Silicon Flatirons
KNOW WHAT'S NEXT.

Outcomes Report

AI Initiative Conference: The State of the Art of Artificial Intelligence in the Practice of Law

November 9, 2021
Zoom Webinar
Event archive: https://siliconflatirons.org/events/the-state-of-the-art-of-artificial-intelligence-in-the-practice-of-law-artificial-intelligence-initiative-conference/

Prepared By

Xelef Botan
JD Candidate, University of Colorado Law School; Researcher, Silicon Flatirons

Harry Surden
Professor of Law, University of Colorado Law School
Artificial Intelligence Initiative Director, Silicon Flatirons

Published on May 2, 2024

Silicon Flatirons Center for Law, Policy, and Entrepreneurship
University of Colorado Law School
2450 Kittredge Loop DR | Boulder, CO 80309
siliconflatirons.org

Contents

Executive Summary of the Conference

Artificial Intelligence (AI) is increasingly associated with enhanced decision-making processes that can reduce costs and increase efficiency within different fields. As such, AI systems have increasingly become important tools used across the private and public sectors. As AI techniques have become more popular, the legal profession has also started to increase its usage to benefit from the services that AI systems can provide.

On November 9, 2021, the AI Initiative of the Silicon Flatirons Center for Law, Technology, and Entrepreneurship at the University of Colorado Law School hosted a conference to discuss and analyze AI's use in the Practice of Law.[1] In the first session of the conference, panelists examined what AI tools were being used by lawyers in the transactional practice, those tools' popularity, and the role and impact of AI in legal transactional practice. In the second session of the conference, panelists analyzed AI in the context of civil and criminal litigation, as well as whether AI was outperforming manual methods used in the discovery process and implications for prosecution and defense branches. In the third and last session of the conference, panelists discussed the likely trajectory of AI and law over the next five years and basing their realistic predictions on the current capabilities that AI technology possesses.

[1] *See The State of the Art of Artificial Intelligence in the Practice of Law, Silicon Flatirons Center* (Nov. 2021), https://siliconflatirons.org/events/the-state-of-the-art-of-artificial-intelligence-in-the-practice-of-law-artificial-intelligence-initiative-conference/

Part I: AI in Transactional Practice

The panel began by observing that AI's reach and actual use is sometimes exaggerated by parties seeking to attract attention or market products. The goal of the session was to provide an honest assessment of the actual capabilities of AI, and its limitations, that is grounded in evidence. Outside of the law, some of the most promising aspects of AI have already been employed in self-driving vehicles and language translations. These tools, which are built upon machine learning techniques, involve identifying patterns in large volumes of data.

Professor Surden sought to situate the actual state of AI based upon the current evidence. In many cases, the public's perception of the capabilities of AI does not reflect the underlying reality. Despite much discussion in the popular press, there is no evidence that AI's current state is approaching 'strong AI'—the state at which AI systems would meet or exceed human thinking. Prof. Surden emphasized that the current evidence indicates AI will not be at that 'strong AI' level of abstract thinking capabilities within the next five years. He noted that predictions about the state of AI beyond the five-year timeframe are highly speculative and unlikely to be reliable. However, he indicated that multiple significant theoretical and technical hurdles would need to be overcome for AI to begin approaching strong-AI status. Rather, AI has achieved significant achievements, but often in much more modest and limited areas than commonly envisioned, involving image classification, pattern matching, business prediction, and large-scale language models for generating text.

Where AI is Being Successfully Used in Transactional Practice

Some legal professionals expected AI to usher in a sea change in legal practice. What has happened instead, is a gradual evolution of some tools that incrementally incorporate AI technologies. One area in which AI-enabled tools are gaining traction in transaction practice involves document review. AI-enabled document review is heavily employed in real estate, mergers and acquisitions, and contracts fields. These documents are analyzed using natural language processing (NLP) tools. Such NLP techniques have been around for a while but have recently become more sophisticated.

However, despite their increased sophistication, these tools still require 'humans-in-the-loop,' whereby humans act as supervisors. In essence, these tools increase efficiency by allowing humans to go over a large volume of documents much more quickly and saving them valuable time.

