

Robo-Lawyers

Your New Best Friend or Your Worst Nightmare?

GARY MARCHANT AND JOSH COVEY

Gary Marchant is a professor at Arizona State University's Sandra Day O'Connor College of Law, Phoenix.

Josh Covey is with Quarles and Brady LLP, Phoenix.

The headlines are ominous:

"Artificial Intelligence Beats Big Law Partners in Legal Matchup"

"The Verdict Is In: AI Outperforms Human Lawyers in Reviewing Legal Documents"

"AI Beats Human Lawyers at Lawyering"

"The Robot Lawyers are Here—and They're Winning"

"Artificial Intelligence Beats Lawyers Again"

"Now That Lawyers Have Lost to AI, What Is the Future of Law?"

Those, and other examples like them, have led many people in the legal industry to subscribe to one of the following two beliefs about artificial intelligence (AI):

1. **Overbuying the Hype:** AI will replace all human lawyers in the next couple of years.
2. **Ignoring the Facts:** What lawyers do is so bespoke—customizing everything precisely for clients' unique situations—that no machine will ever be able to do the legal work that clients need.



Both beliefs are wrong. They bring to mind a phenomenon that Bill Gates addressed in his autobiography years ago: Humans *overestimate* the impact of emerging technologies in the short term and *underestimate* the impact of emerging technologies in the long term.

The first belief makes a great headline. Yet, it fails to account for AI's many limitations, such as a lack of human common sense, bias in algorithms, and

the impact of bad data on AI's outputs. AI is certainly superior at various narrowly defined tasks and at processing vast amounts of information. But humans remain superior in applying common sense and critical thinking to information, leaving key legal skills—like effective client interaction, advocacy, and negotiation—beyond AI's current capabilities.

AI will not evolve quickly enough in the near future to replace all lawyers. Rather, AI is replacing certain *types* of legal work and already has done so. AI's speed in completing certain tasks reduces the time human lawyers need to spend on those tasks. In the aggregate, that phenomenon reduces the need for human lawyers.

So there will be, and already has been, some displacement. But

human lawyers are still indispensable, and technically proficient lawyers likely will be in even higher demand.

The second belief deserves more discussion, however, because many more people subscribe to it. The belief that AI cannot perform legal tasks or automate legal process—that it's all hype, no substance—simply does not align with real experience and current facts. Indeed, AI is already dramatically affecting the practice of law.

Current Abilities of Artificial Intelligence

Two years ago, JPMorgan started using AI-powered software to review commercial contracts in seconds, saving it 360,000 billable hours a year. Moreover, a 2016 study by Deloitte found that AI is already causing job losses in law firms—mostly in lower-skill non-attorney positions—and predicts that over 114,000 law firm jobs will be lost due to automation in the next 20 years in the United Kingdom alone.

Evidence abounds that this trend will continue:

1. In 2017, McKinsey & Company estimated that 23 percent of a lawyer's job can be automated. Deloitte asserts that 39 percent of legal jobs can be automated. We're currently nowhere near those numbers.
2. Economic pressures and increased competition will continue to be an incentive for providers of legal services to find ways to deliver more services at lower costs, which is precisely the type of efficiency that AI can provide. Relatedly, a recent Accenture report asserts that businesses investing in AI and human-machine collaboration could see a 38 percent increase in revenue by 2022.
3. Clients in most industries are actively incorporating AI in their own internal processes and operations, and increasingly ask their law firms why they are not using similar cost-saving and efficiency technology tools.
4. Society has been through this repeatedly. Once machines reach a level of sophistication at which they can compete with humans at specific tasks, machines quickly evolve to become superior in performance, reliability, and associated costs.

So if AI is the future of successful lawyering, what exactly is it?

AI in the legal field is the subject of a lot of hype, as any emerging technology is. What's a realistic understanding of what it can and cannot do?

It is important to note that AI is not limited to machines that move and act as humans do, like the robots in the TV series *Westworld*. Rather, it refers to software performing tasks that historically have required human intelligence. AI is therefore best conceptualized as an umbrella term that covers a nexus of

technological dynamics: Exponential increases in computing power; an explosion of digital data; and the continued evolution of machine learning, sophisticated algorithms, natural language processing, expert systems, neural networks, and similar technologies.

Yet, algorithms, machine learning, and other technologies underpinning AI have been around for decades. So why is AI getting so much attention right now?

It's a simple formula: Big Data + Today's Computing Abilities = A Game Changer.

