Submission to the Commission for Justice in Wales:
Research Summary on Legal Automation

Elliott Ash[1]

While (the option of) replacing a judge with a robot is still many years away, the technological trends that will lead to this eventuality have been in motion for decades. As early as the 1970s, computerized legal research services has begun to transform legal practice. Interestingly, many of the same debates about legal automation we hear today (in terms of replacement of human jobs) were common at the time (McCabe, 1971). For attorneys and for judges, WestLaw and other legal document databases were already becoming indispensable by the 1980s (Harrington, 1984).

Just as the prospect of a robot judge might change what we think of justice, digital research changed views on what "the law" is (Bast and Pyle, 2001; Hanson, 2002). Before the advent of digital databases, legal documents were organized according to key numbers indicating particular areas of the law. This created the perception of law as a discipline made up of timeless categories organized under general principles. Computer-based research allowed search across areas, fostering a new view of law as somewhat disorganized, chaotic, and changing. This new view of law focused more on the facts and doctrines of particular cases rather than on a hierarchical structure of legal principles (Bintliff, 1996).

Automated search through legal databases eventually gave way to automated search through digital document evidence. Recent reviews of this technology include Voorhees (2000) and Roitblat et al. (2010). While human reviewers tend to disagree about the relevance of documents, automated systems are now providing more consistent retrieval.

The most active stage for legal automation number is currently contract review and drafting (Lauritsen, 2007; Manna, 2014). By applying tools from computational linguistics, computers are now able to extract agents and requirements from contracts. These computer systems are faster and are becoming just as reliable as human readers. Ash et al. (2018b), in a social-science application, use a syntactic parser to extract rights and obligations from labor union contracts.

Talley (2017) provides a broader discussion of the potential future of a "driverless" legal system. This might include, as discussed in Love and Genesereth (2005), a"computational law" where legal rules are encoded in a formal logic that computers and A.I. programs can process and implement. The budding industry of smart contracts and cryptocurrencies illustrate the transformative potential of computerized law (Kiviat, 2015; Davidson et al., 2016).

These trends are impacting the judiciary as well. With some measure of evidence and some measure of the decision, it is straight-forward in principle to automate legal decision-making. First, one collects data on some or all of the previous cases, which consist of evidence-decision pairs. The model would then work to produce an approximation of the legal rules used by judges. If the approximation is close to the truth, then many cases will be predicted correctly. If it is far from the truth, then many cases will be predicted incorrectly.

[1]Assistant Professor of Economics, University of Warwick. Email: e.ash@warwick.ac.uk. Web site: elliottash.com.

Machine learning models are designed to gradually update the estimate to move gradually closer to the true legal rule. Eventually, the model will get as close as possible given the available evidence data.

In the case of the law, these models would take a set of case characteristics (evidence) and tell us how a judge would probably decide. This is a small but active research area. These papers include Ash et al. (2018a) (predicting bankruptcy court decisions with 67 percent accuracy and minimal case information), Katz et al. (2017) (predicting U.S. Supreme Court decisions with 70 percent accuracy), Chen and Eagel (2017) (predicting asylum courts with 82 percent accuracy), and Amaranto et al. (2017) (predicting prosecutor charge decisions with 88 percent accuracy).

Practically speaking, I envision a short-term implementation of this tool as a sort of robot clerk. It is an app that takes in evidence data, runs the numbers, and then produces a prediction about what previous judges would have likely decided. Human judges would then use this prediction as an input into their own decision, which could be based on a wider range of factors. In uncertain cases (near 50/50), the decision would still be made wholly by the human judge.

Many, probably most, cases in modern legal systems are clear-cut on the merits (e.g. Posner, 2008). A clerk with an hour or two of time could figure them out easily. But in many legal systems, including many state courts in the United States, the judges and clerks are overwhelmed with a massive caseload, and the clerks do not have these two hours to spare. If the algorithm can figure these cases out in a second or two, that would save a lot of time. If it improved consistency of decisions across judges in the meantime, that would be an added benefit.

These algorithms provide predictions about what the average judge would do based on the evidence. Any individual judge biases would be corrected. A more serious problem with the approach outlined here is that judges are biased *on average*.

Most of the research on systematic bias in the judiciary is on how Black defendants are treated in the U.S. criminal justice system, holding other factors constant (Fagan and Ash, 2017). Black defendants are stopped more often, given bail less often, are charged with more serious crimes for the same acts, are more likely to be convicted by juries, and given longer sentences by judges. These many biased legal decisions, taken together, also tend to result in disparities in socioeconomic outcomes (Alexander, 2012). In consequence a system with initially biased treatment could result in continued biased outcomes even when the initially biased procedures are corrected.

