dri™
Center for Law and
Public Policy

Artificial Intelligence and Legal Practice

Management and Ethics

This publication and the works of its authors contained herein is for general information only and is not intended to provide and should not be relied upon for legal advice in any particular circumstance or fact situation or as a substitute for individual legal research. Neither DRI nor the authors make any warranties or representations of any kind about the contents of this publication, the accuracy or timeliness of its contents, or the information or explanations given. DRI and the authors disclaim any responsibility arising out of or in connection with any person's reliance on this publication or on the information contained within it and for any omissions or inaccuracies. The reader is advised to consult with an attorney to address any particular circumstance or fact situation.

DRI
222 South Riverside Plaza, Suite 1870
Chicago, Illinois 60606
dri.org
© 2020 by DRI
All rights reserved. Published 2020.
Produced in the United States of America

This copyrighted product is provided free to the public for general and lawful use, with attribution, as a public service of DRI—The Voice of the Defense. Sale of this product or any part of it is prohibited.

Artificial Intelligence and Legal Practice

Management and Ethics

DRI Center for Law and Public Policy
Artificial Intelligence Working Group
July, 2020

dri Center for Law and Public Policy™

DRI Center for Law and Public Policy

Kathleen M. Guilfoyle, Chair, DRI Center for Law and Public Policy, *Campbell, Conroy, & O'Neil*

Issues and Advocacy Committee

Jill Cranston Rice, Chair, Issues and Advocacy Committee, *Dinsmore & Shohl*

Artificial Intelligence Working Group

Edward P. Purdue, Chair, *Perdue Law Group LLC*

Julia Brickell, *H5 Technologies*

Renee Dopplick

John H. Dollarhide, *Butler Snow*

Kevin T. McCarthy, *Larson King LLP*

Marta-Ann Schnabel, *O'Bryon & Schnabel*

Sharon D. Stuart, *Christian & Small LLP*

Introduction

We are experiencing a time of unprecedented change. The World Economic Forum has dubbed this era the Fourth Industrial Revolution and predicts faster and more widespread change than in any of the first three. The changes—technological changes enabled by quantum computing and exponential growth in data storage capacity—are predicted to affect virtually every aspect of our lives. AI systems or products (often dubbed "intelligent" or "autonomous"), are the topic of much fascination, human prediction and hyperbole. These processes, built through an integration of human-crafted algorithms trained on amassed data sets, with extensive feedback, are intended to improve efficiency in a range of tasks through use of statistical models. These are intended to replicate the results of human decision-making (whether the task is to "find more of the same," brake a vehicle, report about the cause of a patient's physical symptoms, or make countless decisions, big and small, that humans make on a daily basis)." AI systems are the subject of immense discussion and are predicted to speed work previously done by humans (with an effect on jobs), save time, and enable new insights or actions based on patterns in existing data. They are predicted (eventually) to permeate how our refrigerators order food, televisions select viewing content, medical providers determine and deliver our health care, cars (or someone else's) transport us to work. They may affect how packages are delivered, elderly are cared for, whom we hire and fire, meet, and marry. They may affect how people and companies are regulated, policed, and judged, and how we lawyers advise our business, litigation, family law, and criminal law clients. AI systems already affect the topics on which we need to be competent, with which we need to stay current, and for which we need to engage those with the requisite expertise.

Lawyers encounter artificial intelligence in multiple roles. Both as in-house and outside counsel, they advise clients who are developing (or buying businesses that develop) AI systems or are deploying AI systems in their business operations. As litigators, we may use AI tools to research the law, predict case outcomes or preferred judges, choose legal teams, assess opposing counsel, meet eDiscovery obligations, or draft court papers. We may look to AI systems for insight into jury selection or to help clients defensibly dispose of records. As deal lawyers, we may have such tools help draft or review contracts and assess deals. In managing law practices, we may choose such tools to help predict or meet budgets, better understand market competition, assist with hiring decisions, or benchmark lawyer performance.

The implication of algorithmic contributions is important to lawyers and to the rule of law. As we (and the clients we advise) choose to use (or not use) processes enabled by

artificial intelligence, the expectations set by the model rules of professional responsibility demand that lawyers analyze these new approaches at several levels. The basic duty of competence mandates that lawyers stay abreast of the risks and benefits of technology. (ABA Model Rule 1.1, comment 8.) Therefore, we must acquire a general understanding of AI systems and associated issues. Beyond that, clear communication with clients about the risks of its use (or non-use) requires a solid understanding of what truly can be accomplished with a particular AI systems and the competencies required of the user to achieve that result. It also requires knowledge of potential biases and potential misuses. Competent advice to clients about marketing claims, potential liability, privacy, contracts, due diligence, and intellectual property requires the same. To support a determination that a lawyer's fees are reasonable requires some assessment of the relative merit and cost of alternatives that may be presented by AI systems. Candor to the tribunal and fairness to the opponent requires knowing what has actually been accomplished by AI systems that gave rise to litigation claims or were used in case preparation. And avoidance of discrimination—by the lawyer or the client—requires understanding the biases imbedded in the algorithms and the data sets on which such products are trained, the processes in which they were deployed, and the purposes for which they were truly designed.

It is imperative that lawyers consider the potential impact of AI on our legal systems and social justice as well. The biases embedded in social data that is used to train machine learning and other algorithms, the biases of the trainers, and weighting and rankings in the algorithms themselves have the potential to embed bias in a system that should be blind to color, economics, gender, and more. As AI systems make their way into our legal system for everything from predictive policing to case distribution to sentencing, the affect of the choices made in training and design may systemically tarnish the luster of the rule of law. As these processes are considered for streamlining management of legal practices, lawyers will play a critical role in identifying and ameliorating the potential detrimental effects. Lawyers can only provide this protection by becoming knowledgeable about the standards by which AI systems should be evaluated and demanding disclosure of the resulting information so its trustworthiness can be determined.

Most importantly, as lawyers, we should avoid jumping to the conclusion that AI tools are magic, foolproof, fair, intelligent, thoughtful or even function reasonably. Our ethical obligations are unchanged. We need to understand the tools we are using, the quality of the results we are producing, and the accuracy of the representations that we make about the results. We must calibrate our trust based on actual efficacy, for which we need accurate information not marketing enthusiasm.

Table of Contents

Current Legal AI Technologies

This section provides information on certain applications of AI systems and products specific to the legal industry that are currently available. This section does not address speculative or concept-level applications of AI to legal tasks. Additionally, this section does not address the external subject-matter applications of AI that affect the legal industry (*e.g.*, self-driving car litigation or mergers and acquisitions of companies that produce AI systems).

Before discussing what these platforms can do, it is important to understand how they work. Generally, these new legal technologies make use of machine learning. Machine learning is a decades-old algorithmic approach that enlists one or more algorithms in a process of pattern recognition. The algorithms are exposed to data sets and trained by humans on data that is similar to the data the algorithm is being trained to tag. In the case of the variety of machine learning most commonly used for legal applications, "supervised machine learning," after significant human input to identify for the algorithm the data that is meaningful for the purpose, the software amasses a library of information against which future data sets are matched. Most of the machine learning we will discuss makes use of text processing. In short, the software is fed the text of a document and analyzes it based on the algorithm and amassed data library. How the algorithms work is, being comprised of at least two entire scientific disciplines (computer science and linguistics), beyond the scope of this paper. But it suffices for our purpose to say simply that the algorithms are designed, and the software programmed, to provide users with certain information contained in documents, and—to the extent it matches the amassed training, provide it more quickly than a user could find it alone.

Document Review/E-Discovery

Modern document review was one of the first widespread applications of AI to a task traditionally performed by a lawyer. The information age has committed much of the facts of our lives to recorded form, and most of that is now created and stored electronically. Litigators have ethical and rules-based obligations to identify, review and analyze relevant documents. The boon in the amount of information and communications stored electronically, has increased the burden on attorneys to simply find the documents that will allow them to ascertain the important facts of a litigated matter. Coupled with increasing pressure from large organizations to reduce the amount of money spent in litigation, the seemingly inefficient task of having a human read every page of every potentially relevant document became fertile ground for the implementation of AI-enhanced processes.

Much document review is now done through predictive coding, which may be delivered as a service or as a user-implemented software. One type of predictive coding involves machine learning in which the user "trains" the software to distinguish relevant documents from non-relevant ones (or potentially privileged documents from documents not potentially privileged or documents belonging (or not) to other pertinent categories (*e.g.*, confidential). While, in terms of particulars, there is great variety in the machine learning technologies and workflows currently in use, there are, at a general level, key steps that all have in common. The process starts with the identification of an initial training sample of documents. That sample is manually reviewed by an attorney who is familiar with the matter and the documents in the sample are identified as relevant or non-relevant, etc., according to the categories of interest. From the coding applied to the initial sample, the machine learning algorithm creates a statistical model of the characteristics of documents marked relevant and non-relevant (or in, or not in, another category). The algorithm uses that model to code not-yet-reviewed documents as likely relevant or likely not relevant (with the degree of likelihood indicated by a probability score).

Some predictive coding systems can be instructed that particular words are always to be coded a particular way. For example, some systems can be instructed that if a document includes "mylawyer@mylawfirm.com," all documents containing "mylawyer@mylawfirm.com" would be tagged as privileged. To be sure, this is an oversimplification. In practice, the algorithm performs thousands of analyses on multiple levels (not only whether a word appears in the document but also where it appears in the document and its proximity to other words, and so on).

A reviewer (again, in most cases, an attorney who is familiar with the matter, usually the same attorney who coded the initial training set) then reviews some subset of the documents that the algorithm has newly coded and indicates whether he or she agrees with the system's coding. The results of that manual review are then fed back to the algorithm so that the statistical model for distinguishing relevant from non-relevant documents can be improved. This iterative process of coding, review, and adjustment of the algorithm continues until quality control measures show that the system is achieving the targets set.

