams [4], significantly exceeding the 25%

 X and Y contract over the sale of a Toyota Corolla. Their contract specifies that X has to pay Y one million dollars in damages if X does not deliver the car on time. X does not deliver the car on time. Is Y entitled to damages? If so, in what amount?

 Yes, Y is entitled to damages in the amount specified in the contract, which is one million dollars. If X does not deliver the car on time and Y suffers damages as a result, Y can seek to recover those damages from X. The damages will be in the amount specified in the contract, unless that amount is found to be unreasonable or the parties agree to a different amount. It is important for parties to a contract to carefully consider the terms of the contract, including any provisions for damages, before entering into the agreement.

 The current market price of a new Toyota Corolla is between \$20,000 and \$30,000. Is Y in the above contract entitled to damages? If so, in what amount?

 In the original scenario I described, Y is entitled to damages in the amount specified in the contract, which is one million dollars. The current market price of a Toyota Corolla is not relevant to the damages that Y is entitled to in this case.

The damages specified in the contract are intended to compensate Y for any losses or expenses incurred as a result of X’s failure to deliver the car on time. The damages may be intended to cover the cost of purchasing a similar car from another seller, the cost of renting a car, or any other expenses that Y has incurred as a result of X’s breach of the contract. The damages may also include additional amounts for inconvenience or other losses that Y has suffered.

It is possible that the damages specified in the contract could be considered unreasonable or excessive in relation to the actual losses or expenses incurred by Y. In that case, a court could reduce the damages awarded to a more reasonable amount. However, without more information about the specific circumstances of the case, it is not possible to determine whether the damages specified in the contract would be considered reasonable or unreasonable.

Figure 1: ChatGPT exhibits the sign of preliminary *legal reasoning* in answering a question about “liquidated damages” after additional information is supplied, which is beyond the ability of GPT-3 shown in [5].

baseline guessing rate. It has also been reported that ChatGPT could have passed a practice US bar exam with a headline score of 70% (35/50)⁴. If such general-purpose models are further enhanced by incorporating legal domain knowledge [24], they will be even more proficient at legal tasks.

Therefore, we now have good reason to believe that it would not take too long for a ChatGPT style L³M to carry out the above-mentioned simple low-level legal tasks (see Section 2.1) with human-like or superhuman performance, and thus render the corresponding problems within reach for scientific research. So, what’s next?

2.3 Future Although L³M have come a long way, we think that there is still a steep path ahead of them to master the intricate art of legal reasoning. Other than the commonly desired properties like *fast*, *economical*, *accurate*, *interpretable*, and *responsible*, the next gener-

²<https://www.m11d.cc/>

³<https://openai.com/blog/chatgpt/>

⁴<https://tinyurl.com/2m6yewed>

ation L³M need to make considerable progress on the following problems that are particularly acute for a computational attorney.

- *Updatable.* These models must be kept fresh and current in the relevant legal field if they are to remain reliable for legal reasoning. In particular, common law (case law) systems rely heavily on judicial precedents, so outdated or incomplete models may produce wrong legal analysis and results. Taking steps to ensure the timely updating of the models with new information from Westlaw “Court Wire” and “KeyCite Overruling Risk” etc. can therefore make a considerable difference in their ability to deliver sensible solutions in legal contexts. Since retraining the model from scratch to intentionally forget obsolete knowledge and incorporate additional desiderata would be very expensive (costing weeks of time and millions of dollars), we would instead prefer new methods to update the model *post-deployment*, which could yield significant savings over the simplistic strategy of retraining. There have been some early explorations of editing large language models [11, 12]. The techniques that have been developed for *stream data mining*, *machine unlearning*, and *life-long learning* (aka *never-ending learning* or *continual learning*) could be helpful.
- *Stable.* These models must know their limits and reason within the bounds of existing legal system in the relevant jurisdiction. In other words, we should take measures to prevent the models from “*hallucinating*”, i.e., inventing seemingly-plausible but non-existent responses by themselves. The techniques that have been developed for the quantification of *model uncertainty* aka *epistemic uncertainty* (such as *Bayesian deep learning* and *evidential deep learning*) and *stable learning* for *out-of-distribution* (OOD) generalization based on causality etc. might be useful.
- *Provable.* These models must be able to prove that their legal opinions or judgments are derived by strictly following the relevant rules as defined in the law [26], especially when they are challenged by the other side in court. This is a higher requirement than being explainable or being responsible, as the models not only have to explain what logical steps have been taken in their sophisticated legal reasoning process, they also need to identify the concrete applicable legislation to justify the correctness and fairness of each such step. It would not be good enough to just “get the job done” or find

