ABA Commission on the Future of Legal Services
21st Century Technology vs. 19th Century Law Practice

Michael Mills — Neota Logic

February 27, 2015

Thank you for the invitation to speak to the Commission. My topic today is the role of software in legal services. I am going to attempt to answer the question: does software practice law?

Why does this question matter?

Why do we care about this stuff? Why is the role of software in law important? Let's have a look at the context in which we operate. If you hang around the legal technology industry, if you go to the ABA Tech Show or ILTA or LegalTech, you might think that this silvery fellow on the screen is the modern lawyer.
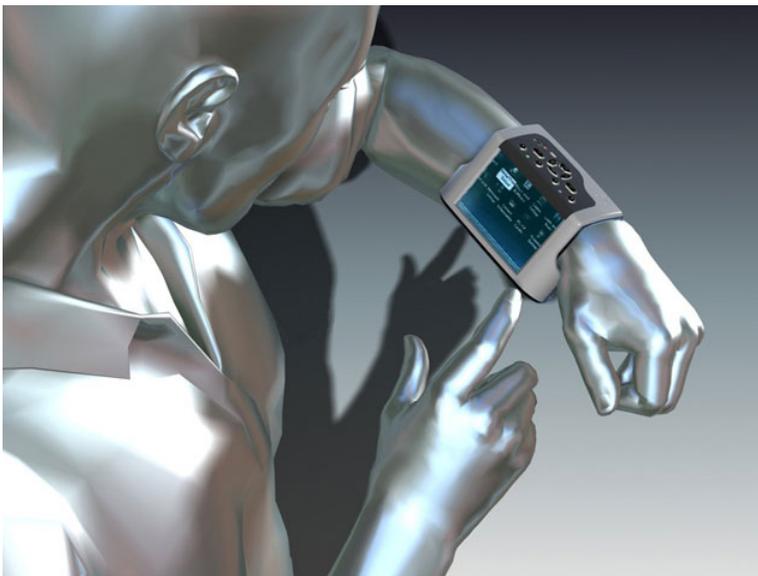
The profession now embraces email and the internet. It uses software rather than typewriters, reads email on smartphones, and even talks to clients by way of Skype.

We are not a fast moving crowd in this profession. In fact, apart from a rather small collection of basic software tools, the profession—as a whole, on average-is more like this fellow.

What is the effect of this old fellow's approach to technology—telephone and Dictaphone at hand? This is a simple economic analysis of the legal services market.
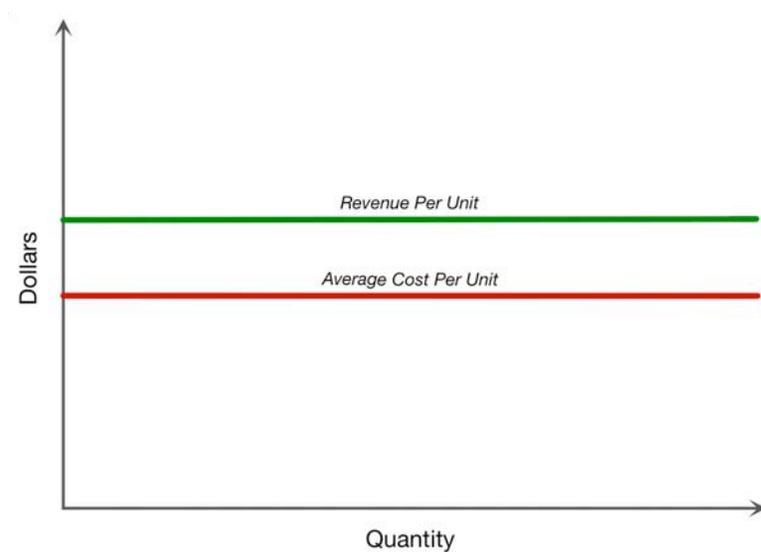

The top line is revenue. The top line is pressed downward by cost-conscious clients and the bottom line, the average cost per unit, is nudged

downward by lawyers trying to preserve margins. Law firms cut support staff and space. Alternative legal services providers hire cheaper labor and put them in cheaper space. All that is better, but it is not good enough. There is something still wrong with the way in which legal services are delivered. The problem is in the straight line.

The cost per unit of output goes down when you manage things effectively with project management, multi-sourcing, and so on. An hour costs $400 instead of 600 and that hour produces two tenths of a contract instead of one. But when quantity rises, total costs rise at exactly the same rate. Every unit of output costs the same as the last one.

This applies to not-for-profits too, if you think of revenue as funding. The lines are closer together and the squeeze is much tighter.

