EMERGENCE OF CONTRACT STANDARDS AND ITS FUTURE IMPACT ON LEGAL EDUCATION

KINGSLEY MARTIN*

1. LAW AND TECHNOLOGY

Let me start with a story. It was almost thirty years ago that I graduated from Harvard Law School and joined my first law firm. I distinctly remember my first assignment, which was to draft a credit agreement. Despite my education, I had no clue how to perform the assignment. I had to do it the same way as everyone else, which was to go around the corridors, bang on some doors, get some examples, sit down and start to read them. I have no doubt that my first few attempts and, in all likelihood, years' worth of drafts, were sub-standard. This is what, in part, motivated me to focus my career on technology and law.

Earlier this week, I gave a presentation to students at a leading law school and I am sorry to say that todays' law students are no more prepared to perform my first assignment. I asked them three simple questions: "By show of hands, how many of you in the room are interested in a career in law?" One hundred percent. "How many of you are interested in combining some aspect of technology with your legal career?" Zero. "How many of you are worried by the rise in technology as a potential threat to your chosen career?" One hundred percent. It just didn't add up in my mind. Using these questions as a springboard, I first sought to show how technology can better prepare our students. Second, I wanted to address their fears. It is apparent that many students are concerned that technology might take away their jobs. My goal was to explain that technology is an incredible opportunity.

Over the past few decades, the legal profession has narrowed to serve just the one percent. Lawyers working in this space, of course, enjoy a very satisfying and financially rewarding life. The 99 percent is also a massive market opportunity. However, we may not be able to meet the needs of this market in the old fashion way, training lawyers by rote learning. We must harness technology to help us serve the broader market more efficiently.

* Chief Executive Operator and President of KM Standards. Thank you Jennifer for inviting me and thank you for introducing me as a Legal Rebel.

2. THE RISE OF THE MACHINE

In order to take advantage of technology, we must better understand how the machine works. We know technology today is getting ever more powerful. Already, a computer can drive a car. A machine can fly an airplane. It can perform brain surgery. The question I raised is whether a computer can draft a contract? This is what I want to explore with you today. How far can it go? What can it do? How can it help? What are its limits?

On the one hand, we have the rise of the machine, exemplified by IBM Watson. Watson first beat the world's best chess player. It has beaten the best contestants on the TV show, Jeopardy. Now, Watson is learning medicine and is already considered to be one of the best diagnosticians in oncology.

On the other hand, we are trained to question. Lawyers are skeptics. One line of doubt holds the opinion that while a machine can perform incredible computations, it cannot think. A good example of this position is described by Professors Dreyfus and Dreyfus in title, *Mind over Machine*.[1] In the book, the authors assert that a machine can never replicate human though and therefore strong AI (artificial intelligence) will never work.

However, we do not need to reproduce human intelligence in binary form. It matters less how we solve problems. We are more concerned with results. Indeed, if we compare results—who won—what is the difference between intelligence, judgment, and brute force?

3. STAGES OF MACHINE LEARNING

We are working our way through training technology to expand our intellectual capabilities. The evolution mirrors the way that we learn. There are three key stages to machine learning.

Stage one is to find the relevant material. In learning a new area, we need source material, just as I gathered as part of my first legal assignment by asking more experienced lawyers. We need the relevant cases or prior contracts. Today, we are less likely to go to the library or form file. This stage is automated. We use a computer to find the sources. The key technology applied in this stage is search.

1. HUBERT L. DREYFUS & STUART E. DREYFUS, MIND OVER MACHINE: THE POWER OF HUMAN INTUITION AND EXPERTISE IN THE ERA OF THE COMPUTER (1986).

Stage two—once we have gathered the relevant sources—is to review the material and identify the relevant elements: the relevant issues in the case, or the relevant elements of a particular transaction. This is how we learn over time: by reading many examples. And, this is the second stage of automation which applies the emerging technology of analytics.

Stage three is to determine the optimal outcome. What is the best argument to make in a case? What is the best structure of a transaction? This capability, on the bleeding edge of automation, is the focus of "big data" and expert systems.

Let's review some live examples. On the screen is a merger agreement taken from the SEC EDGAR database. It is eighty-five pages long and contains thirty-one thousand words. Our software has read in previously fifty merger agreements as a reference set. The application examines the structural organization of sample documents. It creates a table of contents that works like a checklist. It shows you all the provisions found in merger agreements, how frequently they occur, and full range of standard and deal-specific language. The software can review the eighty-three page contract and report back in just a few seconds; orders of magnitude faster than any human reviewer.