The combination of big data and current computing power results in unprecedented data analysis. Such analysis yields insights into meaningful trends in historical data and, increasingly, in real-time data. These insights then feed into machine learning—the ability of machines to use algorithms to learn from data and adapt accordingly, as opposed to simply following what they've been programmed to do.

The result is faster and more accurate data-driven insights propelling machine learning and accelerating the pace of AI's evolution and its practical applications.

Remarkably, with machine learning, computers no longer perform tasks only according to human programming. Instead, the AI system learns to do tasks and improves itself by analyzing data, patterns, and outcomes. This change from the rules-based approach of computer programming to the data-based approach of machine learning triggers important implications for the substance and process of law.

As to substance, individuals may not intend, or even anticipate, some of the actions and decisions that machine-learning algorithms will adopt, making legal accountability a challenge. As to process, a machine-learning system has to be trained and learn over time, much as a child learns, which means that a law firm cannot simply take up a machine-learning product and expect it to operate at peak capability on day one. Rather, it could take months or even years of training to achieve optimal performance, which can create a significant first-mover advantage.

These substance- and process-based dynamics further support the development of AI-powered products and services that improve the efficiency and effectiveness of performing legal tasks. Progress is being made and it is being made quickly. While the prospect of machines automating legal tasks wholesale is still a relatively nascent concept, exponential change dictates that the future will yield substantial progress.

Indeed, some of the litigator tasks in which AI can now assist lawyers—and in some cases already replace them—include the following.

Document review and discovery. Large law firms have adopted AI more in the area of document review than in any other. Recent research shows that, due to automation and outsourcing, attorneys at large law firms now spend only 4 percent of their

time on document review, a substantial reduction from even recent times.

In fact, six years ago, Maura Grossman and Gordon McCormack published an article in the *ABA Journal* about their empirical research, showing that technology-assisted review of electronically stored information is at least 50 times more efficient than manual review. Advances in AI will continue to widen this gap by improving the accuracy and efficiency of document analysis.

For example, in February 2018, LawGeex released a study that pitted AI against 20 highly experienced human attorneys in reviewing five nondisclosure agreements and accurately identifying risks. The average accuracy rating was 94 percent for the AI; the average accuracy rating was 85 percent for the attorneys.

More notably, the lawyers took an average of 92 minutes to review the nondisclosure agreements. The AI took just 26 seconds. That is, AI was *212 times* more efficient, and more accurate, than the lawyers. And that was in reviewing just five documents. When the number of documents increases to the thousands and human fatigue sets in, AI dominates. And AI doesn't mind working weekends.

Legal research. AI improves both efficiency and effectiveness in legal research. AI can read a million pages of legal documents in one second. Advances in fields such as natural language processing are quickly improving its ability to understand the words and meanings of those million pages, helping lawyers complete more comprehensive legal research in less time. A 2018 whitepaper released by Blue Hill Research indicates that legal researchers using ROSS Intelligence's AI-powered tools can reduce time spent on legal research by 30 percent.

Drafting of pleadings. AI systems are being deployed to automatically generate a variety of legal documents—wills, trusts, divorce papers, contracts. The next step is for AI to generate pleadings in litigation, which already has begun. AI vendors like LegalMation can draft litigation pleadings at a fraction of the time and cost that a human lawyer will take. Using the IBM Watson AI system, a company or law firm can just upload a copy of a complaint into the LegalMation program. In just two minutes, it automatically will generate an answer and discovery requests.

Companies such as Walmart have signed on to use this AI system in most new lawsuits against the company, which is expected to result in major savings in costs and time, displacing as much as 10 hours of attorney billable time per lawsuit.

Case analysis. Various AI products also help attorneys make more effective legal arguments. Vendors such as Ravel Law and Lex Machina offer analytics services—using big data analytics and machine learning—to predict how individual judges might rule, identify which precedents they will find most influential, determine the percentage of cases decided on summary judgment, and even detail how to make briefs more persuasive by customizing them for specific judges. Some of these programs

even provide advice on the legal theories a particular judge is most likely to accept and the best phrasings or words to use, or to avoid, in a brief.

These analytics programs also can help with forum selection, providing insights into which forum is most likely to produce a favorable outcome. Other products will help size up opposing counsel, including their track records and previous successful and unsuccessful arguments. Still other services provide analytics on expert witnesses and their strengths and weaknesses for a particular matter.

Technology-assisted review of electronically stored information is at least 50 times more efficient than manual review.