There have also been new improvements in the areas of 'semantic queries' to carry out more comprehensive searches within the patent law

area; 'elements queries,' where text and image as separate elements are used together to do more thorough searches within the trademark area; and billable hours tracking and narratives by attorneys, making them a lot more accurate, in particular, in the in-house counsel departments of large corporations. In other words, most AI systems in law seem to be serving as decision support systems for humans, rather replacing human decision-making with automation.

New tools for 'prediction' are also being developed. Within employment law, for instance, inputting certain facts can engender a prediction on whether someone will be treated as an employee or an independent contractor. Intellectual Property (IP) is another area, especially patent law, where improvements in AI have made it possible to get a prediction more accurately on whether an idea is patentable using various criteria on a granular level. IP and patent law are early adopters in AI.

An area of AI that is not much talked about, for example, is its use in matchmaking between lawyers and clients, which involves using algorithms to match the right client with the right lawyer, as well as matching fortune 500 companies with top 100 law firms.

Computable contracts, though not employed widely yet, is an area where further improvements are expected. This is an area that data-driven machine learning NLP AI is really struggling with, and a rules-based approach could come in handy for insurance policies, for instance, where people can get answers to why their claims were accepted or refused.

To reiterate, smart drafting, contextual searching, guided reading, and matchmaking are some of the most common areas that AI is being employed.

The Limitations of AI

One of the panelists stated that limitations exist more in some attorneys' understanding/thinking that AI systems would function as attorneys, [i.e.: input a question and receive an answer,] rather than in the tools themselves. There have already been great improvements in AI tools. Expecting those tools to replace human attorneys is unrealistic and raises a lot of ethical questions as well.

If the data being put into a machine is misinterpreted or misapplied by people who don't know how to properly handle it, AI tools could be much worse than no tools at all. For example, when a machine translation is presented to a regular human translator, the sentence might be awkward but not incorrect grammatically, so it is left unchanged. However, if humans were the first to translate it, they would translate it differently and more accurately. In that kind of a scenario, not having the AI tool could end up being a much more accurate decision.

In practice there are two types of lawyers: one type consists of skeptics, who do not believe in the power of AI technology at all, while the other type is made up of true believers, who expect AI to be able to do everything. Reality, in fact, is somewhere in between the conceptions of the two groups.

Some Misconceptions of AI

A common misconception is that AI systems are completely automated, whereas, in reality, there are very limited instances where there is no human-in-the-loop. Even the most advanced self-driving vehicles company, Waymo, has a call-center where humans monitor the cars and help them get out of trouble, which they almost inevitably end up in at some point.

A context-aware robot of superintelligence is what most people think of because of the movies, tv shows, and books they were exposed to. However, the reality is that we are far away from that point. In the legal world, there is talk of developing an AI machine that can pass the bar exam, more specifically the contracts portion. On the one hand, there are practitioners who claim they can develop such a machine with a couple of researchers in a few months. On the other hand, there are those who claim that such a machine would take a long time to develop because questions on the bar exam require so much common knowledge that a complete ontology of the world would be needed for the machine to pass the bar exam. The machines can't and won't be able to do what human lawyers can do until there is a complete ontology of everything in the world, which will take a long time to get to.

Difficulty of Developing Objective Measures

Panelists also stressed the need for, and the difficulty of, developing objective measures to test and develop certain aspects of AI. The need for standards is self-evident, but what benchmarks to use to develop those standards is a question that has yet to be answered satisfactorily. AI systems, for instance, are better than human lawyers to issue-spot in non-disclosure agreements within contracts, but that is not the only thing that makes a lawyer a good lawyer. AI systems should be thought of as complimentary tools that facilitate lawyers' jobs, not as their competitor.

Does AI Provide Advantages to Transactional Attorneys?

