

The Promise of Machine Learning for the Courts of India

Sandeep Bhupatiraju (World Bank)

Daniel L. Chen (Toulouse School of Economics, World Bank)

Shareen Joshi (Georgetown University)¹

December 19, 2020

1. Introduction

Artificial Intelligence (AI) and machine learning (ML) – adaptive computer programs that can perform functions typically associated with the human mind – offer new opportunities for improving the productivity of large-scale organizations [1]. Recent advances in data collection, systems of aggregation, algorithms, and processing power have transformed computers from machines that perform rigidly defined tasks to machines that can learn without supervision and adapt to new inputs without the need to be reprogrammed. These advances have already brought tangible benefits in the world of business, medicine and large-scale systems more generally [2, 3].

AI platforms offer a variety of opportunities to improve justice systems [4-6]. The large volume of data from surveillance systems, digital payments platforms, newly computerized bureaucratic systems and even social media platforms can be analyzed to detect anomalous activity, investigate potential criminal activity, and improve systems of justice. AI systems can also reduce the barriers to accessing courts by providing users with timely information directly, rather than through lawyers. In the United States, AI is already used in the processing of bail applications, DNA analysis of crimes, gunshot detection, and crime forecasting. [5, 7]

For the purpose of this article however, we will focus on the potential of AI systems to improve data quality and the processes of justice themselves. This set of applications does not feature AI systems as replacements for human decision-makers. Rather, they help *improve* human decision-making capacity and productivity. A growing body of work demonstrates how small external factors, most of which the participants are unaware of, can leave their mark on the outcomes of legal cases. Analysis of the US courts for example, finds that the tone of words used in the first three minutes of a hearing, the incidence of birthdays, the outcomes of sporting events, and even the time of day of a hearing, affect the outcome of cases [4]. These methods can also identify both conscious and unconscious bias. The analysis of 18,686 judicial rulings, collected over 77 years, by the 12 US circuit courts (also known as courts of appeals or federal appellate court) illustrated that judges demonstrate considerable bias before national elections [8]. Similarly, there is new evidence that sequencing matters in high-stakes decisions: decisions made on previous cases affect the outcomes of current cases, even though the cases have nothing to do with each other. Refugee asylum

¹ **Sandeep Bhupatiraju**, World Bank, 1818 H Street Washington, DC 20433. Email: sbhupatiraju@worldbank.org. **Daniel L. Chen**, Toulouse School of Economics, Université de Toulouse Capitole, Toulouse, France and World Bank, 1818 H Street Washington, DC 20433. Email: dlchen@nber.org. **Shareen Joshi**, Walsh School of Foreign Service, Georgetown University, 3700 O Street, Washington DC, 20057. Email: shareen.joshi@georgetown.edu

judges are two percentage points more likely to deny asylum to refugees if their previous decision granted asylum [9].

AI systems have enormous potential to improve the processes of justice in India. Human capacity has already been identified as a significant constraint in the system. India has just 19 judges per million people, and 27 million (2.7 crore) pending cases [10-12]. The justice system has already made considerable advances in adoption of information technology, released large volumes of data to court users and encouraged them to use electronic systems. Yet, legislative, institutional, and resource constraints have limited their full impact [12, 13].

We argue that the integration of machine learning tools with newly available legal data offers a mechanism to identify biases in judicial behavior and propose real-time corrections for these behaviors, thereby streamlining the system and reducing backlog. Such tools can also identify discrimination and bias, even when these are not evident to the participants in the courts themselves, thereby strengthening the credibility of the judiciary [14-16]. The availability of data enables new varieties of academic research on the efficiency and effectiveness of the legal system itself: macro-level studies conducted as “distant reading” of the system, rather than micro-level “close reading” as commonly done in legal studies.