Those lines, those straight lines, are not how Microsoft did it, not how they became the big successful company they are. The marginal cost of delivering another copy of Windows is close to zero so that as you increase the number you sell the average unit cost declines. The bigger the quantity, the better the profit margin. Economies of scale, as Henry Ford taught.
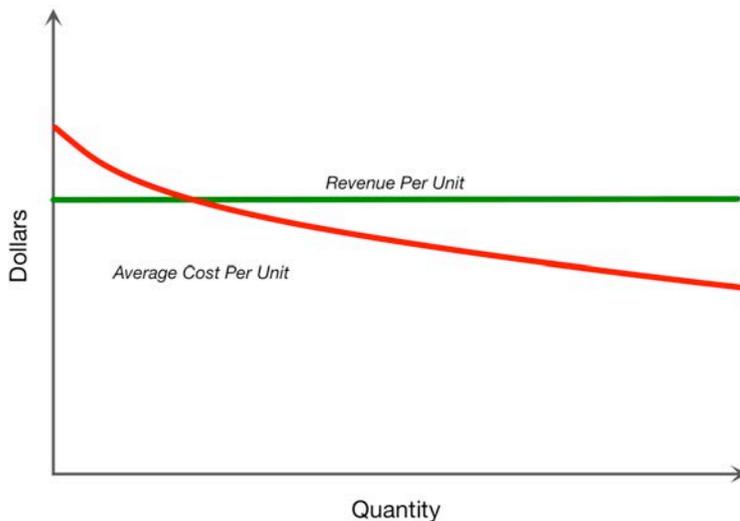
Alas, quantity in legal services is not necessarily a good thing. We have the diseconomies of scale—internal coordination costs, quality variations—but not enough of the economies. In short, legal services missed the Industrial Revolution.

The problem is constant cost. How do we fix that? We use software.

The Computing Context

For many years in computer science, artificial intelligence ("AI") was not altogether jokingly defined as "just the stuff we do not know how to do yet." Thanks to amazing hardware power and some very smart minds, AI is very much for real and very much more ubiquitous than most of us realize.

Indeed, AI is moving so rapidly and becoming so capable that some of the leading minds in science, people like Stephen Hawking and Bill Gates, and Elon Musk, have said that we ought to begin thinking about the social and even moral risks presented by ever-smarter software. That is a longish term problem and beyond the ABA's remit to solve, but perhaps it is now close enough that questions should be framed and answers essayed.

Software runs on hardware, of course, and the astoundingly rapid evolution of hardware has made possible the powerful software we now have that I am going talk about in a moment. This is IBM's latest chip, the TrueNorth.

It has 5.4 billion, b for billion, transistors. It has a design inspired by the architecture of the brain, with "neurons" and 256 million "synapses"— a very powerful device.

Thanks to chips like TrueNorth, computers are growing smarter and smarter year by year. This chart by Ray Kurzweil, a famous computer scientist, argues that computers today are about as powerful as a mouse brain.
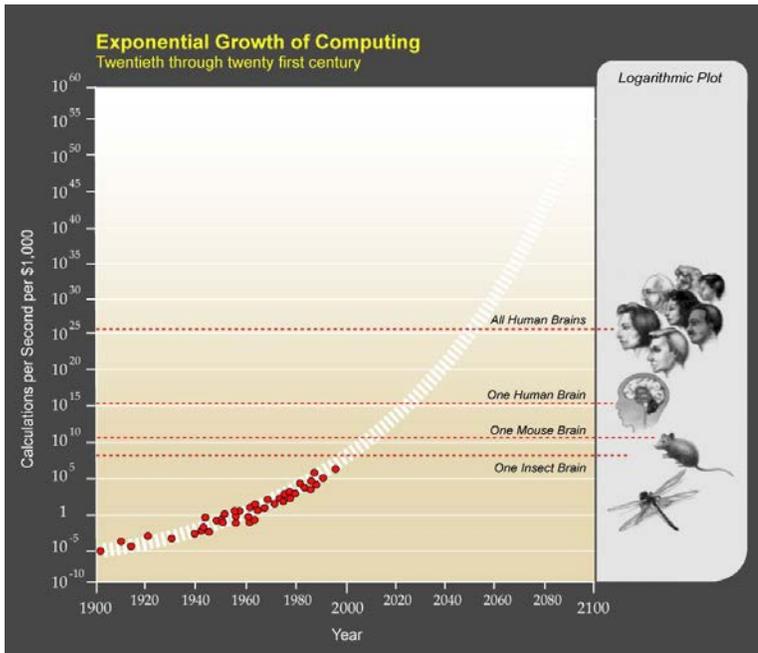
That's enough power for us to do some very useful things in law, which is a remarkably digital product. It is lightweight, malleable, portable, and computable.

Keep in mind that whatever software can do today, more can be done tomorrow and next week and next month. You have all heard about IBM Watson. This is a chart that IBM presented last year of the improvement in the power of Watson over the course of just three years.
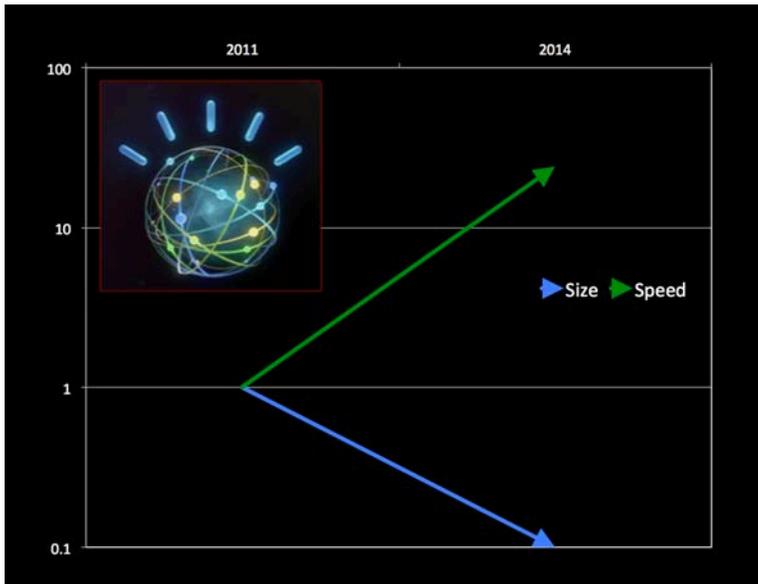
An Artificial Intelligence Taxonomy