Currently, AI is being used to turn legal departments' playbooks on negotiating certain agreements, which usually reside in the brains of attorneys, into computable forms. So, when there is an incoming contract, the AI system can automatically do the redlining based on the company's playbook. This, then, saves many hours of lawyer time, makes the whole department more efficient, and gives it a more competitive

edge. [Smart] contract drafting, or tech-assisted drafting, is another field where AI is being productively used by way of devising a contract that is most friendly to a lawyer's client, or adversarial to the opposing party. The law firms, companies, and lawyers that use these tools enjoy a competitive advantage over their competitors.

Another panelist claimed that despite talk of leveraging some of the latest advancements in NLP technology to develop an AI system that can take the description of an invention and draft a set of patent claims, it will take some time before it becomes a reality. The lack of human intelligence/common sense is an important impediment to making a useful and practically applicable tool that then can carry out more of those high-functioning lawyerly tasks such as drafting patent applications.

Exciting Areas of AI Growth in Transactional World

One panelist indicated that advances in NLP, large-scale language models, such as Muppet Models or Megatron, where they learn off of patterns from huge corpuses of data, is an exciting growth area.

Another panelist further stated that BERT and GPT NLP models were exciting areas of AI to follow. How machines can start to understand the concepts, not just the key words, is a really exciting prospect.

Computable contracts are another very exciting development that is still in the making. These contracts are automatable and operationalizable. Specifically, the insurance industry, where there are areas and parties with different bargaining power and information asymmetries, is ripe for computable contracts. Even though there is still some time until that capability is realized, once developed, it can really empower consumers to better understand their rights in contracts.

Panelists then answered some questions raised by students. Contrary to the conventional wisdom, paralegals and support/administrative staff have more readily welcomed adoption of AI tools than they were expected to because the tools reduce their workload and more efficiently streamline some of the most tedious aspects of their day-to-day tasks. As to the areas where AI platforms outperform humans, the panelists gave examples of area of prediction, for patent law, where humans cannot compete with machines in digesting the amount of data to come to such a prediction, and search queries, where complex algorithms produce results that validate or invalidate something they would have never found using traditional methods. Relatedly, there needs to be a distinction drawn between 'quality' and 'time' aspects of AI systems. Given an infinite amount of time, a human could create a higher quality contract than today's AI systems, but when time is factored in, AI systems are overwhelmingly quicker than humans

because they can read and look across a great set of data instantaneously.

As a principle, where there's large amounts of data that would be too time consuming or cumbersome for humans to look through, it is safe to say computers can outperform humans. Spotting fraud in millions of credit card transactions is a good example.

To prepare the next generation of lawyers for a changing legal marketplace, where AI is occupying an ever-growing space, law schools need to integrate the tech into all the core subjects, so students can also learn about computable contracts in contracts classes, and they can learn about patents analytics in patent classes. This will prepare students for tools that will become more functional and essential going forward. Schools will also need to grapple with the legal and ethical issues raised by AI and its increasing use in all aspects of life.

Part II: AI and Litigation

The second part of the conference kicked off with examining how AI is gaining a foothold in litigation. Prof. Grossman gave a brief history of Machine Learning (ML) in litigation practice. Pre-2006, a case with 35,000 documents was considered big, today such a case could have upwards of 10 million documents. Technology Assisted Review (TAR) was first coined in 2011 to describe the ML process of separating relevant and non-relevant documents. TAR outperformed attorneys in efficacy and efficiency. Grossman noted the lack of desire among legal professionals to fully adopt AI in litigation for fear of losing control, not finding the "smoking gun" in documents, potential dangers of mis-training AI, or small coding errors leading to disasters. Though adoption has increased, it is not yet at a level commensurate with the increasing benefits it offers.

Prof. Linna drew attention to the larger AI ecosystem, and how it has been growing, evolving, and has come to occupy a greater place in legal research tools along with e-discovery. Moreover, there are improvements in AI drafting tools, whereby these tools draft initial answers to complaints, discovery requests, motions, and briefs. Other litigation tools used in legal aid organizations, bankruptcy, divorce, and landlord-tenant disputes are also gaining traction. Similar tools are used by the administrative state to adjudicate social security disability benefits, or by regulatory agencies, such as the SEC. Adoption might not be at desired levels, but it is making steady gains and increasingly becoming indispensable to all parties involved.