The adoption of AI systems however, is not a panacea for a court system. The technological sophistication of the tools creates what is often called the “black box” problem [17]: the technological sophistication of the tools makes them less interpretable to many. The challenge of interpretability also raises concerns about the accountability and oversight for these systems. Furthermore, the gap between those who can and cannot access and understand these technologies also exacerbate existing social divisions and intensify polarization. For all these reasons, we refrain from advocating that the tools of AI and ML be used to replace human decision-making through for example, automation of bail applications. Rather, we believe that the systems be leveraged to *aide* and *improve* human decision-making within the system. We believe that the adoption of AI systems stimulates reflection on questions that are truly fundamental to the systems of justice and the benefits of the system are considerable when leveraged in a careful and ethical way.

2. The Promise of New Data

In the past fifteen years, considerable efforts have been made to adopt and deploy information technology systems in the courts of India. One of the most significant projects, the e-courts project, was first launched in 2005 by the Supreme Court of India through the “National Policy and Action Plan for Implementation of Information and Communication Technology (ICT) in the Indian Judiciary”.

The e-courts initiative introduced technology in Indian court rooms in a variety of ways. The most obvious and observable feature of the system was the deployment of technology within the courts itself. Judges were provided with LCD touch screen machines, screens and projectors were connected via a local network to disseminate information to lawyers, electronic boards at the courts display the queue of case numbers for hearing scheduled on a particular day, etc. Outside of the courtroom, e-filing procedures were established and an architecture of data management was created. These investments have eventually paved the way for the creation of the National Judicial Data Grid (NJDG), a database of 27 million cases available to court users to view the status of pending cases and access information on past hearings.

For our purposes, the most significant resource available through this has been the digital archives of cases. We were able to scrape these publicly available digital archives to construct an e-Courts district court dataset of 83 million cases from 3289 court establishments.² We were able to curate details like the act under which the case is filed, the case type, district where it originates, the parties to the case, and the history of case hearings, in a manner that makes the data amenable to large-scale analysis.

We supplement this database with a variety of other data sources. Three examples of these are below.

Data on Judges: In order to better understand the impact of specific judges – their identity, training and experience – we have constructed a database of judges for the courts of India. We have begun this task by extracting data from the Judges Handbooks, released by the Supreme Court of India, and appended to it information from various High Court websites. Thus far, we have assembled details of 2,239 judges from the handbooks for the years between 2014 and 2020. Most notably, 93.5% of these judges are males and 6.5% are females and their collective experience at the courts of India covers approximately 70 years.

Database of Central Acts: This auxiliary dataset is intended to give a definitive list of standardized act names. This could then be used to standardize the act names appearing the various cases databases. This allows us to analyze all cases filed under a given act. We have, for example, examined all cases related to the Water Act of 1974 and found a total of 978 such cases at the Supreme Court and High Courts of India. The list of central (Federal) acts can be viewed on the Legislative Department website of the Ministry of Law and Justice. There is currently no central source for all state legislation – this needs to be obtained from the state websites separately.

Other Administrative Data: Data on other institutions can be linked to the judicial data at the district as well as the state level. For example, data on Indian banks and their branches is available through the Reserve Bank of India. This database contains information on the name, location, license number, license date, address, and other unique identifiers. We have scraped, cleaned and organized this data for further analysis. It contains about 160k entries. The unique identifiers and location information allow us to merge this data with banks appearing in litigation in courts that are present in the e-courts databases. The merging of this data with the legal data allows us to examine a variety of interesting questions about the development of financial markets in an area, participation in the system of justice, and the impacts of legal rulings.

3. Potential Applications of Machine Learning in the Courts of India

Legal data released by the Indian judiciary are voluminous, messy, and complex [13]. Cleaning and pre-processing the data is critical for any form of analysis, especially so for supervised algorithms trained on this data. Traditional empirical legal studies have typically addressed this issue by relying on small-scale data sets, where legal variables are manually coded, and the scope of inference is related to a small body of legal cases that is pertinent to a single issue [18-20].

