

Artificial Intelligence and the Future of Law: Musings Part 1

Bill Novomisle, Founder and Chief Design Officer at In-Gear Legalytics

One of the best aspects of being a consultant working with emerging technologies to service the legal world is that I get to have numerous conversations with people across the world that are engaged in the practice of law as well as the development of cutting-edge technologies. Perhaps it is because of all the recent press that [Watson](#) has been receiving, and specifically the [ROSS](#) application, one question that has come up repeatedly from the practitioner side is what the future holds for lawyers given the advances in artificial intelligence (AI).

I am not [Nostradamus](#), but it is certainly a lot of fun to take a bit of time and think about what the future may hold for the world of legal practice. What follows are my musings on the impact that AI is having and will continue to have on the practice of law around the world. I have divided these musings into two parts. Here in part one I focus on the state of AI as a technology as well as the specific ways that AI has been deployed in legal services application thus far. This is not the first attempt to take a survey of this landscape and make predictions about the shape of the future. For an excellent, and even more extensive review, please see [Michael Mills' reviews](#).

The Current State of AI Generally

The number of articles on AI is staggering and the advances it is making in fields and industries across the board is equally staggering. As just a brief sampling of developments from June and the first week of July 2016 (!), here are few that caught my eye:

- We already knew that some weather reports and sports reporting were being written by AI, but the Associated Press has now [partnered](#) with Automated Insights to use computers to automatically generate stories covering minor league (USA) baseball games. American [college sports](#)

will likely be next, and the technology is sophisticated enough to present [anecdotes](#) about the event so that each story “reads human.”

- Google, unsatisfied with merely incorporating AI into its Translate, Photos and Gmail products, has now used AI to generate both [visual](#) and [musical](#) art.
- My *alma mater*, [Dartmouth College](#) held a [competition](#) to see if a panel of judges could differentiate between computer-generated and human generated sonnets, short stories and music. While the computers didn’t fool anyone, the judges all agreed that the quality of the submissions suggested that it was very probable that AI would be able to produce art indistinguishable from human-generated art in the foreseeable future.
- While traditional AI has focused on “deep learning” which relies on analyzing massive amounts of data (e.g. Google’s [AlphaGo](#), or IBM’s Watson), a new model has been developed for [analogy-based decision making](#), which allows for AI-based decision making on limited data sets, and allows the computer to reason from past precedent (sounding familiar to any lawyers out there?). This has already been applied to visual problem solving, college-level physics problems, and moral decision-making. And down the road, it could be adopted to account for [cultural differences](#) in values.
- The line between understanding the human mind and programming computers continues to blur as [recent breakthroughs in AI](#) are driven by advancing theories of information processing rooted in cognitive sciences.

That is just what a small sampling of headlines and press releases contained over a five week span in the summer of 2016. As with all scientific research, there will be steps forward and steps back, and there will be overstated claims and there will be under-acknowledged breakthroughs. But the direction is clear enough – AI will impact all of us, in ways both seen and hidden in plain sight.

The Current State of AI in Legal Practice

It is well beyond the scope of these musings to try to compile a comprehensive list of all the activity of AI in the legal space. Here is just a short review of prominent recent developments.

More and more big names from Big Law and business are investing in AI technology, whether it be AI or slightly less-advanced expert systems. Just off the top of my head, I have seen [Kira Systems](#) recently sign pretty high profile deals with [Clifford Chance LLP](#) in document and contract review as well as with [Deloitte](#) for similar uses. Likewise, Allen & Overy LLP (in coordination with Deloitte) has rolled out [MarginMatrix](#), utilizing a type of computer-based expert system to automate application of complex rules for over-the-counter derivatives trading and to eliminate the bottle-neck and errors of human intervention.

Berwin Leighton Paisner (BLP) has used the RAVN artificial intelligence system in its real estate disputes practice to [conduct document and contract review](#). (And was recognized by the British Legal Awards and the BBC for its efforts).