Prof. Surden emphasized that attorneys could gain real competitive advantages by adopting relatively simple technologies such as e-discovery and predictive analytics, and the field might move in that direction.

[Prof. Wexler added that] AI adoption is even lower in criminal cases, where the discovery process involves comparatively smaller numbers of documents than in civil cases. However, criminal cases are also becoming more complex, and AI tools could prove vital to finding pertinent files in 10 hard-drives worth of content, for instance. AI tools might play an even more critical role for the efficiency, accuracy, and fairness of criminal cases, were they to be more widely adopted.

The reason why AI is not as widely adopted in criminal cases, Grossman noted, is the cost. Once the document number goes above 15 to 20 thousand, savings become more significant as the cost is determined based on gigabytes, whose price is based on volume; the higher the volume the cheaper it gets, but if the document volume is small, then the cost associated with AI is greater. Additionally, public defenders and solo attorneys might not have as much of a need or the resources to pay for AI tools in criminal cases.

Researching one million documents does not cost much more than researching 70,000 documents. There are additional marginal costs for hosting etc., perhaps, but what makes it expensive, thus lower adoption rates, is that AI tools have not commoditized enough to bring down costs. Once costs associated with AI tools come down, they will be more broadly adopted by public defenders, small firms, and solo attorneys, thus paving the way for more efficiency, accuracy, and overall improvements in the criminal justice system.

AI's Limitations in Litigation: Where Promises Exceed Reality

The mis/conception of robot judges or robot lawyers is misplaced and far from imminent. The idea of a robot judge in a murder trial is fanciful, for instance, but a robot judge for a traffic ticket or landlord-tenant dispute might not be. Unlike in the autonomous vehicles industry, where billions have been poured into its improvement and a ton of university research has been used, there is no equivalent investment in the legal AI industry. Some of the important questions to ask, Linna noted, are: what is our vision for the future and what these systems could look like? How will we disseminate legal information? What do legal services look like in an emerging digital world?

One of the major problems of AI is its data problem; there isn't enough of good quality court data for instance (partly due to confidentiality issues). Quality data is essential for further development of AI tools, and a lack of it leads to problems in the process as well. For instance, standards for what makes a lawyer's draft of an answer good quality are incredibly important to measure to develop AI tools that can then produce other high-quality works.

With the development of standards, processes, and availability of good quality data, it will become easier to automate and augment tasks and create more beneficial AI tools. This, then, will pave the way for the resolution of harder and more complex problems. A holistic and more proactive approach as to the vision for future AI, quality data, and reliable standards will engender better and more useful outcomes.

The importance of the role university research plays as to the issues of creation the knowledge/data and access to justice warrants repeating because private sector might not be as accessible.

How to Get Good-Quality Representative Data

A lot of legal data that could be used to help society make and improve AI tools are protected by secrecy or confidentiality agreements. This then limits the scope of data to what is publicly available.

Lack of standards, where quality and characteristics of AI tools on the market can be measured, compared, and evaluated, such as Consumer

Reports, coupled with significant misrepresentations of metrics compounds the quality issue. An intermediary entity, like the one producing Consumer Reports, would solve some of those standardization issues but would also require a big legal department for potential suits companies would bring against it. Universities, as highly regarded sources of knowledge creation, could play this trusted intermediary role.

Currently there is distorted access to data and distorted incentives to develop AI technology because it is based on funders' interests. Putting it differently, the AI tool developed will favor the party paying for it and disfavor the party it is being employed against. One way of preventing inaccurate and unfair outcomes is to allow universities to specifically design these technologies to serve interests of innocence and fairness, instead of directing them to prove guilt for example. Universities are better equipped to serve the greater good.

The police's criminal data is distorted because it reflects the bias, prejudice, or inaccuracies of the police officers who input it into the system. For instance, police patrolling certain neighborhoods more than others leads to higher arrest rates, higher correlations of crimes and more data perpetuating support for police patrol in those neighborhoods. So, police practices themselves could and do lead to distortion in data and the models built upon that data.