The world of artificial intelligence is a big one. It covers in computer science a great many interconnected species.
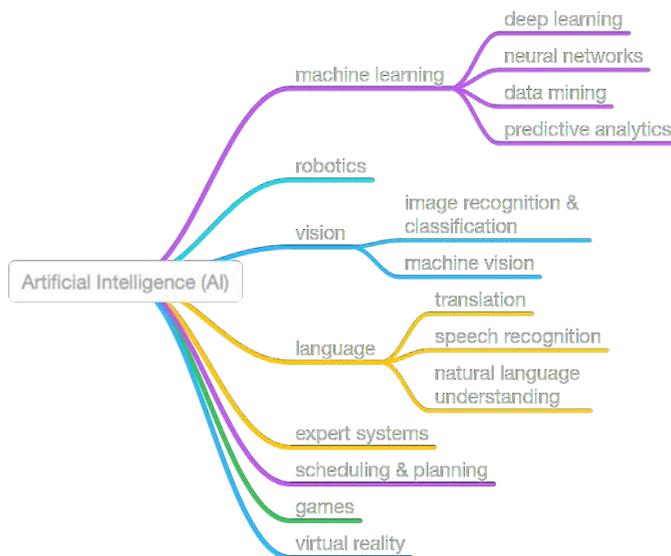
Natural language understanding, games, and maybe even virtual reality will have a role in law. Today I want to focus on just two of the branches of artificial intelligence, machine learning and expert systems.
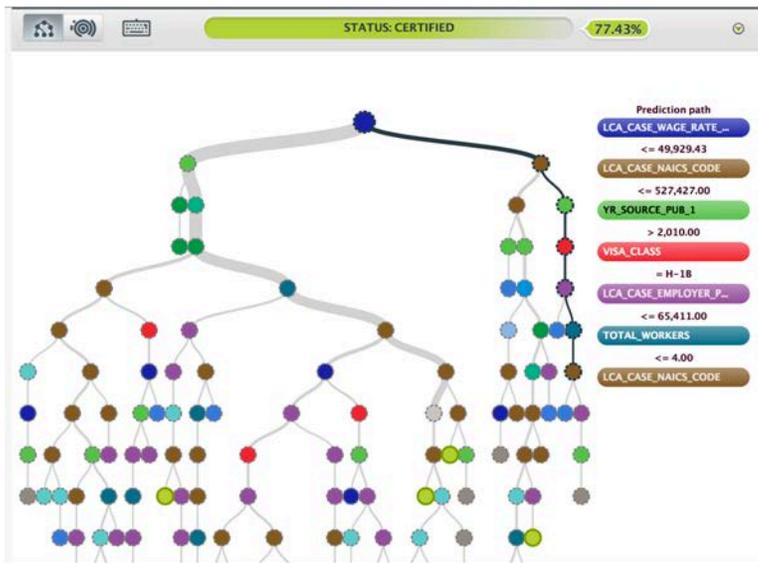
Machine Learning

Machine learning is the most prevalent form of AI in the law today. We need to know about it because the profession is using it today and every day and increasingly so. Machine learning is a set of algorithms. There are many of them, principally designed to cluster, classify, and identify objects and concepts in large data sets.

Machine learning is being used most pervasively in the legal profession to tackle the classification, extraction, and understanding tasks of electronic discovery. It is so good that a number of careful comparative studies have demonstrated that "technology assisted review" is both more effective and more efficient than having humans plow through the documents. See, e.g., M. Grossman, G. Cormack, XVII Richmond J. Law & Technology 11 (2011).

Machine learning should have become almost universal in discovery, yet has progressed slowly in the face of surprising indifference by lawyers and even their cost-conscious clients. We are indeed a slow-moving crowd.

The second major category of machine learning is predictive analytics, for which a shorthand definition is this: apply algorithms against large data sets of inputs and outputs, find patterns, induce associations among data points (call them rules), apply the rules to new inputs to predict outputs. From Amazon to Wal-Mart, every retailer uses predictive analytics to maximize the odds of your buying something, often something you weren't thinking of buying.

This image illustrates a machine learning analysis of the characteristics and outcomes of a set of wage and hour cases, from which the software built a model that can be used to predict outcomes of future cases.