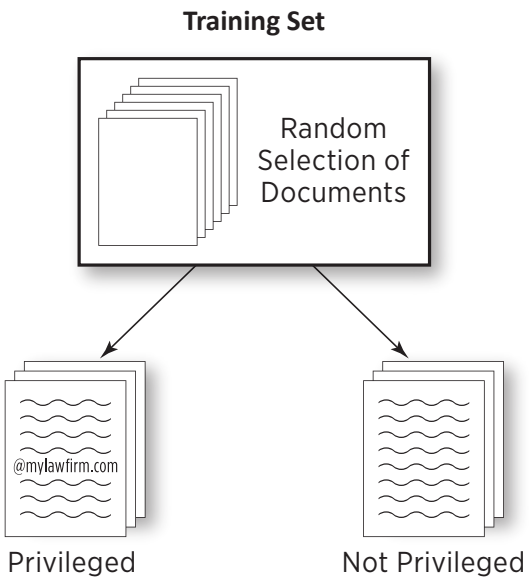

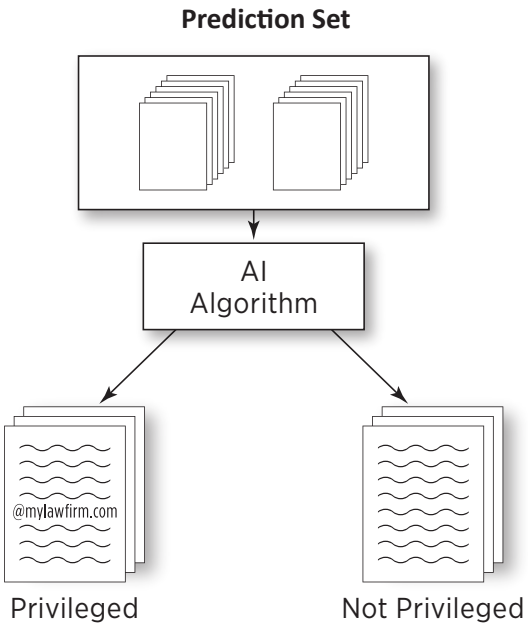

Other forms of predictive coding have different strengths. Some use algorithms to cluster; some, properly constructed, identify relevant content. Others target potential privilege.

Technology-assisted review, including predictive coding, has gained acceptance in many courts, beginning with *Da Silva Moore v. Publicis Groupe*, 868 F. Supp. 2d 137 (S.D.N.Y. 2012). In 2015, in *Rio Tinto PLC v. Vale, S.A.*, 2015 WL 872294 (S.D.N.Y. Mar. 2, 2015), the court opined that it is black letter law that a producing party could use technology-assisted review. Currently, most disputes about technology-assisted review are not whether it should be used, but how transparent the parties should be with each other on the approach and protocols that will be used in execution (ranging from sample set parameters to the extent to which different processes (like key word search and technology-assisted review) could be combined in a single review process to measures of precision and recall of the results. Expect more widespread adoption of technology-assisted review as technological fluency among practitioners increases, data volumes grow, and cost pressure from clients overcome other factors and continues the drive toward efficiency.

Legal Research

Another widely used AI application in the legal industry is in legal research. As with document review, AI legal research tools employ machine learning and other approaches to assist the user. Unlike traditional research platforms that relied on user-created Boolean searches, these new platforms allow users to make natural language searches. In other words, the user can describe the relevant legal issue, relevant facts, and even procedural posture as if he or she was discussing the research topic with a colleague. The artificial

intelligence built into the platform will then rank, weight, and score the inputs and run comparisons to its database cases, statutes, and regulations, and provide results that are specific to the relevant information. As these systems have continued to grow through machine learning, they have become even more discerning in their searches. The new research platforms allow for searches that are more narrowly tailored to specific legal issues as the systems are able to refine searches based on legal issue, jurisdiction, procedural posture, etc. For instance, one of the platforms allows the user to insert a question as specific as, "What is the distinction between independent contractors and employees under the New York 'economic reality test' since 2016?"[1] The user can then make the search even more specific by describing the procedural posture (motion for summary judgment), and some relevant facts for the motion (driver for delivery service).[2]

One of the most significant functions of these new platforms is the improved "Shepardizing" technology. Every practitioner is familiar with the need to Shepardize legal precedent cited in briefs. AI technologies have taken the ability to Shepardize a step further. Specifically, these platforms quickly inform practitioners when a case has been implicitly overruled. This situation arises when a court's holding was based on a prior case and that prior case was overruled by a separate decision. It is best described by the below graphic[3]:

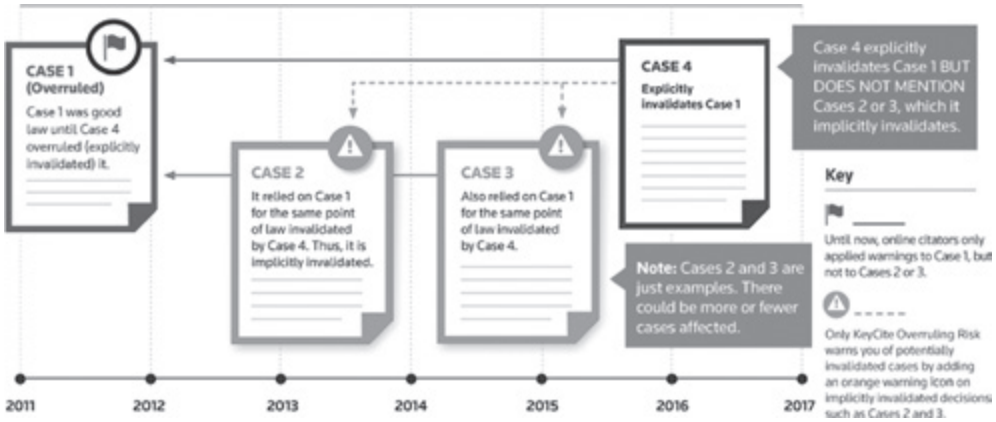

These new systems save practitioners substantial time in reviewing cases to ensure a holding remains good law. They also help to flag for practitioners the described situation where a case has been implicitly overturned because relied-upon precedent has been overruled.

Brief Writing Assistance

AI based research platforms take traditional research a step further by reviewing draft briefs or other pleadings and providing suggested legal precedent for inclusion. These sys-

[1] "ROSS Intelligence Offers A New Take On Legal Research," Above the Law, https://abovethelaw.com/2019/05/ross-intelligence-offers-a-new-take-on-legal-research/?rf=1 (accessed on September 2, 2019).

[2] *Id.*

[3] "Westlaw Edge: KeyCite Overruling Risk, Thomson Reuters, https://legal.thomsonreuters.com/en/products/westlaw/edge/keycite-overruling-risk (accessed September 2, 2019).

tems are easy to use. The practitioner simply uploads the document to the chosen platform and the AI facilitates the research verification. Each platform has differences, but they provide the same general functions. The platforms will identify each cited case; they will flag those cited sources that have been overruled, questioned, or criticized and provide links to negative authority. The platforms provide a clear list of the cases that have been cited negatively. The practitioner can then click on each case and the platform will provide a link to every case that has referenced it negatively.

The platforms also have the capability of identifying legal and factual issues discussed in in a brief and of suggesting authorities that may be relevant to those issues. The platform will also provide links to the suggested cases or other legal authorities. The practitioner can then follow a link to view the suggested authority and determine if it should be included. In the same vein, the platforms provide cases with similar language to the uploaded brief. These functions provide a different approach to research than user-input Boolean searches, which like any search terms depend on the skill of the user and return hits, but not cases that are ranked as "similar" in some way. The platforms also provide a double check of the practitioner's work, potentially identifying certain cases or other authority that is worth considering for completeness.

Finally, the platforms assist practitioners in recognizing weaknesses in an opponent's brief. Once the opponent's brief is uploaded, the platform will engage in the same review of the cited legal authority based on the text in the brief. It will highlight any cases that have overruled, questioned, or criticized the opponent's cited authority. The platforms also return cases that are relevant to the legal and factual issues in the opposing brief which were not cited by the opponent.

Contract Analysis

AI has been implemented in assisting attorneys—outside the litigation context—with various forms of contract analysis. Much of these focus on making the infamously tedious task of due diligence more efficient. Using the text-assessment capabilities of a type of machine learning called natural language processing, AI software can assist with term standardization across contracts, identify redundant provisions or inconsistent language, and extract key data from a contract (like party names, addresses, dates, and term length) and then create reports of that data from many contracts for human review, in less time than a human can.

AI systems can also be used to assist with triage in contract drafting and review projects, especially those involving multiple iterations and revisions. The information tracking functions in these tools can be used to inform the user about changes to a contract and across documents. They can be instructed to examine predetermined portions of contracts to identify if requested prerequisites are present (including opposite party signatures) and will then alert the user that a contract is ready to be signed.

Notably, these tools are primarily useful in high-volume situations and where a certain contract or provision is standardized and used in many different situations. In situations

where there are a fewer number of contracts to be reviewed and where there is much variability among the language, these tools provide less of an efficiency gain.

Outcome Prediction

Machine learning has also been implemented to make predictions about case outcomes. Like with predictive coding, the AI software can be engineered to take a set of given facts, identify patterns or weigh and rank information in a training set, and then output a predicted result for future instances. Instead of a training set of documents and a user manually coding those documents to train the system, outcome prediction software allows the user to input a set of parameters, including outcome, from past known cases, and instruct the software to assess patterns or weigh and rank information about future cases and provide ratings of likely outcomes based on the comparison, somewhat like an actuarial table. For example, in the personal injury context, the software can apply an assessment of key data from past cases on which it is trained, and to a newly given data set (injured person's age, salary, medical expenses, etc.) and make a "prediction" the amount a jury might award.

Some have applied natural language processing—text-based machine learning—to predict how judges will rule on issues put before them. In 2012, researchers published a study where they used an algorithm that assessed general case characteristics to predict how the U.S. Supreme Court would rule on the 628 cases before it in 2002, compared to a team of legal experts who likewise made predictions. The algorithm was correct in 75 percent of the cases, compared with the legal experts' success rate of 59 percent. Of course, the inputs available for modeling US Supreme Court decisions are quite different from the inputs that might be available for modeling decisions in a lower court proceeding (*e.g.*, in a proceeding resolving a landlord-tenant dispute) so care must be taken to ensure that evaluations of the effectiveness of the technology are in fact reflective of real-world applications.

One area where this technology is employed is automated dispute resolution. These are online, AI-based tools where the users input some amount of data, including settlement values or ranges, and the automated system generates an outcome. A form of alternative dispute resolution, this process can achieve a result in minutes that might take a skilled mediator several days to craft through traditional techniques. One simple example is double blind bidding, where parties input a settlement range and, if there is overlap, the system automatically splits the difference of the overlap. One party might input a settlement range of $100,000–$150,000. The other might input a range of $50,000–$110,000. The system would automatically generate an award of $105,000. The complexity of the data that input into the system and the result it generates varies greatly.

Intellectual Property

A final area of law currently being affected by AI is intellectual property. AI tools have been developed to speed up patent searches for "prior art" as well as trademark searches. These tools are essentially search engines specialized for finding the distinctions and similari-

ties in patents and trademarks. These tools have similarly been used to maintain large IP portfolios through automated monitoring of patent applications and alerting the user when there are patent or trademark applications that contain a certain level of visual or linguistic similarity with the user's IP.