Government's increasing use of AI, especially in the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) Recidivism Algorithm, raises important issues. The third parties supplying these tools to the government—tools on which judges rely to decide whether to jail someone or not—do not share the underlying data to be tested and validated for possible biases and/or distortions. Importantly, the government lacks the expertise to evaluate these tools that affect people's lives in profound ways.

A[n AI] tool used to determine whether to imprison someone must be based on good quality data and must be free from bias and distortions.

AI in the Forensic Evidence Realm

Forensic evidence is the realm in which AI makes its most prominent appearance in the criminal legal system. It is widely used in bullet/gunshot matching, complex DNA analyses at crime scenes, fingerprint analysis, voice matching, iris matching, handwriting authentication, authorship authentication for tweets, etc. The issue with these usages is that these AI tools have not been scientifically validated. There are issues of access to scrutinize and evaluate these use cases as discovery requests and subpoenas to do so are blocked by intellectual property or privilege claims. While expert testimony is subject to cross-examination, thus ensuring fairness, accountability, and due process, AI tools used to evaluate or determine evidence are not held to the same scrutiny or standards.

Scripts for defense counsel to cross-examine experts of AI tools on the stand would be an important development, which Academia could undertake more so than any other party. These cross-examination scripts could lay bare whether any subjective [design] choices go into training those tools and how to resolve them.

Grossman added that AI evidence tools are useful; however, the developers should not be allowed to deny requests to disclose the algorithms used or the data it was trained on if they are being employed for sentencing. She also mentioned that two kinds of experts are needed on the witness stand: the person who pushes the button and gets the result, and more critically, the person who trains the AI.

Surden stated that the companies selling AI tools to the government should be required to share the source code and algorithms employed in them, especially in a case where life or liberty is at stake.

Linna gave the example of Robert Williams, who was wrongfully arrested because of a faulty facial recognition match, and stressed that though biases exist, some of the more serious problems could be eliminated with more robust validation and more representative data.

Wexler recommended two concrete changes. First, trade secret evidentiary privilege should not apply in criminal cases, as it is a non-legitimate exercise of withholding relevant evidence from the defense. Second, privacy statutes should not block criminal defendants from relevant evidence. Though not strictly tech specific, these would improve cross-examination and adversarial scrutiny processes.

Predictive Analytics

Predictive Analytics is the use of data to make legal predictions about outcomes of cases, damage awards, etc. Clients of corporate legal departments drive the push to have lawyers use data to determine how much something will cost. There would be AI tools that could reliably predict whether a judge would rule one way or another by now, for instance, according to predictions, which has not materialized. The reality is Predictive Analytics is not where it was forecasted to be. It would be more accurate to label what exists as descriptive statistics rather than predictive analytics.

While agreeing that predictive analytics is not as useful of a tool at the evidence stage, Wexler drew attention to their widespread use in pre-trial criminal prosecutions and post-conviction sentencing stages. There is a need to scrutinize and evaluate these popular tools; however, the lack of procedural rights in those pre-trial and post-conviction settings is impeding the ability to contest those tools.

Q&A

Can AI be used to reduce bias in a courtroom?

Referring to a study of parole boards in Israel, where hungry judges were very unlikely to grant petitions before lunch, and another study showing some American judges approved 40% of asylum cases, while others approved close to zero, Linna stated subjective preferences/biases of judges could be minimized if better quality research, data, and algorithms are employed.

Grossman added, though there is potential for AI to reduce biases, a lot more research needs to be done to determine the effect of these tools on judges in the first place and how numbers and predictions generated, say on recidivism, push in one direction or another. This is especially true if the judge in question is up for re-election for instance.

There is a lot of potential to develop tools that could demonstrate biases, such as a tool that predicts when police witnesses are likely to perjure themselves, but development has not happened because of market failure. Academia could play a critical role here, too. Data offers the opportunity to identify biases, whereas peering into judges' mind to see bias is not possible or as easily discernable.