In the commercial world, you probably know Lex Machina, which is using machine learning techniques to predict the outcomes of patent and other intellectual property litigation. The company had to build a dataset before beginning analysis because, in one of the tragedies of the legal industry, far too much data is unavailable to the public and unavailable in a machine-readable form. Once they did that, they had the power to predict outcomes with commercially useful accuracy and have built a company around that. Dan Katz at Michigan State has done a lot of work in predictive analytics as well, and there are a number of other startups who are doing that sort of thing.

Another example of machine learning in the law is a pair of companies, eBrevia and Kira Systems (Diligence Engine), which have built software that reads and understands contracts. I say "understand" in a limited sense, of course, but it is a powerful and practically useful sense so that during due diligence rather than plow through acres of contracts with young lawyers trying to abstract and understand them, the software does the first cut and does it very effectively and very accurately.

And then there is IBM Watson, which as you all know won the Jeopardy game three years ago and then moved on to do useful work in the real world. The most-cited is a wonderful project at Memorial Sloan-Kettering in which oncologists are working with Watson to improve their ability to diagnose and to plan effective treatments for certain types of cancer.

Watson is nudging into the law too. There are a couple of below-the-radar projects but one has been quite visible recently. IBM conducted an open contest and some students at the University of Toronto came up with an approach to using Watson to do legal research, named ROSS.

The general counsel of IBM said last summer that Watson could pass the multistate bar exam without difficulty. Powerful though Watson is, its ability to pass a bar exam may say more about the exam than about Watson.

The Commission heard recently from Mark Britton, the founder of Avvo, who told you that visitors to the Avvo website have posed 6,500,000 questions about their rights and obligations and strategies for dealing with problems that they perceive as legal problems. Those questions have been answered by lawyers, who participate in Avvo as a form of advertising because it may lead to conventional legal work, just as big law firms publish reams of new developments in newsletters and blogs.

Given access to the right source materials, Watson could answer at least half of those questions as well as or better than the lawyers. Not because the Avvonians are not good lawyers. They are. But because Watson knows more about more topics in a more persistent and pervasive way and can make more comprehensive associations among the various materials that are available to it.

Who owns the law?

"Access to the source materials" is a critical issue here. I do not know whether it is within the remit of the Commission, but it does bear on the ability of technology to improve the future of legal services. There are many bodies of legal materials that are walled off, because public systems are old and budget-starved and in some instances because commercial relationships take the public materials of our practice and turn them into exclusive products.

Even when records are available, they are often not available in the digital form necessary to be computationally productive. PDF images, for example, are not nearly as useful as xml data. Removing the access and format barriers will require public policy changes.

Imagine if a powerful, publicly available search and reasoning tool had access to primary materials (statutes and regulations) and also to briefs and records and decisions in all of the trial and appellate courts and administrative agencies, state and federal. IBM Watson would be doing far more than winning points for guessing "Who was Bram Stoker?"

Google is of course in this game too—indeed literally in the game. Their subsidiary DeepMind announced recently that their deep learning multi-layer neural network software taught itself how to play four dozen computer games from the old Atari 2600, an early game machine. You may think that computer games are indeed just games, but they are actually analytically very hard problems. This software started from zero knowledge of these games, watched them on the screen, played at random, and over a period of time taught itself so well that it beats humans and always wins the game.

The rules of law are much harder than the rules of Breakout and other early Atari 2600 games. The fellows of DeepMind have no illusions about how hard real-world problems are. The rules of law are especially difficult, not only because they reflect real life rather than virtual life, but also because they are made by human political processes and are therefore both accidentally and sometimes intentionally incomplete and inconsistent and obfuscatory.

Look again at that graph of Watson's power growth over just three years. As that moves forward, the challenges presented by the game of law will become much more within the reach of computational power. The DeepMind paper published in *Nature* reports an extraordinary advance in the ability of autonomous software to figure out on its own the rules of games and play them effectively. There are portions of the law in which one could imagine that happening relatively early. Obviously not in the broad spectrum of the law, but in narrow areas of the law.

Expert Systems

Let me turn now to the second major branch of artificial intelligence that is most relevant to law, expert systems. Machine learning is a marvel in all of its manifestations—from reading email (truly one of the most boring tasks that young lawyers are given) to helping oncologists at Sloan

Kettering. However, in some contexts machine learning is not suitable or sufficient. It generally produces probabilistic answers and the algorithms are dark and mysterious and incredibly complex. This is a fragment of one algorithm:
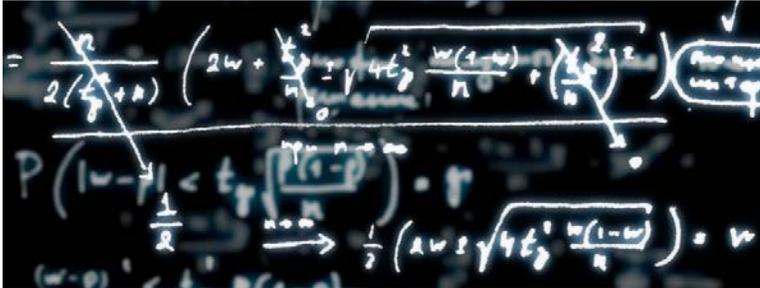
In the legal context, people often want to know *precisely* what they should do *now* about a *specific* problem in these circumstances. Lawyers are generally determinists and they are certainly skeptics. Clients are result-driven and they are very busy. They want answers. They do not want a list of search results, they do not want a stack of pages that they need to read through in order to figure out the answer.