AI Considerations for Civil Defense Legal Practice Management

This section looks at the key benefits, considerations, and challenges for civil defense law firms and corporate legal departments looking at incorporating artificial intelligence solutions into their practice management. The rapid emergence of AI-driven legal practice management solutions supports new ways of doing business with the potential for creating new growth and profitability. The implications for cost management and competitive legal services are wide-ranging, and there is no single model for building a successful AI-powered law firm. To make AI work for your law firm or legal department, delivering value for your core business strategy and operating model is essential.

How Can Artificial Intelligence Benefit My Legal Practice Operations?

A significant force for improving profit margins and cross-practice coordination will be how to cost-effectively deliver AI value for operations. The operational gains from AI within a law firm or legal department are most likely to focus on seven areas: automation, compliance, case management, risk management, digital marketing, efficiency, and productivity.

Automation

The days of handling tedious and duplicative tasks are dwindling, at least for those law firms looking to maintain a competitive advantage. Automating routine tasks, particularly manual ones, can help law firms reduce redundancy and mundane case-related tasks while building rapid-response capabilities. Don't overlook potential opportunities to capitalize on automation gains. Is every case or brief a unique snowflake? If not, AI might be the secret sauce for replacing income currently driven by repetition to income driven by volume.

Compliance

Compliance is among the most time-consuming, people-intensive operational tasks of a legal practice. The power of AI, particularly machine learning, can automate monitoring and reporting requirements, thereby reducing the burden on staff. Although AI is advancing, AI solutions in this area should be paired with human review. AI-assisted compliance solutions can help staff more readily detect and assess conflicts, potential ethical issues, supervisory oversight performance, conformance with litigation holds, upholding of firewalls, human resource issues, and risk management performance.

Case Management

Delivering on effective case management ultimately builds the stronger relationship between your legal practice and the client. Managing cases and clients can be complex, leading to high-volumes of information, data, and interactions and not enough time to document, process, and archive. AI-driven case management solutions can be helpful in automating identification, classification, and clustering of documents and other assets, such as multimedia. Attorneys can search, analyze, and manage large amounts of information in a fraction of the time and at a scale not previously possible for small to large teams of attorneys. Transcription and translation solutions can automatically transcribe and translate documents and audio, making the content searchable and easy for attorneys to explore. AI-powered facial recognition and photo interpretation software can help identify images and videos containing specific people, objects, colors, and features.

Risk Management

AI systems can help law firms manage risks, including as related to asset management, client deliverables, compliance, and business changes. Next generation solutions can help you forecast case outcomes, cash flow, and demand for legal services to help bridge the gap for decision-making of whether to accept a case and, if so, whether it might be strategically beneficial to mediate, arbitrate, or litigate. The software can apply a wide range of input variables, such as case complexity and attorney resource allocations, to a potential case to compare and report where it might fit on the risk spectrum, legally and logistically. Such information can help decision-making on whether to pursue a model of high-volume, low-risk versus high-risk, low-volume workload, and help manage a blended approach. Predictive evaluations of case outcomes could provide another view into risks, forum selection options, attorney assignment, and sequencing of related cases.

Productivity

The core backbone of the civil defense legal business is maximizing use of attorneys' time and expertise to deliver quality legal services responsive to clients' needs. AI solutions can liberate attorneys to pursue higher-value work of client representation, such as legal reasoning and strategic analysis, while spending less time in the lower-level legal research and logistical weeds. This shift to more revenue-generating time will enhance your bottom line. Productivity gains are not limited to attorneys. AI has the potential to promote continuous operational improvements throughout your legal practice, enabling more effective use of employee time and talent. These gains can potentially reduce cost-burdensome overhead while increasing employee satisfaction and allowing your legal practice to accomplish more revenue-related work with the same headcount.

Efficiency

One of the most common ways that AI is pushing law firms and legal departments for the better is in their efficiency. AI tools can overcome traditional obstacles by identifying twisted paths and roadblocks in process workflows, easily evaluating scheduling and calendaring options across a multitude of variables and workflows, and integrating billing codes

on invoices without using attorney or staff time to manually enter them. AI-powered efficiency gains in routine tasks of finance, billing, and compliance can free up valuable time for revenue-producing lawyering.

How Can Artificial Intelligence Benefit Client Representation?

Clients look to civil defense attorneys to provide legal advice, defense options, and strategic representation informed by their deep legal expertise and understanding of mediation, arbitration, and litigation. As discussed in "Current Legal AI Technologies," *supra*, AI-based technologies can help attorneys advance the quality, responsiveness, and power of their client representation through additional options to support e-discovery, legal research, brief writing, contract analysis, due diligence, predictions for litigation and arbitration, and patent and trademark searches.

AI solutions can help small and midsized law firms to overcome growth obstacles. With AI services or tools, law firm size does not necessarily represent significant competitive advantage. Even a relatively small law firm can handle large volumes of information quickly and efficiently without adding additional attorneys. Keep in mind, though, that the use of AI-based technologies does not guarantee success. AI systems require know-how and expertise associated with data and computer science. Further, while certain processes have the capability of providing valuable insights on legal issues, relevant facts, and options, civil litigation can still be unpredictable. Also, AI cannot replace an attorney's solid knowledge of the law, the specific client, and pathways to a strengthened position for succeeding with the legal outcomes and costs associated with a case.

What Are Key Considerations When Choosing an AI Systems Solution?

Rising stars in the AI marketplace are changing all time, in part because the business models for tech and legal services alike are evolving in an increasingly dynamic and innovative environment (and in part based on marketing budgets). Similar to any technical investment, it is important to consider what your needs are and whether the new software will be the right thing. The guiding principle should be whether an AI solution will enable you to sustain or grow your practice. The gains might be in quality, market distinction, productivity, reduced costs, or not losing revenue in the future to another law firm competitor.

Choosing the right AI solution will require asking many questions of the vendors, understanding lots of jargon and concepts to understand their replies, and matching the appropriate solution or solutions with your goals. Consider whether you have the appropriate processes in place within your legal practice to address how to evaluate your needs, assess and select potential AI systems to meet those needs, and conduct ongoing monitoring and evaluation of the benefits and risks, including ethical risks, of any implementation of AI technologies. Before your law firm signs an agreement with an AI solutions provider, prioritize these five key considerations:

Accountability for Accuracy, Reliability, and Explainability

Law firms and legal departments need a meaningful understanding of the AI systems they are using. Specifically, they need to ensure accountability for accuracy, reliability, and explainability.

Accuracy depends on designer and user competence, capabilities of a technology, suitability for purpose, and substantive inputs. Accuracy is measurable. Does your firm have the expertise to deploy the AI tools effectively and measure the results? Is the tool suitable to the purpose for which it is being used? Ongoing effort is needed to ensure that the AI system functions as technically designed and to accurately evaluate the output. Don't hesitate to ask the hard questions and press for answers. What assumptions are used in the algorithms? What assumptions are used in the types of data used to train the algorithm? What types of data were used to create the algorithm? When the algorithms are trained for a specific client, how are they trained? How much training data and staff time to train the algorithm do you need to provide to achieve accurate, reliable, and meaningful results? How do you evaluate outcomes for potential unwanted bias, inaccuracies, inconsistencies, and deviations from expected outcomes, including serious omissions? Is there training on how to evaluate the outcomes? Is the process for evaluating the outcomes a very manual process?

Consider an algorithm trained to identify relevant documents as part of discovery. Does the experience level or type of legal background of your staff influence the training of the algorithm on what is considered "relevant"? Will the documents be classified by degree or probability of relevancy? How frequently will you need to review and train the system? How will you measure the precision and recall of the output—the accepted measures of output in the science of information retrieval to convey the amount of data returned that is on point and the amount of the target data that was actually retrieved? What expertise will be needed to derive those measurements? What precision and recall has the system achieved in independent testing? In what ways can you justify the relevancy of the documents to opposing counsel and the judge? Would that same trained algorithm, with minimal training or additional input, be reliably useful for finding relevant documents for the next discovery project? Can you run simultaneous projects? Would that same trained algorithm's ability to find relative documents work just as accurately and reliably in a due diligence context? How will you be able to investigate whether incorrect or biased patterns emerge based on your data or interactions with the AI software, particularly if driven by machine learning and predictive technologies?

How are updates and corrections handled? If the AI algorithms are trained against legal or regulatory requirements or the tool houses a library of those requirements, what happens when those requirements are updated? How is the corresponding update to the training or library handled logistically? Will you need to pay for those updates? How will you know if the update is accurate and reliable? Will you need to retrain the algorithm and conduct additional audits to verify accuracy and reliability? Given that these retraining and auditing hours might be non-billable, how frequently can you expect to need to do this? How much time might it take?

Closely related is explainability, the extent to which AI solution can be explained and represented so that its functional properties, assumptions, biases, and trade-offs can be accounted for and understood. If you need to disclose your use of AI systems to clients, opposing counsel, or courts, you will want to ensure you can demonstrate your technology competence and that of those deploying these technologies for your firm, be able to explain how results are achieved, and be able to explain your firm's use and reliance on AI.

Additional insights on these topics and the ethical implications for the competent representation of clients are discussed in "Legal Ethics and Social Justice," *infra*.

Information and Data Governance

AI technologies are not only deployed on data for generating results but also are trained, improved, and developed against data. Your data, or your clients' data, will be needed to train and use the system for your particular purposes. Clarifying data format and migration requirements, data rights, and data security should be fundamental and prioritized considerations. Will the AI solution augment, replace, or integrate with your existing systems, documents, records, and processes? Are those systems compatible with the new solution? How much data will be needed to train the algorithm for your law firm's needs? Small and midsized firms will want to understand better if they can gather the appropriate data, and, if so, whether they have enough data. Will your existing data formats be compatible with the new AI solution? If not, who will be responsible for the document and data migration? How will missing information or differences in format be handled? For numerical data, how will differences in precision, particularly rounding, be handled? How will those changes to the data affect the result? If you want to use the algorithm differently in the future, can you change the training data and retrain the algorithm?

Data Privacy and Security

What data privacy and security safeguards are built-in as part of the AI solution? Where is the data stored? Is the data encrypted? Where is it processed? What data is transferred? Is it transferred outside the firm or the client's data storage locations? How does the vendor (and/or your firm) protect data privacy and confidentiality? If your firm or department deals with cross-border data, how does the AI solution comply with non-U.S. privacy regulations? The EU General Data Protection Regulation (GDPR) imposes restrictions on the usage of personal data, including in AI systems without consent and requires companies using AI and other automated processing to provide an explanation to individuals affected by AI decisions. Will your firm or department need to delete data that isn't GDPR-compliant? If so, will the remaining data be sufficient to train the AI and to get meaningful results?