Going Forward

Grants by governmental bodies play an essential role in making possible the research and tools necessary to develop systems and machines that are accountable under the law. Universities need to get better at receiving and applying for those grants. The funding available from private companies or the Department of Defense, for instance, necessarily caters to the needs and interests of researchers, which makes the role universities need to fill that much more important.

Other countries are also showing interest in and development of legal AI systems, which adds an international aspect to all our discussions. This, in turn, gives an urgency to the idea of creating tools that are truly based on rule of law, grounded in principles of fairness, and objectiveness free from biases and distortions.

Part III: AI and Law in the Near Future

The third part of the conference kicked off with Prof. Surden recommending skepticism towards future predictions of AI that are more than five years out and are predicated on technologies that have yet to be developed. By contrast, predictions that are grounded in existing science and are within a one-to-five-year time frame, though still fraught with some inevitable uncertainties, can be more reliably made, given that certain technologies are likely to get faster or cheaper, etc.

Most Promising and Likely Developments

For Meghan Ma the revelation of what information is contained within legal texts is an exciting prospect. People usually have assumptions or intuitions for what the law is, mostly based on experience of it, but now the content of the language of the law can be tested on both computational/rules-based side of law and data-driven side looking into neural networks. That allows for unpacking how law has been captured and manifested in natural language.

Another panelist, Prof. Ashley, identified two areas of law impregnated to promising developments: automated contract review for due diligence purposes and legal text analytics improving access to legal materials. The former will likely happen within the next two years through amplifying the utility of virtual data rooms, thus enabling lawyers to better understand what is in the enormous data repositories of contracts associated with due diligence searches. Lawyers' notes and annotations when they review contract repositories are not used in development of AI tools currently. If or when they are, it will improve the process greatly because it will allow certain more relevant information to be privileged over others. Allowing public users to find sentences that explain statutory terms in cases is one way the latter is taking place.

Halliburton acknowledged document analytics as the area of great potential for the foreseeable future. Document analytics is the ability of machines to extract structured data out of unstructured documents. When humans look at a contract, they see that the lines, bullet points, titles on the document give it structure and clarity and demonstrate different kinds of information. Understanding the structure allows for a lot more information. However, machines cannot see the same structure in documents that humans do. So, getting information from documents and converting them into structured data [such as maps, graphs, prices, dates, etc.] will usher in a new phase in AI.

Limited use of AI tools for social justice issues that have significant impacts on society was a recurring theme that Prof. Surden reiterated, while noting the great potential these tools offer for social justice values.

In-house departments of large corporations and big law firms have the resources to invest in the development and betterment of AI tools; however, those have relatively less social impact.

Incremental Improvements

Within predictive analytics, predictive filing is a process wherein a person can save a couple seconds each time they use it to file an email. This is an incremental improvement, but the totality of those seconds frees up some time that can then be used for more value-added work.

Such incremental improvements lead the way for more substantial improvements. For legal research, for instance, one could learn all the right methods/tools to perform searches that engender useful and accurate information, but it is a steep learning curve that consumes a lot of energy, time, and money. Imagine being able to do very accurate searches by just using regular language, a seamless transition without the need to learn new tools and methods. That is what would be a very substantial change that will be made possible through incremental improvements.

Incremental improvements lead to important changes because of positive feedback loops, where the machines are trained as they function by inputting the new learnings along the way. For example, the more you train your email program, the better it gets at detecting spam. This process makes such systems more efficient and more useful.

Meghan Ma stressed the necessity to marry the two approaches (data driven vs knowledge representation) where one is focused on the structure and framework of information [how to get better data out of (un)structured documents?], while the other is more focused on the context or nuance of the language itself [how do we guarantee not missing important information/aspects?].

Where will AI Add the Most Value?

Automated Contract Review is increasingly being adopted in corporate legal departments, which puts in-house counsels into the position of monitoring the creation of new contracts and understanding existing contractual obligations. This will be the most valuable use of AI in the next five years in the corporate legal AI realm. This is a realistic prediction because incremental developments, and increasing amounts of data available will render it possible.