Expert systems are the branch of AI that deals with explicit and transparent systems. It is a way of building applications, and a formal definition is useful:

**software** that
transforms
experts' **knowledge**
into **practical answers**
that
other people can use

Expert systems are software with which you can acquire and organize expertise from one or more experts—that is, people who know a lot about a subject. Then the software can be used to deliver that expertise to other people, tens of them, or thousands of them, or tens of thousands who

know less about the domain but need a solution to a specific problem within the domain.

This is a way of looking at the law which is different from the way lawyers are traditionally trained. Lawyers see the law as text, all those statutes, regulations, decisions, books, and memos that we write and read. But expert systems see the law as a set of rules and relationships.

There are many ways to create them. The important thing to understand from the Commission's perspective is that expert systems, like machine learning, are everywhere today. As an example they are in the pocket of your doctor's white lab coat. This diagnostic system that runs on an iPhone is used by thousands of doctors.
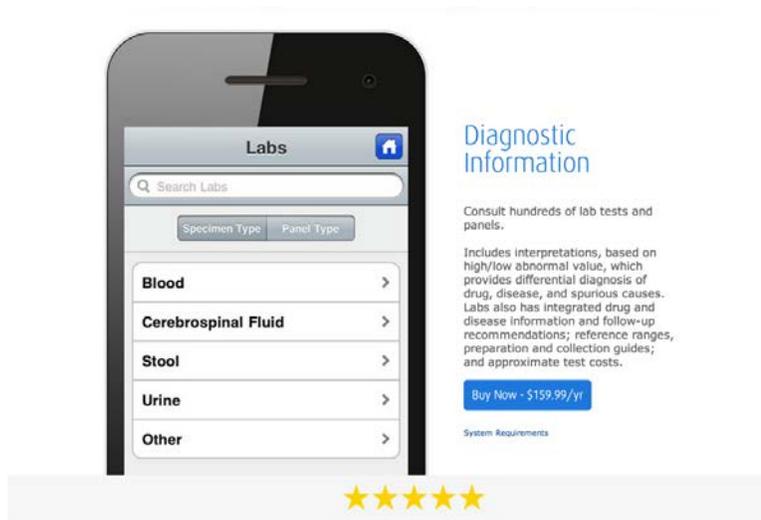
There are many similar medical expert systems.

The IRS spent more than $10 million to build a suite of expert systems called the Interactive Tax Assistant. The original purpose was to help the IRS folks who give telephone advice to give better answers. Then the IRS said, "Wait a minute. Why can't we just open up these systems to taxpayers?" So, if you are a taxpayer with a question about one of the areas in which they have built the expert system you can go to the IRS website and get very precise fact-specific, context-specific guided answers to complex tax questions.

IRS

Am I Eligible To Claim My Charitable Contribution?

This application will help you determine if your charitable contributions are deductible.

Please click the **Continue** button to proceed.

[Continue]

Think of TurboTax. TurboTax is an expert system by the definition that I just gave. It is one of the most complex expert systems ever built because it covers 50 state tax laws as well as federal law. It can provide not only the calculations that allow you to fill out your tax return but also very precise guidance on how to minimize, in a legal sense, the tax obligations you have. TurboTax is built by an army of programmers and has thousands, tens of thousands of rules, in it.

There are expert systems built for self-represented litigants. This is one that is built with my company's software. Think of this tool shown on the screen as one of the elements you would pull out when unbundling legal services. This is for a self-represented litigant who needs guidance on how to handle an appeal.

Instead of handing that person a long memo and a piece of paper that says here's what you do, you instead give them access on a court kiosk, a social service agency's kiosk, or on their phone to an expert system that asks a series of questions about the problem they have, and gives guidance how to handle their specific appeal on an issue that comes up every day for tens of thousands of people across the country and which could never be lawyered cost-effectively. There are not enough pro bono hours on the planet to meet the demand for useful legal advice from people who need it and cannot afford it.

Law Schools & Technology

This application was built by students of Georgetown Law School. I want to take a moment to talk about law schools because in thinking about the future of legal services you need to be thinking about the future of legal education. Are law schools teaching what students are going to need to know in five years or ten?

Richard Granat, who has submitted testimony to you, and Marc Lauritsen, a well-known expert on software and the law, several years ago compiled a list of the law schools that were running significant technology

programs to bring to bear or bring to the attention of and allow students to experiment with the technologies that are relevant to legal services.

There are 205 law schools on the ABA approved list. As far as we can tell, there are fewer than 25 of the 205 who have any consistent education for law students about technologies that are relevant to their practice.

This means that five years from now, when they are young lawyers, they are not going to understand the intersection between law and technology that will be necessary to deliver legal services in a cost-efficient way. It strikes me that one of the things this Commission can do is work with the other committees of the Association focused on the content of legal education, because the lack of education in technology is in my view a form of legal illiteracy.