Case staffing. Analytics products that combine big data, data analytics, and machine learning can help law firms and clients evaluate the optimal staffing for specific types of cases in particular jurisdictions. Wolters Kluwer, among other vendors, offers an analytics platform based on over \$100 billion in billing records. These databases enable both law firms and clients to benchmark lawyer performance in terms of hours billed, staffing approach, and case outcome. It also allows law firms to use comparative and competitive analysis to make their own strategy and decisions about staffing of new matters and to better understand their market competition.

Outcome prediction. In deciding whether to pursue litigation and in developing litigation strategy and settlement options, clients often will ask litigators about the chances of success in a particular matter. Attorneys will usually offer a qualitative assessment or even a range of probabilities, based on their experience and the strengths and weaknesses of the specific case, but their estimates are usually little more than a hunch or rough guess.

AI systems can incorporate and integrate far more data than the human mind can handle and provide quantitative predictions of success that are now consistently outperforming the best attorney estimates.

For example, the vendor Case Crunch conducted a highly publicized head-to-head contest in October 2017 to test whether its AI platform could beat 100 experienced British lawyers in predicting the legal outcome of numerous financial mis-selling claims. In the end, the Case Crunch software outperformed the lawyers, with Case Crunch's predictive algorithms correctly predicting 86.6 percent of the claims versus 62.3 percent for the human lawyers. The margin of victory was attributed to the AI system's better ability to incorporate the importance of non-legal factors in predicting case outcomes.

Law professor Daniel Katz and colleagues similarly used a machine-learning algorithm to predict Supreme Court outcomes. Their AI system was able to correctly predict 70.2 percent of the Supreme Court's decisions and 71.9 percent of the justices' votes. That surpasses the performance of legal experts and other strategies to predict case outcomes, which on average achieve a 66 percent success rate or lower.

Jury selection. Artificial intelligence can also be used to analyze and integrate large data sets to predict juror responses to specific cases and arguments, which can then be used in voir dire and for jury selection. The vendor Voltaire, for example, integrates online research on jurors with machine learning to provide actionable advice on individual jurors based on extensive analyses of personality traits and psycholinguistics. The company provides the attorney a dossier on each potential juror based on the analysis of its AI algorithms.

Online dispute resolution. A number of courts and companies around the world are exploring the use of AI to provide online dispute resolution (ODR). One example is Modria, which originated from the eBay and PayPal dispute resolution system, and it offers private ODR that already has been used more than a million times. Modria was recently purchased by Tyler Technologies, which is working with state courts to develop state and local ODR services that will be integrated with the court systems.

AI-based ODR systems are also being explored in the United Kingdom and the Netherlands, among other jurisdictions. While such AI-based ODR systems have great promise in helping to address the access-to-justice problem, their current objective is to litigate many low-value cases that lack significant roles for attorneys.

This broad range of current AI capabilities in litigation raises important business, strategy, and ethical issues for litigation attorneys and their firms. One issue is the retention of AI vendors. Although some large law firms are developing their own in-house AI capabilities, most firms partner with one or more outside vendors because vendors typically offer products limited to one or two functions, such as legal research, document review, case analysis, or outcome prediction.

Therefore, each firm must conduct its own business case analysis to determine whether to contract with one, or multiple, AI

vendors. If a firm chooses to do so, it must then decide *which* vendors to retain. There currently are multiple competing vendors in each AI legal domain. Yet, as in any such crowded start-up markets, one or a few industry leaders will inevitably emerge, just as has occurred in the e-discovery market.

AI Effects on Firms and Lawyers

But law firms face a Catch-22. If they wait until the industry leaders emerge, they will be put at a significant competitive disadvantage, especially given the learning curve that machine-learning AI systems require. On the other hand, if a firm contracts with a vendor and that vendor loses ground and eventually withdraws from the market, the law firm will be left stranded with a quickly obsolete AI system. It may well even have to start over with a new vendor.

There is no easy solution to this dilemma. And it's easy to say that firms must move forward with careful consideration and consultation, exercising the due diligence they often advise for their own clients. It's more difficult to know how best to do that and what choices to make.

Another issue raised by AI is the effect on young attorneys. Many of the tasks traditionally performed by junior associates—such as document review and legal research—are rapidly being displaced by AI systems. What then do young associates do, given that it is important to keep hiring entry-level associates to maintain a pipeline of attorneys moving up the experience ladder?

While it may seem tempting to just avoid the problem and hire only mid-level associates, that's clearly not sustainable. Not all firms can hire just third- or fourth-year associates. Every lawyer has a first and second year before becoming a third-year attorney.