Likening current AI adoption in the legal world to the lowest, yet most foundational, level of Maslow's hierarchy of needs, Halliburton pointed to the great potential AI will offer once it is more broadly adopted and the basics are fully built. This foundation will lead to more interesting analytical work, like having enough food to survive is to enjoying Shakespeare.

Computable contracts is another area impregnated to important developments because it works well with contract analytics, but also it is possible that contract analytics provide information to then build better computable contracts.

In sum, the next five years will see improvements in understanding documents (formal v. informal, structured v. unstructured, implicit v. explicit) and making them more structured/useful. This infrastructure improvement, in turn, will lead to more powerful analytics down the road.

Limitations and Roadblocks

The problem of incorporating legal knowledge into machine learning processes still persists. Though there have been big advances in representing legal texts with transformer models, such as legal BERT and GPT-3, there still are two challenging roadblocks ahead. First, deep learning has information about a ton of complex correlations but no structure, which makes it difficult to say with certainty that the models will be able to explain and argue legal conclusions without taking conceptual structures into account. Second, the context leading to drafting documents, things left unwritten, and shared human experiences that lead to writing in the first place are all missing from these tools, which then renders them incomplete/defective.

Another limitation, and a recurring theme throughout the conference, is the lack of good quality and representative data to train the models. For instance, all AI models scrape data from contract repositories of large public companies (SEC's EDGAR database). But that data is unrepresentative of a mom-and-pop business, so its utility is diminished and only could cater to a specific clientele. Quality of data and access to a high quantity of good quality data are essential for meaningful improvements and development of AI tools. The secrecy or privacy issues surrounding legal documents continues to function as yet another roadblock.

If I Had a Magic Wand

Halliburton would build a contracts corpus where all kinds of contracts (from companies of all sizes) could be deposited to create AI tools that are more powerful and could be useful to many not just a few. Prof. Surden added that in areas where there are large data repositories outside of law, AI has already made meaningful progress. That needs to happen in law, too.

[Meghan] Ma would make legal texts (case law, statutes, contracts) and legal information in general more accessible. Justice is deeply intertwined with information, and without understanding/accessing that information, justice cannot be served properly. The barriers between the legal information regular people need and the difficulty of obtaining that information need to be demolished.

Ashley would waive his wand to enable deep learning neural networks to explain their predictions in terms lawyers can understand and make use of it. For this to happen, legal text analytics will have to extract aspects of legal knowledge from the legal texts and be able to turn them into explanations.

Prof. Surden, while acknowledging the difficulty of actually carrying it out, raised the idea of making court documents more structured, such as containing definitional terms, which would increase the utility of those documents for AI purposes. He also mentioned how the U.S. House of Representatives, a decade ago, decided to revise and structure titles of U.S. code. Though unlikely, if courts could emulate that step and produce structured documents, it would pave the way for more beneficial AI developments.

Most legislation and contracts are already structured. Case law is where the problem lies because judges write in narrative styles. If there are patterns in these narrative styles, Ma added, structures can be found, though that is unlikely to happen in the next five years.

Q&A

Are we close to developing software that could reliably predict outcomes of court cases based on available data so that a lawyer could determine whether a case is worth litigating?

Lex Machina, on which Surden and Halliburton worked during its origination at Stanford, does useful things like predicting how long a case is likely to last once the judge is known, or where to file for a more favorable outcome that translate into smaller bills. However, this is still relatively surface level [or very contained areas of] work that lacks nuances. There is also Blue J Legal, a project that is narrowly focused on tax law employing predictive analytics. So, there is a great area of improvement/expansion for these tools to truly dive into what a case is really about substantively and to reason on the basis of facts and arguments.

What area is AI most successful in now, and where will its biggest success be in five years?

For Halliburton, e-discovery and due diligence review are two of the most successful areas of AI use in law.

In five years' time, Ashley predicted, technologies that improve access to justice, such as rendering legal texts more understandable and accessible, will be successful.

Legislative Texts & Structure?