Expenses and Capital Investments

The AI solutions industry continues to grow, driving costs both up for newer sophisticated advances and down for some commercial off-the-shelf and open source AI software. Small and midsized firms will want to assess what tools and technologies they currently use and what gains might be achieved with AI technologies. Don't overlook the costs for set-up,

data and content migration, staff training, staff time to train the algorithm with relevant data, post-implementation audit time to ensure accuracy and reliability, and ongoing costs to optimize the solution for changing real-world scenarios and representation of clients. Depending on the licensing terms, additional costs also might arise for new software versions or updates to reflect newly enacted legal and regulatory changes. Similar to other technology expenses and investments, consider the risks of vendor lock-in and whether the data formats and specific AI solution will be a barrier to the integration or adoption of new innovative AI products and solutions in the future.

Training for Attorneys, Paralegals, and Staff

Training is a must. Similar to information security and ethics training, a crash course for using AI solutions opens the door for future failures and risk exposure. Does the purchase of an AI solution come with orientation and training programs? Is the training sufficient to create the expertise needed for competent deployment and assessment of output? Will your firm or legal department need to pay for training? How much time will be involved? Do any of the training programs qualify for CLE? Is one-on-one training required? How effective is the training? What if the training is not effective? Will busy litigators in your legal practice have time to learn and use the new AI software, or will it go unused? Does the solution fit culturally within your practice?

Network Infrastructure

Will implementing a new AI solution require additional infrastructure, such as higher-performance computing power and higher-speed internet access? If accessing a cloud-based solution, will you need to make alterations to your network security architecture? If the AI software runs on your own servers, will it need to connect continuously or periodically to the vendor's servers? How secure are those connections? What user access safeguards are there? What business and technical strategies are needed to mitigate digital infrastructure and physical security risks?

What Are Strategic Factors When Building Your AI-Powered Law Firm?

How can you make sure your transition to an AI-powered law firm is a huge win? Deriving value from your AI solutions requires strategic management choices related to business growth strategies, competitive positioning, how best to finance that growth, return on investment, protecting that investment, and risk-taking.

Practice of Law and Professional Ethics

Any solution needs to be considered and understood within the broader context of the practice of law and professional responsibilities, including the duties of communication, confidentiality, supervision, and competence. With AI advancing so quickly, what is permissible within a specific jurisdiction or the marketplace could change, perhaps quickly, and could affect how, when, and why you use the AI solution. Will the use of an AI solution

for client representation be the advantage, or will it be perceived as creating too much asymmetrical information over plaintiffs? Will the use of predictive analytics be perceived as a tool of justice or as enabling forum shopping and undermining fairness? Notably, France in 2019 adopted a law banning the use of predictive analytics to compare and predict rulings by specific judges. This is a fast-changing landscape, and your AI strategy needs to include flexibility for potential changes in key legal and ethical responsibilities. Chapter 4 provides further discussion of the professional ethics issues.

Managing Innovation Risks from AI Legal Services

Begin with understanding where you want to be on the risk spectrum for vendors and the technology. As was said earlier, AI technologies continue to evolve rapidly. Is your legal practice positioned to be more risk-taking to gain a competitive edge by using smaller vendors with highly innovative software? A start-up or early market entrant could be high-risk, high-reward. A start-up might also be cheaper or offer an opportunity to shape the AI solutions for your specific needs. The risk, however, is that the cost of innovation might be frustration. Early adopters of newly released AI solutions might struggle with balancing strategic marketplace advantage and achieving consistent and predictable legal services for clients and internal operations. Anticipating and planning for backup options can help mitigate those innovation risks. If your legal practice is more uneasy about the innovation risks, major companies with a history of providing legal support services may provide greater reliability. While you may meet your goal of staying afloat in the marketplace, you could impact your goal of competitive advantage because your competitors are more likely to have access to the same AI solution. Some firms have solved this by implementing custom AI solutions. Another conservative approach to manage innovation risks would be to implement a pilot project for a discrete function of your firm before expanding to multiple areas and then enterprise-wide.

Regardless of the AI solution, your reputation hinges on the trust of your employees and clients. That trust can be lost due to perceptions of misuse of and distrust in AI technologies. As such, the use of AI within law firm business models should include traceability and transparency measures, internally and externally.

Data as the New Competitive Advantage

Your law firm and legal department possess in-depth and specialized legal knowledge in civil defense. Similar to your other intellectual property, this expertise is highly valuable in the competitive marketplace. Are you required to share your data with the vendor? If so, consider the effects of sharing your law firm data and specialized legal domain expertise to train an AI solution to be used potentially by your competitor. Strategic decisions need to be made whether the value gained from contributing data to improve a vendor's algorithms and AI-driven solutions will be realized if those improvements also benefit your existing or future competitors. Will the gains and insights from your outputs create business value for others as they transition and rely on the same AI solutions? Are certain types of your legal domain assets more valuable than others? If you decide to share your data, how does this affect your data retention policies?

Data Science Teams

Some law firms and corporations are establishing specialized data science teams to support and implement AI strategies and data analytics. The teams can include lawyers with data science backgrounds, non-lawyers with data science backgrounds, or both. Some teams operate as the specialized back-office AI team, focusing solely on internal firm analytics to address operational issues. Other law firms rely on the data science team also for client-facing support, such as working with the client on discovery, due diligence, or modeling potential litigation or arbitration outcomes. If you decide to create a data science team, consider the responsibilities for the data science team and how the legal staff and data science teams will operate to avoid unwanted risks and potential concerns involving the unauthorized practice of law by non-lawyers. The ethical issues are discussed further in "Legal Ethics and Social Justice," *infra*.

Improving Revenue and Profitability

How can law firms use AI technologies to help address revenue and profitability? Will the burden and expense of the AI solution bring benefit for your top and bottom lines and for your clients' successes? AI solutions could help you increase your value for clients by improving decision-making, creating better and faster legal services and products, and enhancing client relationships. Increasing revenue and profitability depends on the extent to which AI solutions can help achieve cost savings, grow the profitable volume of your business, or squeeze a narrow profit margin to provide an edge in a specific competitive area.

AI-enhanced discovery brings the promise of saving time, money, and effort, possibly considerably, during civil defense litigation but achieving those benefits requires adequate knowledge and expertise. There are many different products and services; they perform differently and may be well suited for one purpose but less so for another. The combination of tools, services, and methods that you can competently deploy and sufficiently explain when required is fundamental to the ability to depend on or justify time and costs estimates related to using AI solutions and avoid headaches in a meet-and-confer, a court, or a client's office. To avoid potential fallout, you will need to provide quantitative, evidence-based support to argue efficacy, negotiate protocols, seek cost-sharing or support proportionality objections of "undue" burden. Gaining an early understanding of efficacy, time, and costs involved will better position you for discussions with opposing counsel. If the assertion is that AI-enhanced discovery would be overly burdensome, can you justify the proportionality objection? You may want to establish ongoing relationships with potential vendors to help your team assemble the estimates more readily and responsively for specific clients. Also keep in mind that opposing counsel and the court likely will not be sympathetic to an argument of cost-sharing or undue burden if there are other readily available technologies and supporting services that could be deployed at lower costs for the same or better results.

Insurance

Does your law firm insurance cover the use of AI technologies? If a client requests access to review and test an AI system for accuracy and reliability, would your insurance, as well

as your agreement with the technology vendor, cover their access and use? Consider if the court orders disclosure of the AI technology, your methods, and all relevant data, including training data, because its use or results are disputed. Can you meet your duty to the vendor? Can you prove you met your duty of competence and candor? Can you meet your legal, fiduciary, and ethical obligations to your client, particularly your duty to protect confidential, private, and proprietary information and data?

Artificial Intelligence as a Legal Practice Area

The business opportunity to grow your revenue might be to establish a new legal practice area focused on artificial intelligence or to integrate it into your existing legal practice areas. AI technology is poised for transforming our economy, societies, and lives. Every major sector of our economy will be affected. The expansion of AI has implications for products liability, medical malpractice, personal injury, employment claims, tort liability, intellectual property disputes, construction defects, environmental disputes, and more. Related technologies, such as automation, robotics, blockchain, and the internet of things, similarly will drive new legal issues in civil litigation and new specialties. Clients increasingly are seeking help with legal issues, regulatory compliance, contracting issues, and liability as related to their creation, use, and implementation of AI and other emerging technology.

Legal Ethics and Social Justice

The implications of AI systems and autonomous systems for the practice of law warrant renewed consideration of the model rules of professional responsibility; they also warrant new appreciation for the roles of measurement and expertise. While AI technologies and processes raise complex new ethical challenges, they also present opportunities for lawyers to serve clients better, while fostering access to justice.

Legal Ethics

In the field of professional ethics, many rules haven't yet caught up with AI technology. Much like the challenges in applying traditional advertising and client communication rules to social media, the legal profession now has to apply rules to AI systems that were written based on very different ways of practicing law.

Competence (ABA Model Rule 1.1)

A fundamental premise of competent legal representation is that lawyers know not only what they need to accomplish, but also the methods available and whether they are in fact accomplishing what they set out to do. They must also recognize when they are competent to advise alone and when they need to enlist the expertise of others. (ABA Model Rule 1.1.) It has been explicit since 2013 that these requirements extend to the realm of technology. In that regard, competence includes, at a minimum, "staying abreast of the risks and benefits of relevant technology." (ABA Model Rule 1.1 comment 8.) Without question, AI—"the most important general-purpose technology of our era"—is relevant technology. Eric Brynjolfsson and Andrew McAfee, "The Business of Artificial Intelligence," *Harvard Business Review* (July 2017).

In practice, competence requires much more. Although use of AI technologies does not yet represent the standard of care, such that its use is necessary in a particular area of practice, if a lawyer chooses to use an AI system, she must know its efficacy. How well is the purpose accomplished? Deficiencies may emanate from the algorithms, the data sets on which they were trained, or the method of training or deployment. These are not matters for guesswork. Competent lawyers whose processes incorporate AI technologies must know the strengths and weaknesses of the technologies for the intended purpose and what they have actually accomplished in each instance. "If a lawyer uses a tool that suggests answers to legal questions, he must understand the capabilities and limitations of the tool, and the risks and benefits of those answers in the context of the specific case they are working on." David Lat, *The Ethical Implications of Artificial Intelligence*, Above the Law: Law2020,

http://abovethelaw.com/law/2020/the-ethical-implications-of-artificial-intelligence/ (quoting David Curle of Thompson Reuters). Those who advise clients that are developing tools must understand the tools' mechanics and capabilities, data dependence, planned and potential uses, risks and accomplishments.