One of the biggest challenges in pre-processing the data is the variability of reporting across states and districts. The quality of the data varies significantly – there is no nationally standardized system for defining variables or reporting on them. This makes it difficult to compare individual case-types across courts and

² The eCourts data is public and can be accessed via the district court websites, the eCourts Android/iOS app, or the district court services webpage.

across states [13]. There are no standardized identifiers within the data to follow a case throughout its potential sequence of appeals in higher courts. In a similar vein, there is no easy way to track a criminal case from its entry into the system as a FIR to its exit as a judgment [13].² There are inconsistencies in identifying information about participants, their attributes and the types of laws or acts that the case relates to. There are also issues of incorrect reporting and spelling mistakes. Entries from one field can sometimes show up in another, requiring careful cleaning and systematic recoding of variables.

Various ML tools can be harnessed to improve data quality and address the issues discussed above. We have constructed a robust pipeline to scrape, clean and prepare this data for analysis. We briefly describe some of these methods below.

3.1 Inference about the identity of participants

Some databases of judgements provide no identifying information on the participants in the cases themselves. To better understand who participates in the courts, we first extract litigant names from the raw text of the judgements, and then using some matching algorithms we identify the type of litigant (individuals, companies or state institutions). Classifying participants can be challenging. If we are interested in cases that involve the government for example, we must be able to identify all the different agencies of the state government, national government and the additional agencies that fall within the definition of “state” under Article 12 of the Indian constitution. Manually tagging these entities is prohibitively time consuming and the existence of latent patterns in the names makes this fertile ground for Machine learning applications.

Once we extract names, we can draw inferences about the characteristics of individual participants. Some obvious attributes of interest are a participant’s gender, caste and religion. These attributes however, are not officially recorded in court proceedings. Again, ML methods provide a possible solution. We have focused here on people first and last names, also known of course as surnames, for illustration.

We first format individual names to ensure that each individual could be identified by an honorific title, a first name and a last name. Honorifics such as Shri, Sri, Smt, Mr, Mrs, and Ms. enabled us to directly identify gender. To extend this classification to names without an honorific, we train an algorithm on a publicly available corpus of labelled common Indian first names. Training this algorithm, often referred to as training a classifier, is the process of learning patterns within the data related to the classification. Here, these patterns are the statistics of co-occurrence of alphabets in names, length of the name and other features which have some predictive information. In order to reduce the generalization error, we use the majority vote from multiple trained classifiers, including a logistic regression model and a random forest classifier to make predictions on gender.^{3,4}

³ The features (x-values) in all the models were hand-engineered co-occurrence statistics of blocks of alphabets in various locations within the names.

⁴ The voting procedure is a way to ensemble models so as to reduce the generalization error. For instance, if we had three prediction algorithms for gender and all of them make a prediction of M or F, then we use the majority vote as the final prediction of the ensemble. In this case, at least two of the algorithms would have predicted the same class and we use this as our final prediction. We used the cross-entropy function as the loss function to quantify how well a given classifier did and obtained an accuracy of 0.92 for the ensemble model.

We have also made predictions of religion and caste using similar approaches. Muslims can be recognized in the data through the distinctiveness of Muslim names: common names such as *Khan* and *Ahmed* can easily be assigned and coded, but for others we utilize the occurrence of specific alphabets (such as Q and Z) through appropriate classifiers, to identify additional names. Caste assignment is more complicated because the same last name can be associated with multiple caste groups. The name *Kumar*, for example, could be the name of a person belonging to SC, ST or Other category. In the case of such names, we generate the distributions of the last name across the different caste categories. We use this distribution to generate a prediction in an ensemble along with predictions from other multi-class multi-label prediction algorithms. We assign a caste to each household based on a simple majority vote between these models.