This matters for the robot judge because any automated decision system that is trained on biased data will also be biased. Whether currently implemented criminal risk scores, such as COMPAS, are in practice biased is an area of active investigation and debate. A major challenge is that many currently used risk metrics are proprietary, closed-source, and developed by for-profit companies motivated to defend the validity of the scores. Skeem and Lowenkamp (2016) examine a risk metric used in federal courts and find that while blacks and whites who are otherwise identical will be treated the same, blacks tend to be rated as more risky due to longer criminal histories. Given the pre-existing biases in the criminal justice system, this is exactly how automated risk metrics can reproduce bias.

In the case of both criminal and civil litigation, an automated decision procedure might allow legal actors to game the system. Savvy attorneys (including prosecutors) could learn to produce evidence that appears innocuous but fools the algorithm into deciding one way or the other. In terms of down-stream actions, individuals might learn that some types of (legally relevant) actions are ignored by the algorithm.

The brave new world of legal automation is upon us. These developments are both thrilling and unsettling because they attack the core of our humanity: Is not justice what distinguishes man from machine? These types of questions should be the subject of democratic debate, as should the other questions raised in this report. The brief history of automated legal decision-making points to democratized, non-profit, and open-source solutions.

References

Alexander, M. (2012). *The new Jim Crow: Mass incarceration in the age of colorblindness*. The New Press.

Amaranto, D., Ash, E., Chen, D. L., Ren, L., and Roper, C. (2017). Algorithms as prosecutors: Lowering rearrest rates without disparate impacts and identifying defendant characteristics noisy to human decision-makers.

Ash, E., Chen, D. L., Shang, F., Guan, X., and Yanchao, N. (2018a). Judge writing style predicts decision tendencies in u.s. bankruptcy courts. Technical report.

Ash, E., MacLeod, W. B., and Naidu, S. (2018b). Optimal contract design in the wild: Rigidity and discretion in collective bargaining. Technical report.

Bast, C. M. and Pyle, R. C. (2001). Legal research in the computer age: A paradigm shift. *Law Libr. J.*, 93:285.

Bintliff, B. (1996). From creativity to computerese: Thinking like a lawyer in the computer age. *Law. Libr. J.*, 88:338.

Chen, D. L. and Eagel, J. (2017). Can machine learning help predict the outcome of asylum adjudications? In *Proceedings of the 16th edition of the International Conference on Articial Intelligence and Law*, pages 237–240. ACM.

Davidson, S., De Filippi, P., and Potts, J. (2016). Economics of blockchain. In *Public Choice Conference*.

Fagan, J. and Ash, E. (2017). New policing, new segregation? from ferguson to new york. *Georgetown Law Journal*.

Hanson, F. A. (2002). From key numbers to keywords: How automation has transformed the law. *Law Libr. J.*, 94:563.

Harrington, W. G. (1984). A brief history of computer-assisted legal research. *Law. Libr. J.*, 77:543.

Katz, D. M., Bommarito II, M. J., and Blackman, J. (2017). A general approach for predicting the behavior of the supreme court of the united states. *PloS one*, 12(4):e0174698.

Kiviat, T. I. (2015). Beyond bitcoin: Issues in regulating blockchain tranactions. *Duke LJ*, 65:569.

Lauritsen, M. (2007). Current frontiers in legal drafting systems. In *11th International Conference on AI and Law*, pages 1–15.

Love, N. and Genesereth, M. (2005). Computational law. In *Proceedings of the 10th international conference on Artificial intelligence and law*, pages 205–209. ACM.

Manna, D. (2014). Artificial intelligence insourcing: Why software technology will dominate legal process outsourcing for routine document drafting. *Canadian Journal of Law and Technology*, 12.

McCabe, D. F. (1971). Automated legal research: A discussion of current effectiveness and future development. *J. Mo. B.*, 27:608.

Posner, R. (2008). *How Judges Think*. Harvard University Press.

Roitblat, H. L., Kershaw, A., and Oot, P. (2010). Document categorization in legal electronic discovery: computer classification vs. manual review. *Journal of the Association for Information Science and Technology*, 61(1):70–80.

Skeem, J. L. and Lowenkamp, C. T. (2016). Risk, race, and recidivism: predictive bias and disparate impact. *Criminology*, 54(4):680–712.

Talley, E. (2017). Is the future of law a driverless car? assessing how the data analytics revolution will transform legal practice.

Voorhees, E. M. (2000). Variations in relevance judgments and the measurement of retrieval effectiveness. *Information processing & management*, 36(5):697–716.