Efficacy: NIST Studies Show Results Vary Widely

There is an inherent allure in seeing expeditious results generated by an AI technology, particularly if it purports to provide information that would be difficult to obtain a different way. However, decades of research in the field of information retrieval teaches that actual efficacy of tools and methods vary widely. Studies conducted under the auspices of the National Institute of Standards and Technology (NIST) from 2008 through 2011 evaluated the performance of search tools and methods in a litigation context. The studies were open to the public and attracted entrants from business, law, and academia. The objective was for the participant to locate all information responsive to particular document requests from within publicly available sets of electronic documents. Results were measured by recall (what percent of responsive documents were identified) and precision (of the documents submitted as responsive, what percent were truly responsive) for each document request. When posted on a precision/recall measurement grid, the results were splattered across it. They ranged from a low of less than 5 percent recall (finding less than 5 percent of the responsive documents) and under 5 percent precision (meaning fewer than 1 in 20 documents submitted as responsive was truly responsive) to highs in the range of 80 percent recall, 90 percent precision.

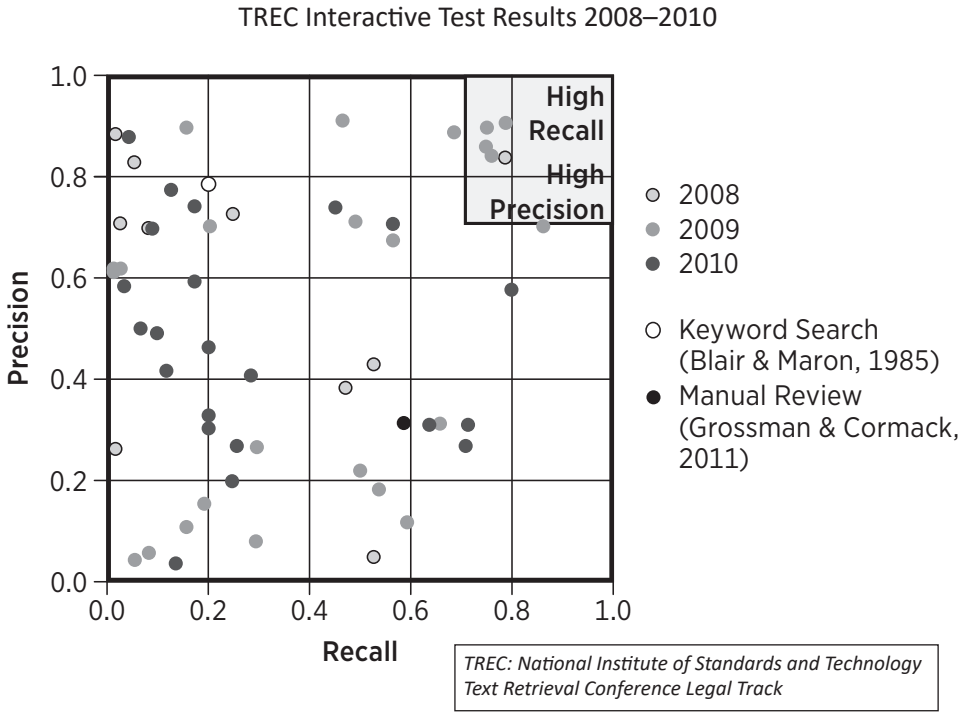

Efficacy: NIST Studies Show Participants Were Inaccurate in Estimating Their Own Results

Importantly, we also know from the NIST research that participants were largely inaccurate in estimating how well they had performed with their search methods and tools. A comparison of participants' predicted success to their actual success in the 2011 NIST studies, as measured under the NIST protocols, showed that participants' predictions of their accuracy were off by as much as 90 points; the average difference between the predicted results and the actual results was 34 points. Accordingly, using this average, a participant who estimated they had found 50 percent of the relevant information may well have actually found as little as 16 percent of it or as much as 84 percent of it, but would be completely blind as to where their results fell in that range (or indeed if they fell below it).

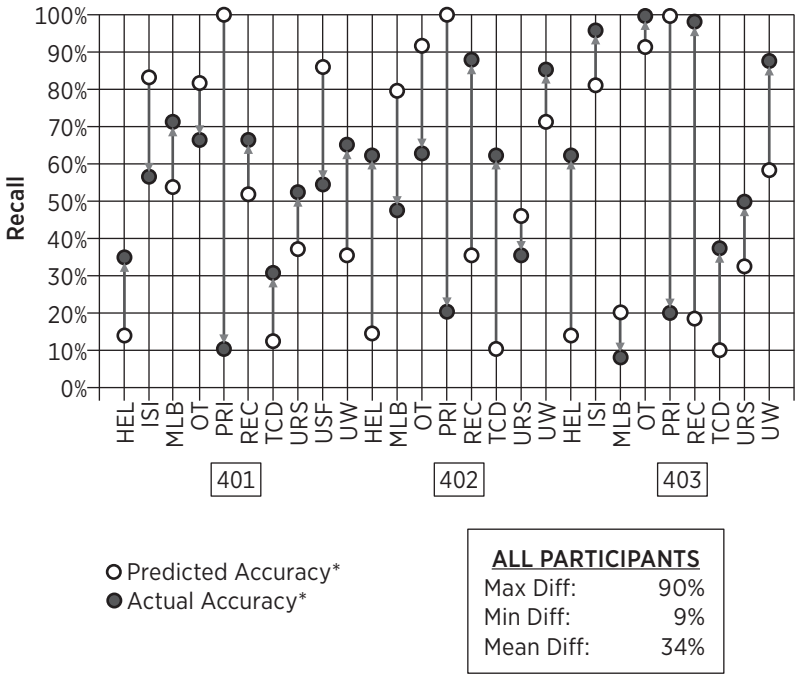

Evidence of effectiveness requires more than just having a result. In essence, what is required is:

1. a clear objective that is reflective of real-world objectives,
2. effective measures by which to judge to what extent an AI process has met the objective,
3. scientifically sound methods to measure effectiveness, and
4. valid data against which to measure effectiveness.

From whom or where does this information come? Ideally, a certification process will develop for the emerging AI technologies appropriate for their purposes, providing lawyers with meaningful information about theoretical capabilities and how to measure efficacy in practice. Short of that, because efficacy varies with data sets, processes, and goals, we will need to increasingly consult with statisticians and data scientists to help us evaluate what

we are actually accomplishing. And, conversely, data scientists and statisticians need to consult with legal professionals in order to ensure that any evaluations of the design of AI technologies and processes do in fact model real-world conditions and objectives.

Evolving Expectations

Another aspect of competence is understanding evolving expectations. What are clients and courts and society expecting from lawyers and businesses in connection with the use of AI? Expectations evolve, and at times even eclipse common usage. In 1932 in *The T.J. Hooper*, the U.S. Court of Appeals for the Second Circuit set an expectation (in essence, a strict liability) relating to the failure of a tugboat company to equip tugs with the latest technology—a radio receiving set—to warn of a coming storm. Tugboats equipped with radios—which were not yet in common use—pulled into safe waterways to ride out a storm. The T.J. Hooper and Montrose tugs, both towing coal barges, rode out the storm in the open water. Two barges sank. While there was no loss of life, the second circuit declared "there are precautions so imperative that even their universal disregard will not excuse their omission." *The T.J. Hooper v. Northern Barge Corp.*, 60 F.2d 737, 740 (2d Cir. 1932). In the AI arena, computer-based legal research and technology-assisted review are well ensconced in expected usage, albeit not yet universally or competently deployed.

New Client and Law Firm Dynamics

AI tools are creating new social and business dynamics with which lawyers will have to be familiar. A powerful example: an algorithmic tool used for pricing became the basis of a price-fixing charge: co-conspirators agreed to price using the same algorithm and were able to set common pricing without ever having to discuss prices. https://www.justice.gov/opa/pr/former-e-commerce-executive-charged-price-fixing-antitrust-divisions-first-online-marketplace (announcing plea deal in *United States v. Topkins*, CR. 15 201 (ND CA 2015) https://www.justice.gov/sites/default/files/opa/press-releases/attachments/2015/04/06/topkins_information.pdf.)

Legal Advice

As noted in prior sections, the responsibility to advise clients who are embedding AI technologies in their products is opening new legal frontiers, and the law has not yet caught up. While lawyers are currently drawing on old paradigms for accountability, the legal frameworks are beginning to change—consider the imperative to address evolving laws governing the privacy and security of data that may be subject to algorithmic processing, for example. In the same vein, the purchase and sale of AI technologies raise the bar for understanding what is being transferred—and whether embedded data implicates rights of data subjects that are not parties to the transaction. How do risks, representations, and warranties change accordingly? As uses of AI technologies increase, lawyers must stay abreast of the technical changes and collateral implications. In addition, as discussed below, imbedded bias in algorithmic output raises ethical questions arise that go beyond solely legal ethics.

Trial Evidence

AI processes will demand new skills from lawyers involved in litigation as well. As AI technologies essentially become the products at issue in product cases (think: medical devices, drones, autonomous vehicles) or contain the evidence in those and other cases (think: Alexa, Echo, drones), competent counsel must be able to sort out the metadata, logs, and results to understand the facts. Trial lawyers will need to become adept at establishing (or rebutting) authenticity and relevance. Will the outputs be self-authenticating under Federal Rule of Evidence 902 (13) or (14)? Is the associated certification under FRE 902 (11) going to apply to the metadata and logs in the tool that may be requested by the opposing counsel? To meet that standard, the electronic record would have to be certified as accurate. Under FRE 902 and 702, trial lawyers will need to challenge assertions that are not supported by valid tests and metrics; they will need to determine and prove (or rebut) measured effectiveness and presence of inherent bias in the tools. Will the bias in the seemingly "impartial" output create prejudice on which the evidence could be precluded under FRE 403?