3.2 Identification of Laws and Acts

Legal texts do not currently employ any standardized citation style for referring to acts or laws. For example, the Hindu Marriage act may be referred to in a variety of ways: “u/s 13 clause 2 of hindu marriage act,” “u/s 13 {b} hindu marriage act,” or “u/s 13 of hindu marriage act 1995”. Again, ML tools can also be used to address this issue.

In our work, we are using a set of tools that create mathematical representations of the text in the form of vectors. "Term Frequency - Inverse Document Frequency" (TF-IDF) is one such popular method to represent a string of words as a vector of scores that reflects how frequently a word is used within a given text and how infrequently it appears in the corpus.⁵ With this representation of the act names as vectors we use unsupervised clustering algorithms like ‘k-means’ and agglomerative hierarchical clustering to infer similarities between act names from closeness of the associated vectors.⁶

Identification of specific laws and acts opens up new opportunities of legal analysis. We can, for example, compare the types of cases across courts, and the time it takes to process them. This is critical in being able to understand where the biggest delays in the system come from, and how they may be resolved.

3.3 New interpretations of text

The ideology of judges, and how this affects their rulings on specific cases has fascinated legal scholars for a long time. [21-23]. Empirical estimates of ideological imprints on cases and rulings however, are difficult to estimate. Recent work in natural language processing and computational linguistics however, makes this a new possibility [24]. Instead of looking at clusters of specific words, as discussed above for acts and laws, algorithms within a class called neural networks can examine the contexts in which words are used and produce representation of these words as vectors such that words used in similar contexts are close together in this vector space. Interestingly, mathematical relations between these vectors encode semantic relations between the encoded words. For example, Word2Vec is one such algorithm that “learns” conceptual relations between words; a trained model can produce synonyms, antonyms, and analogies for a given word. These vectors, often referred to as “word embeddings” can be used to identify structure in the data and make predictions. More recently, “document embeddings” have built upon the success of word embeddings

⁵ “Term frequency” measures the number of times a term appears in a document. Inverse Document Frequency refers to the $\log(N/D)$, where N is the total number of documents, and D is the number of documents that contain that specific term. The TF-IDF is the product of these two terms.

⁶ It first assigns "K" random data points (TF-IDF vectors generated from text of Acts in our case) as cluster centroids (or means) and then categorises rest of the data points to the means closest to them. After this, a new position of the centroid (mean) is calculated, taking the average of the data points categorized in that mean. This process goes on iteratively till the centroid stops moving, or say is at the centre of the final cluster.

to represent words and documents in a joint geometric space [25]. Like word embeddings, these document embeddings can be used to interpret and classify large volumes of text.

The application of these algorithms on a corpus of cases allows for the identification of important patterns. In the United States for example, they uncover the distinctive ideology of judges, and the variations in this ideology based on birth cohort, partisan affiliation and or legal training [26]. For the embeddings approach to the citation network, we can identify similar cases based on how often they are cited together. A more intriguing example is identifying legal analogies through document embeddings. Word embeddings have been documented to know that man is to woman as king is to queen, from the way the four words are used in the English language. A document embedding might identify when the latter case is a related application of a legal principle articulated in the former case.

3.4 Identification of Discrimination and Bias

Concerns about stereotyping and discrimination in the courts of India typically cite a single case or small set of cases to highlight the issue. Bias shown by a court, or even a single judge, is difficult to identify and analyze rigorously. This challenge is of course not unique to judiciaries. A large academic literature demonstrates that bias by a human decision-maker can have conscious as well as unconscious drivers and may manifest in complex ways than can be difficult to prove in a variety of contexts [27-30]. In other settings, such as labor markets and educational institutions, algorithms – rules that take “inputs” (like the characteristics of a job applicant) and predict some outcome (like a person’s salary) – have been shown to create new forms of transparency and serve as methods to detect discrimination [31, 32]. In the courts of India, algorithms could help judges make some critical decisions about cases (for example, dismissals or bail applications). They could also help courts evaluate a judge’s performance.