what training instances or features are more important than others. Probably we can borrow and adapt some techniques from the field of *abductive reasoning* and *neuro-symbolic inference* that have been successful in the mathematical and medical domains.

- *Communicable.* These models must communicate effectively with fellow lawyers as well as legal clients to capture the subtle details and nuances of their requirements which may turn out to be crucial in court later. Furthermore, they should be able to learn not only from legal text documents but also directly from human instructions (i.e., be taught by human lawyers). ChatGPT has demonstrated the power of *reinforcement learning* from human feedback (RLHF) [13], but it is mostly passive rather than active/proactive. In addition to reinforcement learning, *active learning* and *socially situated AI* [10] have potential to play a role in this context as well.
- *Predictable.* These models must be predictable in the sense that they should not just be able to provide valuable legal advice based on sophisticated legal reasoning, but also to anticipate the implications of their outputs to stakeholders and the potential risks or liabilities associated with them in a legal ecosystem probably consisting of both human lawyers and computational attorneys. In order to win cases in court, the models should also be able to imagine the responses from the opponents and negotiate with their lawyers. This would require the combination of language models with strategic thinking, as illustrated by Meta’s CICERO [1]. Moreover, the models could suggest the timing and strategies for the optimal settlement of a dispute or the best terms of a contract [7], based on their calculations of *Nash equilibria*. The techniques from *algorithmic game theory* and *multi-agent systems* may come in handy.

3 Conclusion

In summary, we argue that despite the recent big advances in L³M for the automation of simple legal tasks such as legal text classification, the vision of a computational attorney capable of complex legal reasoning still serves as a “lode star” for data mining and machine learning in the legal domain and beyond. It is time for researchers to explore how to make legal or general AI models updatable, stable, provable, communicable, and predictable. The first two objectives of this endeavor seem to be in sight, while the remaining three could take many years before coming to fruition.