Communication (ABA Model Rule 1.4)

The model rules of professional conduct require client communication. Lawyers are to "reasonably consult with clients about the means by which the client's objectives are to be accomplished." (ABA Model Rule 1.4) A lawyer should obtain informed consent from the client before using AI. ABA Resolution 112, III, A, 2 (Aug. 12-13, 2019). For that conversation to be meaningful, lawyers will need to learn if there are AI processes that could expedite or improve a current process, and whether they would be effective for the intended purpose, considering both skills required for competent operation and expected performance. What about risk? What are the limitations of the AI process? On what data was it trained? What is the risk of bias? What skills are needed in the process including to evaluate the results? Will the new process increase speed or decrease cost or will required training and validation slow the project and overshadow the benefits? Will a judge or counterparty reject the AI process or demand it? Conversely, a lawyer may need to communicate with the client a decision *not* to use AI systems, particularly where that decision could result in an unreasonable fee, thereby implicating ABA Model Rule 1.5. Clients will increasingly assume that their counsel's discussion of the means by which objectives will be pursued will include discussion of the risks and benefits of use (or non-use) of AI processes.

Fees (ABA Model Rule 1.5)

Closely related to the obligation to communicate is the obligation not to charge an unreasonable fee. (ABA Model Rule 1.5) Undoubtedly, lawyers have broad discretion in setting fees, so long as they communicate clearly with clients. But reasonableness also requires selection of reasonable means. Where technology-assisted technologies, when competently deployed, can produce results as effective as (or more effective than) manual approaches and/or can produce such results at lower cost, those tools will need to be discussed with clients. At times, choices may redirect lawyer hours: rather than spend time on tasks that can more effectively be performed with the assistance of AI technologies, their time will return to the work that requires legal expertise, likely to the benefit of client relationships.

On the other hand, there will be types of AI technologies that require lengthy training on local data sets. The training required might exceed the know-how of the lawyers or require time commitments that will consume any contemplated savings from use of the tool. That additional expenditure should be discussed and factored into decisions both under rules 1.5 and 1.4.

Ethics and fees are often discussed in the context of algorithm-based case analysis and litigation finance. A lawyer's duty of loyalty is to the client, not to the person paying the legal fees. (ABA Model Rule 1.8(f).) Information that algorithms may add to assessment of case merits, likely outcomes, and settlement posture may be valid for consideration in counseling the client, but that counseling must be kept independent of viewpoints or motivations of a funder who covers the legal fees.

Confidentiality (ABA Model Rule 1.6)

Lawyers' duty of confidentiality to clients requires that they "make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to representation of the client." ABA Model Rule 1.6. Lawyers must take appropriate action to make sure that client information is safeguarded. In the AI world, where client confidences sometimes must be shared with AI technologies, or generated, used, stored or even comingled with other clients' confidential information, lawyers face new challenges to ensure that information remains secure and confidential. *See* Drew Simshaw, *Ethical Issues in Robo-Lawyering: The Need for Guidance on Developing and Using Artificial Intelligence in the Practice of Law,* Hastings L.J., Vol. 70:173, 198. The constant risk of data breach makes these challenges all the more significant. In fact, Comment 18 to ABA Model Rule 1.6 now specifies the factors to consider in determining the reasonableness of the lawyer's efforts:

> the sensitivity of the information, the likelihood of disclosure if additional safeguards are not employed, the cost of employing additional safeguards, the difficulty of implementing the safeguards, and the extent to which the safeguards adversely affect the lawyer's ability to represent clients (*e.g.*, by making a device or important piece of software excessively difficult to use.)

ABA Model Rule of Prof'l Conduct 1.6 cmt. 18.

Because use of AI systems often requires sharing confidential information with vendors or other third parties, including those providing hosted solutions, it is necessary to minimize risk imposed by such sharing by ascertaining how the information will be safeguarded. ABA Resolution and Report 112, §III, A, 3. At a bare minimum, a lawyer must become knowledgeable about the vendor's information security policies and practices and should ensure that vendors employ security measures in line with the NIST Cybersecurity Framework or a similarly stringent framework, which will entail asset and access management (including multifactor authentication), firewalls, anti-malware protection, timely patching, logging and alerting, and a host of other safeguards. *See, e.g.,* NIST Cybersecurity Framework; NIST 800 053; Payment Card Industry Data Security Standards; CIS Top 20 Critical Security Controls. And while security is critical, it does not represent the full pan-

oply of confidentiality concerns posed by AI processes. A lawyer must also assure herself that confidential information will be protected from unauthorized or inadvertent disclosure to unintended third parties as it is being gathered, datafied, formatted and used, as well. Simshaw, *supra*, at 200. The ABA report suggests that to inquire "what type of information is going to be provided, how the information will be stored, what security measures are in place with respect to the storage of the information, and who is going to have access to the information." *See* ABA Resolution and Report 112, §III, A,3). It is important to note the expectations are evolving toward evidence-based assessments. *See* FTC commentary by Andrew Smith, Director, FTC Bureau of Consumer Protection (Jan. 6, 2020) https://www.ftc.gov/news-events/blogs/business-blog/2020/01/new-improved-ftc-data-security-orders-better-guidance (point *Second*).

The lawyer's due diligence with respect to security and confidentiality should extend to both her own environment and that of third-party vendors. The due diligence should include inquiry into reputation, security, and treatment of confidential data once received (access, sharing with other third party vendors, who they are, status of nondisclosure and nonuse agreements, data ownership, and notice and disposition in the event of sale, merger, retirement or bankruptcy, and subpoena.. Simshaw, *supra*, at 200. The lawyer must explain to the third parties charged with handling this information the critical importance of keeping it confidential.

Finally, beyond confidentiality, the lawyer must have sufficient information about often opaque AI technologies to meet any obligations regarding privacy and data subjects' rights with respect, most notably, to the EU's General Data Protection Regulation and the California Consumer Protection Act; this all requires understanding what data is exposed to the technologies, gathered by the technology, becomes embedded in the technology, or is transmitted to the lawyer? Commitments should be by contract.

Jury research. Consider, for example, biometric data gathered on shadow jurors and held by law firms. Trial lawyers already gather vast amounts of personal information about potential jurors, including personal history, lawsuit information, and social media usage. The television show *Bull*, which aired for three seasons, was the story of a jury consulting firm that used biometrics, web applications, predictive analytics and social media mining, among other tools, to "know jurors down to their neurons" and help their clients win cases. In many respects, the technology used on the show is already readily available and used by lawyers to profile jurors. The show also spotlighted the use of shadow jurors—individuals who mirror the empaneled jurors physically and demographically, who guide the trial team in its strategic decisions. While traditional shadow juries provide feedback by talking with the lawyers, on television the trial consultants connected the shadow jurors to biometrics such as galvanic skin response iPads and biometric watches) in order to gauge how the real jurors would ultimately vote. Although seemingly farfetched, this technology exists to some extent. What are the ethical responsibilities of lawyers who gather and keep biometric data of shadow jurors? What duties do lawyers have to prevent identity theft by ensuring such data does not fall into the wrong hands?

Consider, as another example, AI tools educated on one client's data that might be primed for deployment on another client's data. What might the results reveal?

After asking the right questions, a lawyer must communicate appropriately with the client about any required sharing of confidential information and how it will be safeguarded. ABA Resolution and Report 112, §III, A, 3. She must also communicate about the client's preferences and expectations. The requirement of Model Rule 1.6(a) that lawyers "shall not reveal information relating to the representation of a client unless the client gives informed consent" no longer merely suggests a "negative obligation to avoid actively revealing client information." Simshaw, *supra*, at 198. Rather, lawyers have "a more positive obligation" to "act competently to safeguard information relating to the representation of the client…against inadvertent or unauthorized disclosure." Simshaw, *supra*, at 198 (quoting Model Rule of Prof'l Conduct r 1.6 cmt. 18 (Am. Bar Ass'n 2016). "When transmitting a communication that includes information relating to the representation of a client, the lawyer must take reasonable precautions to prevent the information from coming into the hands of unintended recipients." Model Rules of Prof'l Conduct r. 1.6 cmt. 19 (Am. Bar Ass'n 2016). And our duties may extend to former clients, as well. (ABA Model Rules of Prof'l Conduct 1.9 cmt 1 (continuing duty of confidentiality).

On a different tack, efficacy of an AI tool may affect the ethics of sharing results. Consider, for example, low-precision results from an AI tool used to find data in response to a regulatory request or document production. Is production of unnecessary data in those contexts a violation of the rules (at least if the client is not fully alerted to the choice in advance)?

Conflicts of Interest (ABA Model Rule 1.7, 1.9)

Extending the concept of confidentiality, what new considerations arise in the re-use of an AI technology trained for a first client? May it be re-used by a lawyer who takes the client and the technology to another law firm? May it be reused to benefit a second client? May it be reused for the first or second client against a former client without running afoul of a duty of fairness to a former client under ABA Model Rule 1.9? If the AI technology were re-purposed for use *against* the first client on a related matter, it would likely pose a conflict of interest under ABA Model Rule 1.7, but what if it were for an unrelated matter? Would any of these analyses depend on how dependent the technology was on the first client's data? Suppose the tool came with fundamental training, and was only honed for the first client. Or would the nature of the data matter? Was it trained using client business data (its contracts, for example) or on publicly available information about the client's case (for selection of judges or jurisdictions or trial teams)? Would the ultimate output of the AI technology or process belong to the client or the lawyer or the law firm—an analysis that would likely affect the portability and reuse by a lawyer switching firms.

Candor, Fairness, and Truthfulness (ABA Model Rules 3.3, 3.4, 4.1)

The obligations of candor to the tribunal and fairness to the opponent (ABA Model Rules 3.3 and 3.4) and truthfulness in statements to others (ABA Model Rule 4.1), intersect yet again with the importance of knowing the results of an AI endeavor (which again requires competence). Before making a representation under FRCP 26(g) that a document production is reasonably complete, for example, or making a similar assertion to a government

regulator, one would have to know if that were true. A statistical journey through the discard pile, sampling to find the percentage of responsive documents that had landed there, would be necessary. The percentage of relevant documents in the discard pile would be multiplied by the size of the discard pile to find the total number of responsive documents not produced. Adding that number to the responsive documents produced would provide the denominator of the recall calculation. The number produced, divided by the total responsive set (produced and not produced) would provide the percent recall. If that percent were high enough, the lawyer could fairly make a representation of completeness. In discovery or trial, in investigations or in negotiations, in order to satisfy her duty of candor, fairness, and truthfulness, a lawyer must understand the AI technology her client creates or the AI process she uses well enough to discuss it accurately, likely to explain to the court how the technology and associated process works and why her argument about it is well-founded. In order to truthfully represent the effect of use of an AI system, a lawyer must understand the actual effectiveness of the system and the metrics of its results and the presence or absence of bias.