Building such algorithms requires a feature-rich dataset typically consisting of variables that include litigant characteristics (caste, gender, location, type of crime committed), lawyer characteristics, court characteristics, case details (filing details and evidence provided), additional variables (day, month, year, weather, etc.) and case outcomes (such as granting of bail or dismissal of a case). An algorithm builder would write a “learning procedure” that would aim to provide a predicted outcome from a broad range of inputs choosing from a variety of models like support vector machines, decision trees, Bayesian networks and neural networks. The architecture of these models feature multiple sequential layers of intermediate variables that connect input features and output classes, so that the outputs of one layer serve as inputs of the next layer [33]. These models stand in sharp contrast to traditional statistical methods such as linear regression.

These insights could be invaluable not only within the court room itself, but also in judicial education. Experiments are currently underway, in the Judicial Academy of Peru for example, to assess methods to improve case-based teaching [34]. The data is also suitable for creating personalized dashboards and interfaces that provide judges, mediators and other decision-makers with real-time information about their own performance relative to their own previous decisions and others who are comparable [34]. This information can be used to augment the capabilities of judges and lawyers, increase their productivity and devote greater attention to complex cases.

3.5 Identification of Causal Effects of Legal Rulings

Legal scholars and judges have long made arguments about laws and regulations and justified their arguments with theories about the effects of these legal rules. This situation resembles the field of medicine a century ago: prior to the advent of clinical trials, there were only theories without rigorous causal evidence.

A growing body of empirical research now demonstrates that causal inference is possible in judicial studies. For example, in situations where cases are randomly assigned to judges, the random assignment itself can be used as an exogenous source of variation to evaluate the impact of judicial decisions – since judges in this scenario do not choose their cases and end up with them randomly, observed rulings reflect their deliberations in the case rather than the process of justice that led them to be assigned the case.

Randomizing cases to judges predicted to be harsh or lenient generates the inference on the long-run causal impacts of the length of sentences [35]. The same framework can examine the causal effects of debt relief on individuals' earnings, employment, and mortality [36]. Discretion in decision-making sheds light on a myriad of topics where judicial habits exist such as patent protection [37]. Machine learning can be used to create predictions of judges, which can then be used to estimate the causal effects of their predicted decisions on long-term outcomes

4. A New Relationship Between Human and Machine

Thus far, we have argued that ML presents a powerful tool for improving the systems of organizing and interpreting the voluminous, unstructured and complex data that has been released from the Indian judiciary over the past 15 years. Algorithms can be written to draw inferences about the identity of participants and study the deliberative processes they employ within court rooms. ML tools can also convert high volume of textual data to numerical estimates that can be used for understanding the processes and outcomes of the systems of justice themselves. Analysis of texts for patterns of bias and discrimination for example, can augment the capabilities of judges and lawyers and systematize processes of review.

These tools however, have several limitations and requirements that need to be addressed before they can be effectively deployed in the courts. At the very outset, there are significant issues related to the privacy, security, and control of legal data. Next, the algorithms require data pre-processing, training on large and high-frequency datasets, and iterative refinement with respect to the actual cases where they are deployed. This requires strong pilot programs that are studied as part of randomized control trials (RCTs). Insights on data privacy, cost as well as outcomes require these pilots to be constructed on a reasonable scale.

Only technologies that aid humans, rather than replace them, should be adopted in the courts. This is for a variety of reasons. As noted earlier, algorithms have the “black-box” problem of interpretability, i.e. it is not easy to trace the output of complex algorithms to the data inputs themselves [17]. The word-embeddings algorithms discussed earlier for example, learn biases existing in the corpora. Using these in downstream tasks and decisions without critical oversight raises the risk of replicating these biases elsewhere in the system. The inherent choices of model architecture can also reinforce existing biases by decision-makers within the system. Addressing these issues requires a participatory and deliberative approach towards the design, implementation and evaluation of the adoption of these technologies.