In a different application of AI to legal-related issues, Yahoo! has [used the AI platform Brightflag](#) to review all of its detailed hourly billing invoices from law firms to enforce its billing guidelines and to benchmark its legal spending across law firms and regions.

ROSS, billed as “the world’s first artificially intelligent lawyer” has been [hired](#) by Baker & Hostetler LLP to work with its bankruptcy practice. A few weeks after that announcement, ROSS [revealed](#) that Latham & Watkins LLP and a 150-lawyer Wisconsin-based firm was also leveraging the platform to assist with legal research.

It is entirely safe to say that AI is not currently widely adopted for use in corporate¹ legal services. Nevertheless, the examples above show that it is

¹ AI and law are intersecting outside the realm of corporates as well. For example, take a look at what [DoNotPay](#) is doing to parking tickets in London and New York City.

gaining attention, gaining acceptance, and the technology is steadily improving from a novelty to a routine practice tool – at least in the North American and UK markets.

What is AI, Anyway? How is It Being Used Today?

Before we go any further, a working definition of AI is in order. AI is a machine doing a [cognitive task](#) that we previously thought only a human could do. AI analyzes words, pictures and numbers and is able to extract “understanding” or “meaning” out of this data with either no, or limited, intervention from a human operator. We already expect machines to do lots of tasks that a human could never do (ever try to run at 100 mph?) but there were always tasks that we believed only a person could perform and many of these tasks revolve around our ability to comprehend and synthesize written language. While the application of AI to the law has only begun, one of the key skills of a lawyer is to comprehend and synthesize a specialized form of written language and thus AI developments squarely impact the legal profession.

The key forms of AI that are currently affecting the legal profession include [natural language processing](#), [supervised learning](#), and [anomaly detection](#). To understand what this means for a lawyer practicing today, we need to start with the question of what AI can actually do right now.

Broadly speaking, AI in the law currently performs two tasks in law firms (corporate law departments will be addressed shortly). One is legal research. AI has now reached the point where a user can input question using natural language into a system such as ROSS and the machine can then search all potentially informative sources of information to assemble an answer as well as a confidence rating in the answer it has produced. The speed with which a machine can search and process regulations, statutes, case law, scholarly articles, and news reports simply cannot be duplicated by a human being. As of right now, this technology has primarily been applied in the US, leveraging the massive amounts of published regulations, statutes, judicial opinions, commentary and scholarly work.

The second task AI has been directed toward in law firms is analyzing and extracting relevant information from unstructured data. This is a time-consuming (and often tedious) task that comprises much of the work of due diligence, contract review and document review. AI can analyze massive amounts of data and extract relevant information, even if the language used for the relevant information is not identical in all of the documents. For example, AI can find all contract provisions that affect a party's rights to assign or novate a contract, even if the exact wording used is substantially distinct across the set of contracts reviewed. This is analogous to what the RAVN system does for BLP with quite a bit of success. This is also what [technology-assisted review and predictive coding](#) is being used for in e-discovery on a regular basis in the US.

Turning to the corporate side of the legal services value chain, there are a couple of additional applications of AI within corporations and their in-house legal departments. First, there is a growing but still nascent area of AI that is being used to practice preventative law. This application of AI analyzes massive amounts of unstructured data, similar to the way AI is used for document review by law firms, but by tying together data from across the organization and even from outside it, the tool (such as [NexLP](#)) can proactively detect potential instances of fraud or even predictive intelligence for the business. Because of the speed of AI and the continuous operation of the system, threats and opportunities can be identified and flagged in near real time to a human operator who then decides what actions are appropriate based on the information culled by the machine.