I mentioned the law schools in the context of expert systems also because these students are building applications that actually get used by not-for-profit organizations. Ron Staudt, a member of the Commission, was pioneered Apps for Justice at Chicago, which has spread to other law schools.

Another form of expert system originally came out of an ABA-funded project by Jim Sprowl many, many years ago to build a document automation tool that in some sense was the predecessor to HotDocs, one of the pioneers and still one of the leaders in document assembly.

Document automation is a form of expert system. Technologists may delve the details of inference engines but in fact what is embedded into a document automation system is the expertise of true experts, people who understand how transactions or court situations or other events with consequence to clients should be structured and documented. There is tremendous legal expertise that goes into building these things and delivering them well. They provide through a completely interactive interface with no human intervention very valuable legal advice.

A document automation service that the Legal Services Corporation has funded is LawHelp Interactive, a project of Pro Bono Net. It is used in about 40 states—more than 3,000 different document forms for many different problems that people face when they are self-represented. Last year 500,000 people used Law Help Interactive. That is a huge legal

activity in and of itself. It is the delivery of legal services in a completely automated way using a simple form of expert system, document automation.

The Regulatory Context

There are many people who are very experienced in professional ethics and regulation who have spoken to this Commission. I am not one of those people. You have heard from Charles Rampenthal at LegalZoom, Richard Granat, Ken Grady, Slater and Gordon, and a number of others.

As a person who builds software that can be used to improve the delivery of legal services, I share the views of those who have presented to you that the time has come for what Mr. Rampenthal calls right regulation.

The fact is that Australia's citizens have been well served by that country's reforms. The legal system of the United Kingdom has not collapsed since the Legal Services Act of 2007 was adopted. With the new initiative in Canada, no one is forecasting earthquakes or the imminent demise of the legal profession or the collapse of standards in service to clients.

The big issues those folks have spoken about, ownership and capital and new classes of legal problem solvers, those are important questions and they are ripe for new answers. I very strongly believe that, and would expect the Supreme Court's decision in [North Carolina Board of Dental Examiners](#) to remind us of the framework in which exclusionary practices will be judged.

Today, I want to look at one aspect of the issue—the scope of practice.

Prescriptions from Medical Practice

There is an expression that you may be familiar with in the medical field: "working at the top of your license." The medical field is of course a licensed and regulated one. Expertise no doubt counts. You do not want people making mistakes. But the medical profession has been very open to figuring out how to deliver services using the largest available set of resources in a prudent way.

- MD Degree
  - Intern
  - Resident
  - Attending
  - Chief
  - Speciality & sub-specialty

- No MD Degree
  - LPN
  - RN
  - APN
  - Nurse Practitioner
  - Physician's Assistant
  - EMT (4 levels)
  - Pharmacist
  - Midwife
  - Imaging Technician
  - Respiratory Therapist
  - Physical Therapist
  - Occupational Therapist
- Not Human
  - Clinical decision support
  - Drug recommendations & controls
  - Watson

There are of course folks with MD degrees. There is a bit of hierarchy. You go from being an intern on up to a chief of service and there are dozens of specialties, and those are in fact certified themselves, so there is quality control built into that process.

But over on the right hand side of the chart are all those folks without an MD degree who provide useful service in the medical context, probably to every person on this Commission at one time or another. You have dealt with a registered nurse, an advanced professional nurse, a nurse practitioner, you have certainly dealt with technicians of various kinds.

All of those people practicing appropriately within a suitable scope contribute to the delivery of medical services and do so effectively and do so safely. I do not mean to suggest that the medical profession is without errors, obviously it has errors, but there is a structure in place for the effective regulation of delivery of services by folks who do not have MD degrees.

That allows cost-efficient service, that allows patient-efficient service, that allows high quality service by people who are specialized in doing the things that are necessary at the specific point of service. It is a peculiar conceit of the legal profession that in our model there is one category of person called a lawyer. And that that person will do everything from the most menial to the most complex. That is not an economically sensible model. It is not a client-friendly model. It isn't even a lawyer-friendly model.

The medical profession has understood that there are things that are not humans that can contribute materially to the delivery of medical services.

There are clinical decision support systems of dozens of different kinds. In a hospital, every time a drug is prescribed there is an expert system which evaluates whether that prescription is a sensible one in the circumstances based on the lab tests, based on the patient's history, drug interactions, and so on.

In our profession we divide the world into lawyers and "non-lawyers." Indeed the State of Texas decided that firms could not have chief technology officers or chief executive officers and operating officers because that suggested control by the dreaded non- lawyers. In the medical profession, there are people with recognized professional expertise of a particular scope. There is software and hardware with expertise of a particular scope. All of that can be brought to bear in delivery.