Moreover, how will this affect the leverage models used in many mid- and large-sized law firms? And how will law firms bill clients for AI systems that can do in five minutes what it used to take an associate two weeks of billable time to complete?

Law firms will need to adjust their normal work assignments to give entry-level associates new types of meaningful work and maintain a healthy pipeline of developing attorneys. It is also incumbent on law schools to change their training models to produce graduates who can function and thrive in the increasingly technology-driven legal profession.

AI also will present novel malpractice risks. AI systems will make many mistakes in their early stages, both as a technology and as adopted as specific programs by law firms. Thus, when can a lawyer sufficiently and reasonably rely on an AI system's analysis and recommendations?

Even when an AI system has reached some level of maturity, it still will be important for human lawyers to check on its output. For now and for the foreseeable future, AI systems will lack the common sense of humans and may miss obvious details.

Problematically, many lawyers lack the technological experience and expertise to determine if and when AI output is reliable. Even technologically sophisticated attorneys may be challenged, given that deep learning algorithms are often “black boxes” that are not transparent in their decision making.

Yet, attorneys will nonetheless be responsible for an AI system’s mistakes, especially if the reliance prejudices the client’s litigation outcome. In the initial stages, therefore, premature reliance on AI systems may increase malpractice risk.

As AI systems continue to learn, though, they will become more effective and accurate than human lawyers. That already has been demonstrated in narrow tasks related to legal research and document review. For example, a recent study by an AI firm found that 83 percent of judges reported lawyers had failed to cite important cases in their briefs. The implication of the study was that an AI system would not have missed those important legal precedents and would have surpassed the abilities of human lawyers to find relevant case law.

At some point, it may be malpractice *not* to rely on the AI system and instead use only antiquated human judgment. When that tipping point will come is difficult to predict, especially given the enormous speed at which AI is developing.

AI also will present additional legal ethics issues. Given that AI systems must be trained on data, who has the rights to use a system trained on a client’s data? If a law firm contracts with a vendor to use an AI program for a major client representation and trains the AI system on a client’s data, what happens when the case is over? Can the law firm use the AI system trained on the first client’s data with a second client? Can the client use the trained AI system in a subsequent matter when represented by a different law firm?

Does the vendor have the right to use its now-trained system with other law firms? What if an attorney representing a client using an AI system decides to move to another law firm? If the client moves with that attorney to the new firm, can the attorney take the AI system that the original law firm trained for that client to the new firm?

There are, of course, many possible permutations of these issues, raising tricky questions of legal ethics and duties.

Conclusions

So, bottom line, there is both good news and bad news about the role of AI in the practice of law and litigation. On the positive side, AI is not going to replace lawyers, at least not completely. The adoption of AI into the practice of law will be evolutionary, not revolutionary. AI will displace primarily the more routine and tedious aspects of practice, leaving human lawyers to do the more skillful and professional parts of the job, such as advising clients, developing litigation strategies, negotiating with opposing counsel and parties, brief writing, and oral advocacy.

Attorneys have the opportunity to improve their craft by integrating AI into their practice, helping lawyers to become more efficient, more accurate, and more persuasive. In turn, it may be feasible to represent more clients at lower costs, helping to address the serious access-to-justice problem in the United States. And attorneys who are early adopters have the opportunity to become leaders and pathbreakers in the digital future practice of law.

At some point, it may be malpractice not to rely on the AI system and instead use only antiquated human judgment.

On the negative side, billable attorney hours are being lost to AI. That trend will only increase in the future. It will affect young attorneys disproportionately, creating challenges in how to incorporate entry-level attorneys into the practice. AI also will present important challenges to law firm business models, billing practices, malpractice liability, and legal ethics.

It will require changes in legal training. AI demands knowledge and abilities outside the existing skill set of most attorneys. Lawyers and law firms that fail to incorporate AI will quickly be left behind. And incorporating AI into the legal practice will soon be a matter of keeping up rather than being a leader.

The most important take-home lesson is that the successful lawyers of the future will not be humans acting alone *or* machines acting alone, but humans *and* machines working together.

The human/machine combination is proving superior over humans alone or machines alone in many different fields. Just as AI can beat the best human chess player, a human/AI team can beat both the best human chess player and the best AI chess player. Similar instances of the superiority of the human/machine combination have been seen in research, medicine, business, and the arts.

The same will hold true for law. The successful lawyer of the future will be a human attorney working closely with AI. ■