Reading long legislative texts of thousands of pages is a herculean task. Ashley, while admitting to not having a lot of experience with legislative

texts, said that the prospect of developing AI tools able to summarize these long texts was slim. Ma, who is more heavily involved with linguistics and texts, stated that though there are efforts to translate heavy and long legislative texts, such as Roses Code and Catala, currently AI in the legislation realm is far from being deployable, largely due to the opacity of language.

More structure in legislative texts would make it easier to translate, thus more useful for AI tools. However, agreeing to legislative text is already a difficult and time-consuming task; to make it even more structured will exacerbate those aspects. Additionally, caution is needed to not overdo it because too much structure for legislative texts could actually yield negative results; the text still needs to reflect flexibility and elasticity that is present in natural language. Otherwise, it could lead to other issues and endanger compromises needed to pass legislation.

Halliburton suggested the idea of using Simplified Technical English, a controlled natural language where each word has one meaning, in legislative drafting and contracts to make them more structured and translatable and avoid ambiguity entirely.

Responding to Surden's suggestion of developing new conventions in legislative text, Ashley underlined the importance of developing such conventions by pointing to the lack of conventions on even such simple things as subsections and subparagraphs. Standardizing even those simple conventions, Ashley added, would be very helpful.

Many of the technologies talked about already exist; the issue is with adoption, especially in the law firms. Do you see that changing?

Motivations and working practices vary greatly within different areas of the legal sector. In-house departments are seen as cost centers because they don't make money, so anything they can do to cut costs is an incentive for them. Law firms, on the other hand, focus on billable hours because that is what makes money for the company, which creates its own kind of power structures and incentives. This partly explains the disparity in adoption rates.

Education, at all levels, will play a critical role in changing the attitude to adopting these tools. Adoption will be easier if students, starting from young ages, are exposed to these technologies and/or are interacting with technologists more frequently. Once that is embedded in our education system, there won't be as much resistance to adoption.

Participants

Jason Adaska
Holland & Hart LLP

Kevin Ashley
University of Pittsburgh School of Law

Stephanie Curcio
Legalicity (NLPatent)

Maura R. Grossman
University of Waterloo

Ansel Halliburton
Align Technology, Inc.

Daniel W. Linna Jr.
Northwestern Pritzker School of Law and McCormick School of Engineering

Megan Ma
CodeX - the Stanford Center for Legal Informatics, Stanford Law School

Harry Surden (Moderator)
University of Colorado Law School

Roland Vogl
CodeX - the Stanford Center for Legal Informatics, Stanford Law School

Rebecca Wexler
UC Berkley Law

About Silicon Flatirons Center

What We Do

Silicon Flatirons' mission is to elevate the debate surrounding technology policy issues; support and enable entrepreneurship in the technology community; and inspire, prepare, and place students in these important areas.

Founded in 1999 at the University of Colorado Law School by Phil Weiser, Silicon Flatirons is a recognized leader in interdisciplinary events and programs. We serve students, entrepreneurs, policymakers, and professionals, and support the joint missions of Colorado Law on teaching, scholarship, and public service.

Although technology and innovation have evolved radically over two decades, our purpose remains the same: to convene multi-stakeholder discussions, support innovation, and develop the next generation of technology lawyers, policy experts, and entrepreneurs.

Our initiatives are hubs for pivotal issue areas at the intersection of technology and law, led by internationally recognized experts who facilitate programming and convenings and generate thought-provoking scholarship within those areas.

Learn more at siliconflatirons.org/about-us/.

Our Team

For more information about center leadership, faculty, staff, fellows, and advisory board, visit siliconflatirons.org/about-us/our-team/.

Our Supporters

Silicon Flatirons exists thanks to the generosity of our supporters and the strength of our community. We rely on their contributions to advance our mission to catalyze policymaking and innovation and to develop the next generation of tech lawyers, policy experts, and entrepreneurs. For more information on current official Silicon Flatirons Supporters, visit siliconflatirons.org/about-us/supporters/.

Publications

We promote thought leadership and intellectually honest discourse not only in our events, but in publications from our team, our roundtables, and scholars presenting at our conferences. See more at siliconflatirons.org/publications/.