Software has a scope of practice too. This conceptual chart lays out the ambit of problems that are faced in legal services. Frequency is the vertical axis, complexity the horizontal.
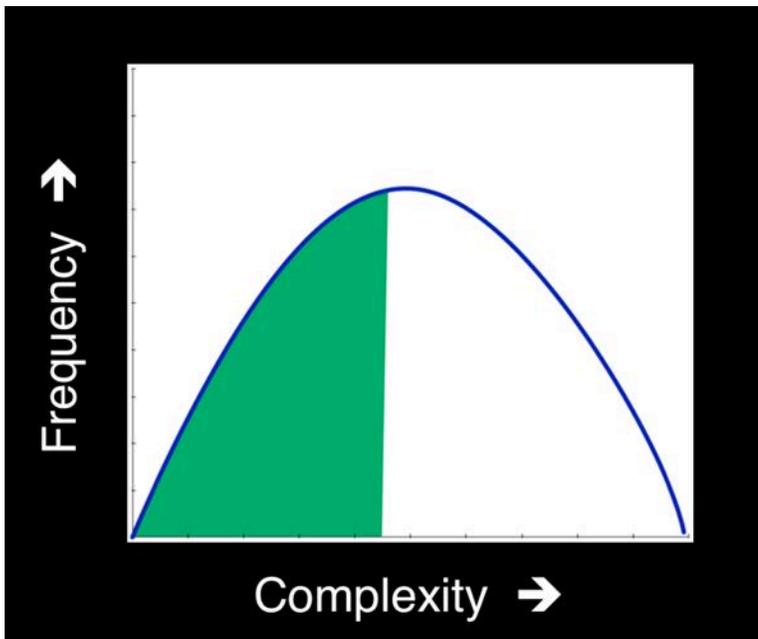
Complexity really means several different things. It can mean the degree of uncertainty about the problem, the dependence on witness veracity, the

need for creative thinking, empathy, subtle judgment and so on, but I shorthanded that as complexity. Most of the world is a bell curve, most problems fall in the middle. There is a tail of very simple problems on the left and extremely complex ones on the right.

Just as the medical profession has done, and the State of Washington and a few other states with additional legal services categories have figured out, there is an appropriate scope of practice. I suggest that software has an appropriate scope as well.

In my view, with the present state of software, it is on the left-hand side of the graph. Software does not handle criminal cases. It does not judge intensity of feeling or veracity; it does not have empathy or intuition. But there is an enormous amount of what lawyers do that lends itself to analysis and service by software.

It is one of the ironies of the unauthorized practice issue that in so many states there is no definition of what constitutes practice; it is a Potter Stewart definition.

There are at least two different approaches to regulating the scope of practice for software.

One is the Texas view: publishing software that is clearly labeled as not a substitute for legal advice is not practicing law. Period.

Marc Lauritsen's illuminating article, Liberty, Justice, and Legal Automata, 88 Chicago-Kent L. Rev. 945 (2013), provides the analytic and constitutional foundation for the Texas approach to speech in the form of software. See also, M. Lauritsen, Are We Free to Code the Law?, 56 Communications of the Ass'n of Computing Machinery 60 (2013).

For the reasons Lauritsen states, we believe that this is the correct approach, both because it is correct on the law and because it matches the profession's duty to serve fairly and effectively all those in need of its services.

An alternative is more akin to what medicine has done with medical devices—to say there is a scope of practice that is appropriate for software, there is a regulatory scheme that is appropriate for it, and that

software is a category that may be regulated like legal assistants or bankruptcy filers or any other new category of legal problem solver that one can imagine. It is subject to prudential rules and can be regulated in that fashion. Andrew Perlman's forthcoming analysis, Towards the Law of Legal Services, 37 Cardozo L. Rev. (2015), suggests this approach. No doubt other "right regulation" approaches may be suggested.

With the driving forward force of improvements in technology, software is a part of the solution to the problem of legal services, and this Commission needs to think about how software is to be accepted into the regulatory scheme.

I agree with the Association's Standing Committee on Technology, which recommended that you look at the professional and ethical implications of artificial intelligence in law. However, I disagree with the Standing Committee that the rollout of these technologies in the provision of legal services would be premature.

In fact that rollout began a while ago and is continuing apace. AI is being used every day in the law, as I have tried to illustrate to you. Technology does not wait for the profession. The robots are at work and they will roll silently past regulators, if the regulators, like Canute, attempt to command the tide to recede.

In summary: artificial intelligence can radically enhance and extend legal services, to the benefit both of the profession and its clients.

Q&A

After the presentation, members of the Commission were invited to ask questions.

Question:

One of the things that you said that really struck me as quite interesting is this point about how information is walled off publicly available information about court cases and the like and it is walled off and that is making legal services more expensive in some sense because it inhibits the ability to develop technology that can analyze that information and perhaps bring down the cost of legal services.

Two related questions. First, can you give us any specific areas where some policy change would be especially helpful, and then two, can you draw that line for us between the inhibition of access to this information and actual lower cost in the delivery of legal services.

Answer:

Let me take the first part of the question about policy recommendations. I am not the expert in this field. Carl Malamud and a group of people who work with him, Ed Walters at Fastcase, and others are true experts, but there is a policy principle that public records should be made available, subject to appropriate privacy and anonymization. But public records should be made available in machine-readable form. There are some emerging standards around legal information, which would facilitate access if appropriate policies were adopted by courts and agencies. In some of the hackathons around the country there have been some quite dramatic examples of using publicly accessible data to do creative things by way of analysis and presentation.