References

- [1] A. Bakhtin et al. Human-Level Play in the Game of Diplomacy by Combining Language Models with Strategic Reasoning. *Science*, 378(6624):1067–1074, Dec. 2022.
- [2] T. Bench-Capon, M. Araszkievicz, K. Ashley, K. Atkinson, F. Bex, F. Borges, D. Bourcier, P. Bourguine, J. G. Conrad, E. Francesconi, T. F. Gordon, G. Governatori, J. L. Leidner, D. D. Lewis, R. P. Loui, L. T. McCarty, H. Prakken, F. Schilder, E. Schweighofer, P. Thompson, B. Verheij, D. N. Walton, and A. Z. Wyner. A History of AI and Law in 50 Papers: 25 Years of the International Conference on AI and Law. *Artificial Intelligence and Law*, 20:215–319, 2012.
- [3] F. Bex, H. Prakken, T. van Engers, and B. Verheij. Introduction to the Special Issue on Artificial Intelligence for Justice (AI4J). *Artificial Intelligence and Law*, 25(1):1–3, Mar. 2017.
- [4] M. Bommarito II and D. M. Katz. GPT Takes the Bar Exam, Dec. 2022.
- [5] R. Bommasani et al. On the Opportunities and Risks of Foundation Models, July 2022.
- [6] I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, and I. Androutsopoulos. LEGAL-BERT: The Muppets Straight Out of Law School. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online, Nov. 2020. Association for Computational Linguistics.
- [7] J. G. Conrad and K. Al-Kofahi. Scenario Analytics: Analyzing Jury Verdicts to Evaluate Legal Case Outcomes. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law, ICAIL '17*, pages 29–37, New York, NY, USA, June 2017. ACM.
- [8] G. Governatori, T. Bench-Capon, B. Verheij, M. Araszkievicz, E. Francesconi, and M. Grabmair. Thirty Years of Artificial Intelligence and Law: The First Decade. *Artificial Intelligence and Law*, 30(4):481–519, Dec. 2022.
- [9] P. Henderson, M. S. Krass, L. Zheng, N. Guha, C. D. Manning, D. Jurafsky, and D. E. Ho. Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open-Source Legal Dataset, July 2022.
- [10] R. Krishna, D. Lee, L. Fei-Fei, and M. S. Bernstein. Socially Situated Artificial Intelligence Enables Learning from Human Interaction. *Proceedings of the National Academy of Sciences*, 119(39):e2115730119, Sept. 2022.
- [11] K. Meng, D. Bau, A. Andonian, and Y. Belinkov. Locating and Editing Factual Associations in GPT, Jan. 2023.
- [12] K. Meng, A. S. Sharma, A. Andonian, Y. Belinkov, and D. Bau. Mass-Editing Memory in a Transformer, Oct. 2022.
- [13] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe. Training Language Models to Follow Instructions with Human Feedback, Mar. 2022.
- [14] G. Sartor, M. Araszkievicz, K. Atkinson, F. Bex, T. van Engers, E. Francesconi, H. Prakken, G. Sileno, F. Schilder, A. Wyner, and T. Bench-Capon. Thirty Years of Artificial Intelligence and Law: The Second Decade. *Artificial Intelligence and Law*, 30(4):521–557, Dec. 2022.
- [15] D. Song, S. Gao, B. He, and F. Schilder. On the Effectiveness of Pre-Trained Language Models for Legal Natural Language Processing: An Empirical Study. *IEEE Access*, 10:75835–75858, 2022.
- [16] H. Surden. Computable Contracts. *UCDL Rev.*, 46:629, 2012.
- [17] H. Surden. Machine Learning and Law. *Wash. L. Rev.*, 89:87, 2014.
- [18] D. Trautmann, A. Petrova, and F. Schilder. Legal Prompt Engineering for Multilingual Legal Judgement Prediction, Dec. 2022.
- [19] S. Villata, M. Araszkievicz, K. Ashley, T. Bench-Capon, L. K. Branting, J. G. Conrad, and A. Wyner. Thirty Years of Artificial Intelligence and Law: The Third Decade. *Artificial Intelligence and Law*, 30(4):561–591, Dec. 2022.
- [20] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus. Emergent Abilities of Large Language Models. *Transactions on Machine Learning Research*, Aug. 2022.
- [21] C. Xiao, X. Hu, Z. Liu, C. Tu, and M. Sun. Lawformer: A Pre-trained Language Model for Chinese Legal Long Documents. *AI Open*, 2:79–84, Jan. 2021.
- [22] B. Yang-Stephens, M. C. Swope, J. Locke, and I. Moulinier. Computer-Assisted Classification of Legal Abstracts. In D. J. Hand, J. N. Kok, and M. R. Berthold, editors, *Advances in Intelligent Data Analysis*, Lecture Notes in Computer Science, pages 437–448, Berlin, Heidelberg, 1999. Springer.
- [23] F. Yu, L. Quartey, and F. Schilder. Legal Prompting: Teaching a Language Model to Think Like a Lawyer, Dec. 2022.
- [24] C. Zhen, Y. Shang, X. Liu, Y. Li, Y. Chen, and D. Zhang. A Survey on Knowledge-Enhanced Pre-trained Language Models, Dec. 2022.
- [25] L. Zheng, N. Guha, B. R. Anderson, P. Henderson, and D. E. Ho. When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset of 53,000+ Legal Holdings. In *Proceedings of the 18th International Conference on Artificial Intelligence and Law, ICAIL '21*, pages 159–168, New York, NY, USA, July 2021. ACM.
- [26] H. Zhong, C. Xiao, C. Tu, T. Zhang, Z. Liu, and M. Sun. 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*, pages 5218–5230, Online, July 2020. Association for Computational Linguistics.