How may AI and ML methods aide human decision-making? An AI based recommender system might start by offering a judge the best prediction of themselves, based on the specific judge's previous decision-making, from a model using only legally permitted features. At the very least, it can help judges be consistent across similar cases by offering the most relevant reference points—and to cabin the influence of extraneous factors. Deviating from defaults can facilitate more conscious Type II deliberation. Moreover, by showing the judge statistics about oneself, it explicitly leverages motives for judges to want to self-improve.

Showing how other judges might make the decision—based on the machine's model of the other judges—offers yet another way for AI to assist rather than dictate decision-making, this time by creating a customized community of instantly available experts (trained on data from the behavior of other experts, possibly over time and across geographic and subject matter contexts). Showing statistics of other judges may also leverage self-image concerns of being a better judge that can motivate moral behavior.

A predictive algorithm to detect judicial error might nudge judges into more appropriate decisions. The predicted errors allow for analysis of the extent to which systematic factors like affect the predictability of judges' errors. By predicting error, we might highlight areas in which decision-makers potentially need more decision support. Of course, it is always possible that the AI system's suggestion would fail to take into account some reliable private information that the judge might have access to. Where this happens, the AI system would be steering the judge off course rather than correcting for their inconsistencies. Therefore, a dialectic, so to speak, can be encouraged between the judge and the AI system, allowing for the AI to learn from the judge as well.

A reasonable demand to guarantee trust and fairness is that algorithms be interpretable. A judge may request a reason for why the deviation may lead to mistakes. Such an incremental integration leverages judge's normative commitments and self-perceptions of being a good judge to facilitate the adoption of these systems.

Given the complexities of working with AI and ML algorithms, it is essential that any rollout be preceded by a phase of comprehensive study and rigorous testing of the systems themselves. Randomized controlled trials that carefully estimate the direct and indirect impacts of the adoption of these algorithms, as well as their cost-effectiveness, are essential. A carefully constructed trial can provide important benchmarks on cost, efficiency, user satisfaction and outcomes, all essential for a justice system to credibly serve citizens. Overall, we believe that when leveraged in a careful and ethical way, AI and ML offer significant potential for the courts of India.

References

- [1] S. Russel and P. Norvig, "Artificial intelligence: a modern approach", 2013.
- [2] E. Brynjolfsson and A. McAfee, "The business of artificial intelligence," *Harv. Bus. Rev.*, pp. 1-20.
- [3] J.H. Chen and S.M. Asch, "Machine learning and prediction in medicine—beyond the peak of inflated expectations," *N.Engl.J.Med.*, vol. 376, no. 26, pp. 2507.