A specific example of the connection between access to data and lower cost of legal services: I'll give one first in the corporate context and then, with a bit more of a stretch because it is less clear, in the consumer context. In the corporate context, the company I referred to, Lex Machina, has poured large amounts of investors' money into building a database of very detailed information about patents and patent litigation. All of that information was in the public domain in the sense that it was filed in

papers in courthouses or in the patent office. But pulling that apart and putting it in an electronically readable form so that it could be analyzed by software was a time-consuming and costly venture.

Once that is done, every company with a patent portfolio can now subscribe to Lex Machina and do very subtle and deep analysis of patent portfolios and the risks of patent litigation. I am not privy to Lex Machina's proprietary data but anecdotally there have been more settlements at earlier stages in litigation. Surely it saves money.

In the public context, I think the policy issue, is not so much a federal one as a state one. The ABA can make recommendations, but I think it is more in the domain of the National Conference of State Courts, for example, that within states there are tremendous inconsistencies from county to county, from courthouse to courthouse, in the extent to which data is available electronically at all and the extent to which it is available in machine-readable form. Uniformity of standards across states would be an enormous policy step forward. It probably isn't really true that every county in California needs its own form to initiate a case since they are, after all, operating under state law.

Those would be two recommendations. One of the consequences in the public service or consumer context is that it is very difficult to create data about outcomes. What Lex Machina did with the patent litigation is to figure out what the outcome of the case was. But in the kinds of cases that consumers face, that data does exist in the courthouses but is not available in any form that we can make use of. You cannot build predictive models unless you have access to data.

Question:

You gave the example of how software or computer intelligence could be regulated. I had asked Chas Rampenthal and I am going to ask you as well, can you think of any useful models either in the United States or around the world where "right regulation" is right in your view, that that balance has been struck appropriately? You mentioned Texas. Are there any other examples?

Answer:

I think he was speaking far more broadly than I about regulation of legal services, and there I plead inexpertise. I am not aware of any jurisdiction—including Australia, Canada, and the UK, which have done the most recent re-thinking about regulation of legal services—who have given any attention to the role software plays and will play in the delivery of legal services and how it should be regulated.

The Texas statute came from a particular circumstance. There was a company that had produced a product and got into trouble with the UPL committee and the legislature solved the problem by modifying the UPL statute. But I do not know of any other states and I do not know of any other countries that have addressed the issue. Yet, I do think it is one that over time will become significantly important.

Question:

To what extent a lawyer should be trained on how to leverage the technology in a particular sense as opposed to really just being informed that it exists so they know where to look? What level of training and understanding do you think the next generation of lawyers really need to have?

Answer:

For the day they go to work, there are some things, very specific things that would be handy for them to know. Suffolk Law School has been working with Casey Flaherty on some of the things that go into that package of immediately useful skills. However, I think the real skill that law schools need to be teaching is not how to use any particular piece of software.

Software changes every three years so learning any particular tool except the ones that you need the day you go to work is probably not necessary. But understanding at some level what software is and what coding is I would argue is essential. Anyone with a general education ought to have done a little computer programming. It can be in the highest level language and it can be the simplest possible program. The notion of piling things up to do things in an analytically rigorous way strikes me as a form

of literacy, and in the law it is particularly so because so much of the law is in fact computable.

Software is what the world speaks and what makes the world work. But I think it is particularly true in the law because we are a rules-driven, rules-based profession. There is an extraordinary synchronicity between law and code that seems to me important for lawyers to understand.

Whether they do that by automating a court form or building an expert system, or writing code or participating in a hackathon, it doesn't matter. They need to understand the concept and get it sufficiently ingrained in their worldview so that as they address legal problems and the problems of practice and the profession over their careers they say, "Aha, this is part of how we think."

Question:

You spent some time talking about Watson and all of the developments. I am wondering if you have any specific predictions or any solutions that are currently in the works that you think are going to have the most dramatic impact on the accessibility of legal services in the future. In other words, draw that line for us. You drew it before with Lex Machina. Do you see particular applications of Watson that you think are going to be especially useful in the years ahead?

Answer:

Watson is very well known. IBM is a big company and Watson won Jeopardy so we all know who Watson is and IBM has encouraged us to know who Watson is, but there are in fact a lot of companies who are doing other very similar things, the Google DeepMind subsidiary for example. There are many smaller ones. There are startup companies like Ravel Law and Judicata and others who are doing remarkable things in presenting and accessing legal information, Fastcase and some of the other legal publishers or legal information companies in fact are doing all kinds of very interesting work themselves.

I think the first benefit is enhanced search capabilities. As that data becomes available and gets fed into those smarter and smarter machines, then it becomes more and more useful and it gets closer and closer to

answering questions in a useful way rather than just presenting a list of search results. All of the folks who sell legal information now are investing in smarter and smart search engines, and there are startups who are doing it and IBM is beginning to do it as well. They are partnering not only with the University of Toronto students but with some others who will bring that technology to bear. But it is not just Watson. There are lots of folks doing that stuff.