Supervision (ABA Model Rules 5.1, 5.3)

In addition to supervising their staff, lawyers also have the duty to supervise the AI processes the firm uses (both those inside the firm and those implemented with vendors and consultants) to ensure that the activities are consistent with the rules of professional conduct (ABA Model Rules 5.1, 5.3). This duty is reflected in the 2012 change to ABA Model Rule 5.3, formerly titled "Responsibilities Regarding Nonlawyer Assistants," but now entitled "Responsibilities Regarding Nonlawyer Assistance." *Variations of the ABA Model Rules of Professional Conduct*, ABA CPR Policy Implementation Committee (Sept. 29, 2017) https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/mrpc_5_3.pdf. ABA Resolution and Report 112 notes that this change is intended to "clarify that the scope of Rule 5.3 encompasses nonlawyers whether human or not." ABA Resolution 112, p. 6. Model Rules 5.1 and 5.3 obligate lawyers to supervise the work of AI processes used in the provision of legal services and to "understand the technology well enough to ensure compliance with the lawyer's ethical duties." *Id.* Included in this obligation is the duty to make sure that work product produced by each AI process is accurate, complete, and does not risk disclosing confidential client information. *Variations of the ABA Model Rules of Professional Conduct, supra*. Likewise, supervisory lawyers must know what tasks appropriately can be enhanced with AI technologies and which cannot, David Lat, *The Ethical Implications of Artificial Intelligence*, Above the Law: Law2020, p. 8, http://abovethelaw.com/law/2020/the-ethical-implications-of-artificial-intelligence/. Striking the balance between over-reliance on AI technology on the one hand, or under-utilization of it on the other, will be critical to fulfilling the duty of supervision.

To carry out their obligations, supervisory lawyers will need to understand what is being accomplished by AI processes the firm uses—including how well they perform and with what potential biases. Drawing on external expertise as warranted, they need to learn what questions to ask, what methods of validation are adequate, whether that validation occurred, and with what result. They need to understand the competencies required for deployment of AI processes and the accuracy and behavioral measures of the results that

will serve as the basis for representations to courts, communications with clients, legal advice, and billing decisions. In addition, supervisors should ensure there is adequate training for associates and nonlawyers to understand the importance of these factors. They need to become facile at translation between law and engineering.

Nondiscrimination

Among the strictures in ABA Model Rule 8.4(g) is the prohibition against discrimination. AI technologies contain built-in limitations and biases. They are built by people, after all, who make assumptions and overlook things. And they are trained on data that may be skewed and will likely carry with it long-standing societal biases. Depending on the data selected, developer biases can be exacerbated. Take the example of AI technologies used to sort job applicants. Bogen, Miranda, *All the Ways Hiring Algorithms Can Introduce Bias* (Harvard Business Review 2019), https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias**.** If programmed to look for patterns in candidates with long tenure or frequent advancement, they will replicate biases from the past, without specific instruction. Suppose algorithms are trained on the backgrounds and credentials of employees who have successfully advanced to partner at large law firms. Candidates who most closely resemble current and past partners will be ranked most highly. Immediate bias. The pool that advances is unlikely to represent a diverse cross section of law school graduates, much less a representative cross section of the diversity that makes up society. *See* https://www.brookings.edu/research/challenges-for-mitigating-bias-in-algorithmic-hiring/; *see also,* https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G. Discrimination based on age and disabilities are emerging as issues as well, in part because the lack of accessibility of the tools to those potential applicants. https://news.bloomberglaw.com/daily-labor-report/ai-creeps-into-hiring-and-some-ask-if-its-injecting-new-biases. Similarly, AI technologies trained on one race, ethnicity, or gender will result in tools that are not inclusive, disparately affect one or more groups, and thus are potentially unfair.

ABA Model Rule 8.4(f) also prohibits harassment. AI technologies can pick up abusive language and attitudes from the data on which it is trained. Consider the circumstances of Microsoft's 2016 debut of Tay, a text based chat-bot. While the idea was that Tay would learn to converse in a sophisticated way, with an emotional dimension, the developers failed to account for the possibility that the bot would interact with those with discriminatory viewpoints, and would begin to mimic their hateful speech. Jonathan Vanian, *Unmasking A.I.'s Bias Problem*, Fortune, June 25, 2018. It is the results and output of AI processes that must be evaluated, not just the purpose or intent.

For lawyers, judges and others in the legal profession using or defending AI technologies—including those used in employment decisions and criminal sentencing—this means determining that the extent to which technologies and processes that have built in biases, that the biases have been addressed, and taking care to ensure that the outputs and uses comport with a lawyer's ethical duty to avoid discrimination.

Independent Judgment (ABA Model Rule 2.1)

ABA Model Rule 2.1 provides for lawyers to use independent judgment to render advice: "[i]n representing a client, a lawyer shall exercise independent professional judgment and render candid advice." This exercise may require a lawyer not only to refer to the law, but also to moral, economic, social, and political considerations, as necessary to adequately address both the client's legal and non-legal needs. Model Rules of Prof'l Conduct, 2.1 (Am. Bar Ass'n 2016). Lawyers using AI must be mindful of its limitations. Unlike human beings, AI has no instincts, no ability to analyze relevant non-legal factors, and no ability to take information into account beyond the observational data to which it has access. Simshaw, *supra*, at 203. Lawyers must take care not to marginalize these human factors as they adopt AI. *Id. See* Catherine Nunez, *Artificial Intelligence and Legal Ethics: Whether AI Lawyers Can Make Ethical Decisions*, 20 Tul. J. Tech. & Intell. Prop. 189, 204 (2017) (noting that attorneys must utilize research skills along with their individual professional and moral judgment—qualities with which AI software is not yet equipped).

Lawyers should also be cognizant of the extent to which their professional judgment remains truly independent as they become more and more reliant on AI technologies. In order to remain independent when using such technologies, lawyers must understand the system's design and the efficacy of the results. To merely rely on the outputs of AI technologies to the exclusion of our independent judgment is dangerous and potentially unethical because it makes us "dependent on the judgments of a technological apparatus." Katherine Medianik, Note, *Artificially Intelligent Lawyers: Updating the Model Rules of Professional Conduct in Accordance with the New Technological Era*, 39 Cardozo L. Rev. 1497, 1517 (2018). After all, legal advice is not merely pattern recognition for prediction. "Our legal system is about reasons as well as outcomes—reasons, asserted by lawyers and memorialized in judicial opinions, which provide a continual opportunity through which to debate and potentially change the law." Simshaw, *supra* at 204, quoting Dana Remus & Frank Levy, *Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law*, 30 Geo. J. Legal Ethics 501, 548–49 (2017).

AI in the Legal System

AI technologies are poised to impact the legal system. Efficiency and reduced cost would be desirable. But if the legal profession doesn't exercise due care, that potential may be overshadowed by institutionalizing bias and error at a secret and unprecedented rate. Indeed, the August, 2019 ABA Resolution 112 calls the attention of the profession to the heightened learning and scrutiny needed in considering use of AI technologies in legal practice. https://www.americanbar.org/content/dam/aba/images/news/2019/08/am-hod-resolutions/112.pdf. After all, these pattern-seeking algorithms are written by people who design them to extract certain information. What they are instructed to extract (or weight heavily) and to ignore (or rank lightly) is, fundamentally, a choice made by the programmer.

Further, AI technologies must be trained. At times that training is handled by the purveyor; at other times it needs to be performed by the customer. Data sets used for

training—compiled or selected by the technology purveyor or the customer—are prone to limitations and biases, at times because those biases are enmeshed in our society, at other times because of the coincidental or skewed method by which the data was collected or selected. The technologies often amass libraries of answers about whether particular pieces of information should be ranked high or low, relevant or not relevant to the purpose. Then they designate if the next piece of data is more like or less like data seen before and rank or treat that next piece of data accordingly. Depending on the technology and on the skills of those engaged in the training and deployment processes, it may take a significant amount of time to reach optimal performance—or optimal performance may never be reached.

Accordingly, the efficacy and trustworthiness of a particular tool has many dependencies. The inherent optimal performance of the algorithm (given quality and biases), the competencies of the trainer, the quality of the training data, and whether the tool is being deployed for its intended purpose all affect output and trustworthiness. In addition, the efficacy of an AI system depends on the competency of the operators and the process in which the technology is deployed. These users should possess the requisite skill expected in one skilled in the art of its operation. Given all these variables, lawyers must be cautious not to overestimate the performance capabilities of AI processes.

Underpinnings of Trustworthiness

Confidence that an AI process has worked *fairly* requires more. And without fairness—an inherent value for operation in a trustworthy legal system—it is reasonable to posit that the AI tool should not be deployed. The Institute of Electrical and Electronics Engineers (IEEE) established a global initiative on Ethics of Autonomous and Intelligent Systems. The law committee of that IEEE initiative has considered ethically aligned design and identified four principles that underpin informed trust. The principles comprise transparency, measurement, competence, accountability. They should be present for any AI process or autonomous system before adoption in the law. The Council of Europe has generated compatible principles for the ethical use of AI and autonomous systems, adding a fifth principle: security.

Efficacy

Knowledge of effectiveness is a fundamental tenant of trust. Effectiveness for a purpose must be measured reliably in a scientifically valid manner, and disclosed, for confidence in the legal system to persist. It is not clear that measurement at optimal operation is sufficient. How did the AI technology and process function in this instance for this use as deployed by these operators?

Competence

Clarity on the skills needed for competent operation of an AI enabled technology and participation in AI processes—which will be the standards against which to evaluate the skills of the operator in a particular instance—are core to evaluating the trustworthiness of a particular system in a particular situation.

Transparency

Trade secret protections are important to the continued development of new systems and uses. Transparency, however, is fundamental tenant of a trustworthy legal system. The dimensions and methods to exhibit appropriate transparency will continue to evolve. At the same time, the need for transparency is paramount, particularly as engineers work to design systems that keep learning in ways that are opaque even to the very engineers that develop them. Purveyors of AI technologies should be able to disclose at sufficient depth how the technology is designed and trained and how it operates (including what information is gathered, used, stored, and/or transferred). These operations should be logged to create an auditable trail. Purveyors should also disclose what is unknown—including as best they can how the systems may evolve in ways beyond human control. They should disclose known biases and aspects of training or use that may create bias or lead to biased outputs. This transparency is important for users and other stakeholders in just legal systems.

Accountability

Designers, purveyors, trainers, and users of AI technologies and processes each have accountability for the impact of the technologies and processes in the legal system. Those whose inputs have effects that may affect fairness and justice should be available to provide the information by which those effects can be known and evaluated. In addition, each should be accountable for the transparency needed for the fair operation of the legal system.