- [4] D.L. Chen, "Judicial analytics and the great transformation of American Law," *Artificial Intelligence and Law*, vol. 27, no. 1, pp. 15-42.
- [5] C. Rigano, "Using artificial intelligence to address criminal justice needs," *National Institute of Justice.Issue*, no. 280.
- [6] D.L. Chen, "Machine learning and the rule of law," *Computational Analysis of Law, Santa Fe Institute Press, ed.M.Livermore and D.Rockmore, Forthcoming*.
- [7] W.J. Epps Jr and J.M. Warren, "Now Being Deployed in the Field of Law," *The Judges' Journal*, vol. 59, no. 1, pp. 16-39.
- [8] C. Berdejo and D.L. Chen, "Electoral cycles among us courts of appeals judges," *The Journal of Law and Economics*, vol. 60, no. 3, pp. 479-496.
- [9] D.L. Chen, T.J. Moskowitz and K. Shue, "Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires," *The Quarterly Journal of Economics*, vol. 131, no. 3, pp. 1181-1242.
- [10] V.A. Kumar, "Judicial Delays in India: Causes & Remedies," *JL Pol'y & Globalization*, vol. 4, pp. 16.
- [11] M. Chemin, "Does court speed shape economic activity? Evidence from a court reform in India," *The Journal of Law, Economics, & Organization*, vol. 28, no. 3, pp. 460-485.
- [12] A. Amirapu, "Justice delayed is growth denied: The effect of slow courts on relationship-specific industries in India,".
- [13] D. Damle and T. Anand, "Problems with the e-Courts data," *Working Paper No. 314*.
- [14] K. Kannabiran, "Judicial meanderings in Patriarchal thickets: Litigating sex discrimination in india," *Economic and Political Weekly*, pp. 88-98.
- [15] M. Galanter, "Competing equalities: law and the backward classes in India,".
- [16] P. Bhushan, "Misplaced priorities and class bias of the judiciary," *Economic and Political Weekly*, pp. 32-37.
- [17] F. Pasquale, "The black box society," 2015.
- [18] A. Bhuwania, "Courting the people: Public interest litigation in post-emergency India," vol. 2, 2017.
- [19] V. Gauri, "Public interest litigation in India: overreaching or underachieving?," 2009.
- [20] S. Krishnaswamy, S. K Sivakumar and S. Bail, "Legal and judicial reform in India: a call for systemic and empirical approaches," *Journal of National Law University Delhi*, vol. 2, no. 1, pp. 1-25.
- [21] U. Baxi, "Towards a sociology of Indian law," vol. 77, 1986.

- [22] G.H. Gadbois Jr, "Judges of the Supreme Court of India: 1950–1989", 2011.
- [23] G.H. Gadbois, "Indian supreme court judges: a portrait," *Law Soc.Rev.*, pp. 317-336.
- [24] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado and J. Dean, "Distributed representations of words and phrases and their compositionality," *Advances in neural information processing systems*, vol. 26, pp. 3111-3119.
- [25] Q. Le and T. Mikolov, "Distributed representations of sentences and documents," pp. 1188-1196.
- [26] E. Ash and D.L. Chen, "Case vectors: Spatial representations of the law using document embeddings," *Law as Data, Santa Fe Institute Press, ed.M.Livermore and D.Rockmore*.
- [27] D. Kahneman, "Thinking, fast and slow," 2011.
- [28] A. Banerjee, M. Bertrand, S. Datta and S. Mullainathan, "Labor market discrimination in Delhi: Evidence from a field experiment," *Journal of comparative Economics*, vol. 37, no. 1, pp. 14-27.
- [29] M. Bertrand and S. Mullainathan, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *Am.Econ.Rev.*, vol. 94, no. 4, pp. 991-1013.
- [30] M. Ewens, B. Tomlin and L.C. Wang, "Statistical discrimination or prejudice? A large sample field experiment," *Rev.Econ.Stat.*, vol. 96, no. 1, pp. 119-134.
- [31] J. Kleinberg, J. Ludwig, S. Mullainathan and C.R. Sunstein, "Discrimination in the Age of Algorithms," *Journal of Legal Analysis*, vol. 10.
- [32] J. Kleinberg, J. Ludwig, S. Mullainathan and C.R. Sunstein, "Algorithms as discrimination detectors," *Proceedings of the National Academy of Sciences*.
- [33] Y. LeCun, Y. Bengio and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444.
- [34] D.L. Chen, M. Ramos and S. Bhupatiraju, "DATA AND EVIDENCE FOR JUSTICE REFORM (DE JURE)," *Development Impact Evaluation (DIME) Group, World Bank, November*.
- [35] J.R. Kling, "Incarceration length, employment, and earnings," *Am.Econ.Rev.*, vol. 96, no. 3, pp. 863-876.
- [36] W. Dobbie and J. Song, "Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection," *Am.Econ.Rev.*, vol. 105, no. 3, pp. 1272-1311.
- [37] B. Sampat and H.L. Williams, "How do patents affect follow-on innovation? Evidence from the human genome," *Am.Econ.Rev.*, vol. 109, no. 1, pp. 203-236.