Machine Learning in Legal Practice: Governance Models from Civil Discovery

Fernando Delgado
Information Science, Cornell University
September 20, 2018

Although often framed exclusively as a relatively new and controversial phenomenon, the application of machine learning techniques to legal analysis and decision-making in the US justice system has a rich yet under examined history. My research, a work-in-progress, examines how machine learning (ML) came to be adopted as a standard tool for automating manual legal document review and analysis workflows for complex, high-stakes civil litigation.[1] From 2008-2015, a period of robust ML experimentation, evaluation, and early adoption in the US civil justice system, issues surrounding efficacy, accountability, and transparency of algorithmic decision systems were hotly debated in and across the courts, legal scholars, professional standards organizations, as well as scientific research communities. This extensive interdisciplinary collaboration involved not only judges and lawyers, but also statisticians, ML experts, and other data specialists. Remarkably in line with recent calls for an interdisciplinary approach to better structure and critically examine the design and use of ML systems (Crawford 2016), the civil side of the US justice system underwent a robust process through which judges and lawyers learned from information retrieval and ML experts regarding the opportunities and limitations afforded by algorithmic decision-making systems; and conversely, technical experts learned from the legal profession how best to tailor and evaluate the systems they were developing to the specific problems of legal analysis and decision-making.

To best situate and articulate the controversies and consensus that emerged during this time, my work focuses on an analysis on 1) case law where ML is questioned or contested for civil discovery, 2) scientific literature evaluating the efficacy of ML for automating expert assessments on document relevance, 3) law review articles addressing the use of ML for civil discovery, and 4) governance and standards literature related to civil discovery published by legal professional associations. The analysis of case law in this study helps clarify the legal doctrine that emerges for the use of predictive coding, while the summary and analysis of both the law review and technical literatures shed light on the legal and scientific discourses both informing and reacting to the emerging case law. A supporting discussion of the standards literature provides a view into discovery and litigation practice more broadly as it existed prior to the emergence of ML applications for legal discovery, as well as how it evolves thereafter. These four sources in tandem allow for a fleshed-out picture to emerge illustrating the key interplay between judges, scholars, practitioners, and technologists that structure ML adoption and governance in the US civil justice system. In addition to clarifying the scientific, judicial, and litigation practitioner concerns on the eve and early phase of ML adoption, I examine the

[1] The most common designation for machine learning applied to discovery in civil litigation is *predictive coding*. *Technology-Assisted Review* (TAR) is also often used interchangeably with predictive coding yet technically speaking TAR refers to a broader set of information retrieval technologies used to manage electronic discovery efforts including but not limited to supervised machine learning. For terminological consistency, I will avoid the use of these industry terms and refer to the technology in question generally as ML for discovery.

role legal think tanks as well as institutions like the National Institute of Standards and Technologies (NIST) had on defining evaluation benchmarks and relevant data sets in the evaluation and early development of legal machine learning tasks.[2]

What we stand to gain from a comprehensive reassembling and critical analysis of ML adoption and governance in on the civil side of the justice system is an understanding of how best to leverage this rich precedent for current discussions regarding algorithmic decision-making on the criminal justice side (Angwin 2016; Koepke and Robinson 2018) as well as other public sector systems (Eubanks 2018). From a comparative perspective, it is worth noting that many of the questions and arguments raised about ML in civil discovery map closely to those raised in more recent discussions regarding the opacity and inscrutability of "black-box" algorithms deployed in criminal justice and public sector settings. One set of early findings in my research suggest that a key mechanism (and possible leading practice) for engendering trust in algorithmic decision-making was a dedicated, committed focus on establishing a shared set of meaningful metrics and transparent methods to evaluate ML system performance through which a diverse group of stakeholders could structure and communicate their findings, critiques, and ongoing dialogue.

From a theoretical standpoint, this work is grounded in the disciplines of science, technology, and society (STS) as well as critical legal studies (CLS). Methodologically, the current stage of my research is structured as a comprehensive analysis across a diverse corpus of court filings, technical journals, law review articles, and professional codes of conduct. Future planned stages of research extend out to ethnographic field methods including in-depth interviews with the key scientific and legal practitioners who were involved in the development of ML for civil discovery. The social scientific approach to this work is informed by over a decade of experience as a technologist in the legal sector developing and testing ML systems for automating legal review and analysis workflows. My ongoing research naturally leverages the insights and experiences from experience in industry, yet focuses on qualitative, critical, and interpretive modes of analysis examining ML technology not exclusively as isolated algorithms, but rather as sociotechnical systems in which algorithmic intelligence is embedded and deployed in larger institutional structures and professional practices (here specifically the US justice system and the professional practice of law). As such, although sensitive to the technical nuances of algorithmic and methodological designs in the development of ML systems, this work extends the state of current research not in terms of technological innovation for ML, but rather the legal structures of technology governance and ethical questions regarding the fairness, accountability, and transparency of ML systems.

[2] See TReC Legal Track. https://trec-legal.umiacs.umd.edu

References

Angwin, Julia et al. (2016). There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks (May 23, 2016). *ProPublica*.

Crawford, Kate (2016). There is a Blind Spot in AI Research. *Nature* 538(7625):311-313.

Eubanks, V. (2018). *Automating Inequality*. St. Martens Press.

Koepke, John Logan and Robinson, David G., Danger Ahead: Risk Assessment and the Future of Bail Reform (February 19, 2018). *Washington Law Review*, Forthcoming.