Not all purveyors of AI technologies and processes will endorse these principles. But that is not the question. The question is what should be known and transparent before the government and participants in the legal system use technologies and processes; the effects should be transparent and measurable, and the uses monitorable. Purveyors are not forced to play in this space, but lawyers and the legal system should require this clarity as a ticket to entry.

Fairness and Responsibility

A legal system functioning under the rule of law requires transparency and fairness. Liability for the intended and unintended effect of AI is insufficient protection and is costly and difficult to impose. It is not a remedy for a legal system distorted by unintended influences or a failure to deliver justice to those who have entered its realm. And liability is too easily transferred by contract to unwitting buyers, who lack awareness of efficacy or effect, lack competencies, and are bereft of the benefits of transparency. When those buyers then deploy AI processes in the legal system without adequate understanding of effects, the results can chip away at the rule of law. Lawyers, of course, will continue to consider risk of malpractice, with the unnerving uncertainty presented by whether and how to use new AI technologies and attendant capabilities for which skill-building is still underway. What if a client's outcome is prejudiced by lawyer's premature reliance on AI technologies or processes? On the other hand, the prejudice may stem from a failure to rely on AI and those with needed expertise in its deployment. For example, a lawyer might be faulted if key documents are overlooked that a properly deployed AI process could have located or key

AI metadata is missed in a product case. Indeed, those overseeing legal ethics at the state court levels periodically implore lawyers to keep up with changing practices in order to meet ethical obligations, and that in turn will gradually influence views of what constitutes the standard of care. *See* California Standing Committee on Professional Responsibility and Conduct Formal Opinion (2015-193); *cf. Wm. Grossman Construction Associates, Inc. v. American Manufacturers Mutual Insurance Company,* 256 F.R.D. 134 (2009).

How serious is the threat that the use of AI-enable processes in the administration of justice or the practice of law will create distrust in legal systems or ethical pitfalls for lawyers? Very serious, already.

Due Process

Sentencing. Fast forward to the case *Wisconsin v. Loomis,* a criminal case in which the Supreme Court of Wisconsin, on a certified question, approved the sentencing court's use of the prosecution's submission of a biased algorithmic ranking of likelihood of recidivism as a factor in sentencing. Studies had shown that the algorithmic tool, COMPAS, over-predicted the likelihood of recidivism for people of color, and under-predicted the likelihood for whites. The defendant was not permitted to interrogate the algorithm or its training for reasons of trade secret. The due process and social justice concerns are palpable. For most protectors of the legal system, those concerns are unlikely to be assuaged by the Supreme Court's stricture that in future cases the sentencing judge must be warned that:

- the tool's embedded factors and scoring are unknown;
- the tool's scores don't predict behaviors of an individual;
- studies found the tool biased against minorities;
- national data, not Wisconsin data, was used; the tool's accuracy may vary as population changes;
- the tool was not developed for sentencing; it was intended for treatment, supervision, and parole.

While paling in comparison to due process concerns, it is noteworthy that if the defense counsel did not know to ask about effectiveness, data sets, bias, or the intended purpose of the tool, or the prosecutor (or civil equivalent) was inaccurate in what it represented, there would be competence, candor, and truthfulness issues as well.

Discrimination

Another area of concern emanates from use of facial recognition software. Use of AI technologies that deploy facial recognition algorithms is continuing to expand, including in policing. It is reportedly in use at border crossings and in police stations, as well as by landlords in low and moderate income apartments. According to a series of studies performed by NIST, these tools have widely varying efficacy and are prone to bias, incorrectly matching those of African and East Asian descent with much greater frequency than those of European descent (although for algorithms developed in China, those of European descent are misidentified more often than individuals of Asian descent). The young and the old are mismatched more often than those of middle age. Grother, P., Ngan, M., Hana-

oka, K., NISTIR 8280 Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects NISTIR 8280 (National Institute of Standards and Technology) (December 2019), https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf. The risk in practice is that dark-skinned minorities are more likely to be unjustifiably arrested and subjected to our criminal justice system than whites. And poor outcomes stem from mere introduction to the system. The Brennan Center for Justice, "Mass Incarceration Gets Attention as an Economic Issue (Finally)(Sept. 13, 2013) https://www.brennancenter.org/our-work/analysis-opinion/mass-incarceration-gets-attention-economic-issue-finally.

Depending on use case, outcome prediction technologies discussed earlier can raise fundamental ethical issues. We must ask, for example, whether the input features used to construct the model are features to which society has agreed the judicial system should be blind. We must assess the extent to which there are biases in the case law on which the system is trained; if there are, then the system can be expected only to perpetuate those same biases. To the extent such systems are used by judges to support their decision-making processes, the principle of judicial independence may be undermined and rights of appeal weakened. (*See* European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment adopted by the Council of Europe European Commission For The Efficiency Of Justice (December 2018) at 8-9, 23-24, 55-56, https://rm.coe.int/ethical-charter-en-for-publication-4-december-2018/16808f699c).

Predictive Policing

AI technologies have been in use for some years in predictive policing. Edwards, E., *Predictive Policing Software Is More Accurate at Predicting Policing Than Predicting Crime* (ACLU)(August 31, 2016) https://www.aclu.org/blog/criminal-law-reform/reforming-police/predictive-policing-software-more-accurate-predicting. Ostensibly, these tools assist police in deciding where to deploy resources with the proffered benefit of reducing crime. The Economist: *How data-driven policing threatens human freedom* (2018) https://www.economist.com/open-future/2018/06/04/how-data-driven-policing-threatens-human-freedom. But there is no evidence that that purpose is served. Instead, these tools cause police to deploy to high crime neighborhoods; their presence increases the likelihood of arresting residents for minor infractions, which in turn delivers more minorities into the criminal justice system and perpetuates the metrics that the area is "high crime." Edwards, *supra;* The Economist, *supra*. Consider, for example, the disparate impact of predictive policing on arrests for underage drinking. For upscale teens, that drinking occurs under the bleachers or in college dorms. For minority teens, it is more likely to occur in the neighborhood, and to be observed by the hyper-vigilant police deployed at the behest of the predictive policing tools. *See*, O'Neil, C., *Weapons of Math Destruction* (2017) (data bias creates inevitable, biased outcomes). The negative consequences of exposure to the criminal justice system, both direct, Digard, L., Swavola, E., Justice Denied: The Harmful and Lasting Effects of Pretrial Detention (April 2019) https://www.vera.org/downloads/publications/Justice-Denied-Evidence-Brief.pdf, and indirect, Martin, E., *Hidden Consequences: The Impact of Incarceration on Dependent Children* (March 1, 2017) https://nij.ojp.gov/topics/articles/hidden-consequences-impact-incarceration-dependent-children, is then disproportionately experienced by minorities.

Advising in the Face of Bias in AI Processes Outside the Legal System

Bias is increasingly reported in uses of AI tools outside the legal system as well, which will affect both social justice and lawyer responsibilities in advising clients who create, buy or sell, or deploy these tools. AI technologies have been widely touted and bias has been reported with increasing frequency in the popular press. In fact, there has been little properly designed research, and even less that has been subject to peer review.

Housing

Algorithmic tools used by landlords and lenders in housing/tenant/credit decisions are reported by studies to demonstrate socio-economic and racial bias. Sisson, P., Housing Discrimination Goes High Tech (Dec. 17, 2019) https://www.curbed.com/2019/12/17/21026311/mortgage-apartment-housing-algorithm-discrimination; National Fair Housing Alliance, "Defending Against Unprecedented Attacks on Fair Housing: 2019 Fair Housing Trends Report" (pp. 45-46) https://nationalfairhousing.org/wp-content/uploads/2019/10/2019-Trends-Report.pdf. Under the Equal Credit Opportunity Act, 15 U.S.C. §1691, *et seq.*, it is unlawful for a creditor to discriminate against an applicant for credit on the basis of race, color, religion, national original, sex or marital status, or age, or because the applicant derives any income from public assistance. Thus, it is critical for creditors to employ algorithms trained to eliminate these biases.

Employment

As noted previously, algorithms used to weed through online job applications are often programmatically designed to match or rank based on characteristics possessed by those previously successful within a company (by any of a variety of measures). To the extent the company has historically favored males from certain schools or zip codes or who played certain sports, for example, the algorithmic ranking will drive the company toward applicants with similar characteristics and suppress diverse candidates. Bogen, Miranda, *All the Ways Hiring Algorithms Can Introduce Bias* (Harvard Business Review 2019), https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias.

Education

Certain organizations have recognized the need for school systems to monitor and share with the public the algorithms used to assign children to public schools and monitor them in school. Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force (December 2019) https://ainowinstitute.org/ads-shadowreport-2019.pdf; s*ee* Elizabeth Zima, *Could New York City's AI Transparency Bill Be a Model for the Country?* Government Technology (Jan. 4, 2018), https://www.govtech.com/policy/Could-New-York-Citys-AI-Transparency-Bill-Be-a-Model-for-the-Country.html (establishing a task force).

Healthcare

Healthcare products raise a number of issues. First, many (probably most) do not have the data to support efficacy that we have come to expect for medical products. Many are

approved without clinical trials or peer review of data, through a fast track approval path under FDA's 510(k) pathway based on similarity to products introduced decades ago. Second, the data sets are predictably biased, resulting in disparate accuracy and utility in diagnosis or treatment predictions for ethnic, racial, gender, age, socio-economic populations that are under-represented. PBS Health, *"Are health care claims overblown about artificial intelligence?"*(December 30, 2019), https://www.pbs.org/newshour/health/are-health-care-claims-overblown-about-artificial-intelligence (discussing inter alia results published by Matheny, M., S. Thadaney Israni, M. Ahmed, and D. Whicher, Editors. 2019. *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril.* NAM Special Publication. Washington, DC: National Academy of Medicine., https://nam.edu/wp-content/uploads/2019/12/AI-in-Health-Care-PREPUB-FINAL.pdf.

Broader questions of fairness are arising as well. The capabilities of AI technologies—and the means by which data is collected to enable them—is increasingly opaque to the user. If a camera on our TV is recording us, our refrigerator is reporting our eating habits, and our health tracking watch is providing our health statistics to insurance companies, we are becoming the product of our own tools. And the situation extends to the bystander. When Alexa is in use in our doctors' offices, ostensibly to turn on and off the equipment but nevertheless transmitting our medical concerns and diagnoses, we have moved well beyond the caution of "buyer beware." We are being unknowingly affected by those who purchase and deploy AI technologies, who in turn may not know the impact or legality of the uses they are making of them. Similarly, when facial recognition software is deployed toward a Planned Parenthood (or any medical) facility, our core privacy is at risk.