Finally, there is the use of AI (or some tools that are technically not AI) to predict future outcomes (which can be useful to corporate legal departments, outside counsel, and even [others](#)). [Lex Machina](#), which was [acquired](#) by Lexis Nexis in November 2015, uses data analytics to predict outcomes in US-based intellectual property litigation, and is expanding into other areas. [LexPredict](#), founded by Professor Daniel Martin Katz (Illinois Institute of Technology – Chicago Kent College of Law), uses a combination of experts, crowd-sourcing

and technology to predict outcomes of Supreme Court cases. These tools offer two interrelated promises for lawyers and their clients. First, they offer a tool that may allow for deeper examination of which arguments and legal strategies may be most effective for a litigant, even offering confidence intervals on the prediction made by the tool. This information can then be synthesized with the practitioner's own views and experience to guide litigation strategy. Second, these tools offer a chance to anticipate significant legal outcomes and proactively prepare for the future, even before it arrives. However, none of these tools is perfect, and both require sophisticated human operation to be deployed effectively.

What are the Current Limitations of AI in Legal Practice?

AI works for legal research in jurisdictions where there is a lot of certainty surrounding the state of the law. In the USA, Canada, UK, Australia and much of continental Europe, law has developed into a well-defined and well-structured set of rules and tests. This also applies to jurisdictions such as Hong Kong and Singapore, which have been heavily influenced by the common law traditions with transparency built in as part of the socio-economic environment. While the sources of the law may be diffuse, the actual legal regime tends to be clear and generally either boils down to a rule (or complex patchwork of many rules), or a standard (or a combination thereof). In the event of a rule, the user is certain that X is [or is not] permissible. In the event of a standard, the user can determine what factors will be considered by an adjudicator and then see how those factors have been balanced against each other. Because these jurisdictions favor predictability and consistency, lawyers, and now AI, can mine a given jurisdiction's applicable sources of information and legal authority to make reasonably confident predictions about outcomes.

In contrast, much of the world outside of those jurisdictions does not have nearly that level of certainty in the law and practice. No jurisdiction has absolute certainty in its laws, so this comparison is necessarily relative. In the US, judicial precedent, common law traditions and lobbying from business interests has resulted in a set of interconnected statutes and regulations that

are generally unambiguous. The relationship between laws and certainty is comparatively looser in other jurisdictions. For example, China unsurprisingly has a law that prohibits “state secrets” from leaving the country. While a US lawyer may expect that the definition of “state secret” would be well-defined, it is [purposely left vague](#) so that even experienced practitioners are unable to define the term. There is also a general absence of official secondary materials to rely upon for guidance on how certain legal terms will be interpreted or enforced. In situation such as these AI doesn’t offer much of an improvement over a basic internet search because the corpus of relevant authorities is limited to the statute itself.

[Sidebar: This statement is not entirely true, as close China-watchers will assuredly recognize. Government and party officials regularly make speeches or give other public announcements that are regularly scrutinized to gain clues into how the laws will be applied or enforced. Nevertheless, these pronouncements are never binding, can easily be reversed, and are beyond the scope of what AI – at least presently – can interpret and apply in a meaningful way.]

AI also works exceedingly well when you have a massive pile of unstructured data and the client needs to extract certain well-defined information out of that data. This need occurs while conducting due diligence on behalf of a potential purchaser in an acquisition. As part of the diligence, a mountain of contracts needs to be reviewed for so-called “change of control” provisions. While the term may be phrased in any number of ways the point is simply that the potential purchaser of this company wants to know which contracts may be impacted/jeopardized by the act of acquisition, and how important those impacted contracts are to the target business. This is a task far better suited to AI than a human. And given that natural language processing algorithms now have been developed for most major languages of business (English, Chinese, Korean, Japanese, Hindi, French, German), this is an area ripe for the application of AI well beyond typical geographic borders.

At present, AI is getting more press than adoption in legal services. However, similar to predictive coding five years ago, the technology is both improving and gaining acceptance as a tool to enhance and improve the skills of lawyers. The critical point however is that AI as it has been applied to legal work is just a tool that still requires a skilled human operator to be effective. Areas where AI has operated entirely autonomously has yet to touch upon legal services. In part two of these musing, I will present my best guess as to how this tool will impact the practice of law in the medium term and in the long term.