The software automatically creates a table of concordance. The left column displays the contract and compares the terms in the contract to the "reference standard" on the right. The analysis shows which clauses are matched (and the degree of match), which clauses are divergent, and which clauses are missing. The report is color-coded. Let's open up the representations and warranties article. The clauses in black text are matched and conforming. The clauses in red text are in your document, but we don't typically see them in other merger agreements. For example, we do not typically see a "FDA and Regulatory Matters" clause that is specific to the transaction under review. Finally, the clauses in blue text are missing in your document.

How does it work? Let's go behind the scenes. The key technology is described by Dan Katz and others from Michigan State as Quantitative Legal Prediction.[2] It is the application of inductive reasoning. There are two fundamental ways to perform the analysis. We can do it from the top down through deductive reasoning, a la Sherlock Holmes. Or, we can also do it from the bottom up by examining all the building block elements and

2. *See generally* Daniel Martin Katz, *Quantitative Legal Prediction-or-How I Learned to Stop Worrying and Start Preparing for the Data-Driven Future of the Legal Services Industry*, 62 Emory L.J. 909 (2013).

organizing them by identifying the common patterns. Originally, we attempted to solve the intelligence behind chess, or Jeopardy, or legal analysis using deductive reasoning techniques and attempting to define all the logical rules to capture intellect. But to each rule there are limitation and exceptions, and exceptions to the exceptions, and so forth. It cannot be done by deductive logic. It is too complicated. Moreover, while much more work needs to be done in the field of human intelligence, it is not clear to me that humans are logic engines. It is just as likely that we use pattern recognition. Indeed, chess masters do not reason through a series of if-else-then possibilities. They "see" the best moves: the patterns. This technique, which is applied by Google, predictive coding applications, and our technology, identifies and organizes the underlying patterns. Rather than seek to define the rules, we give the computer representative samples, and train the machine to find the matching patterns.

We can use machine learning to examine millions of samples and create contract standards. Rather than viewing the technology as a threat, I see contract analysis as similar to the surgeon's MRI machine. It's a tool to examine large amounts of information and see things that we could not otherwise detect in massive volumes of data. To take a stack of paper—hundreds of documents—and be able to analyze them in minutes and determine what is in those documents, and then be able to compare any other agreement. Applying these techniques, we can educate the machine along the lines of the three stages of learning.

4. CONTRACT ANALYSIS

In the first stage, we find the relevant material. It does not matter whether it is on EDGAR, Google, or in private libraries. With today's search tools, we can quickly find sample material. We do not even need accurate search tools. We can cast our net very broadly, because the next stage automatically filters the source to winnow out irrelevant data. But, many will say: what about the problem of garbage in; garbage out? All systems—both human and technological—draw information from a variety of sources: some good, some bad, some indifferent. Indeed, Watson draws principally from Wikipedia and Google, frequently criticized as being suspect. However, like us, we can build machines with filters. We can take a wide range of inputs and filter out the irrelevant and the potentially inaccurate. We don't have to vet the sources. We can review the entire collection, and the tools will help us find the best results.

In the second stage, we analyze the building blocks. We start by deconstructing them into their base elements, paragraphs, and individual

clauses. I used to do this manually. I started this work at a law firm where I developed standard forms of wills, trusts, license agreements, and many others. I began by asking the partners for all of their templates and exemplars. The result was a pile of documents, many feet high. And, in order to review them manually, I had to read the documents multiple times in order to develop a good standard. My technique was to print all the documents, examine the stack, start with the longest document (because it probably contained the most clauses), and collect all the clause captions. Then, I went through the documents one-at-a-time, and recorded the names of the clauses I hadn't seen before. Next, I attempted to organize those clauses in some kind of logical manner. Finally, I went through the stack multiple additional times to identify the standard language for each term and all the variations of that clause. Often, the only way to do it was to lay the pieces of paper out on the floor of the office and try to play a memory game. This would take months. Of course, this is one of the main reasons why we have not been able to apply technology broadly to transactional practice. The manual processes behind contract analysis and document assembly are simply too time consuming and expensive. For me, something like a will would take me two to three months to fully analyze. The software, as you just saw, can review a merger agreement in just a few seconds.

One of the most powerful results of the analysis is a checklist. It shows all the clauses found in the source set of agreements, organizes the terms into a single table of contents, and offers analysis with respect to the commonalty and consistency for each term. For example, if we analyze a set of employment agreements, we will likely find a base salary clause in every contract. However, we may find an overtime clause in only a small percentage of the samples. We can deduce that the base salary is a standard or required, while the overtime clause is optional, to be used in specific situations. The software can't determine when to include the clause, but it can raise the question.

The second key element of the analysis is the examination of clause language across the set of documents. Using the power of statistics we identify the baseline. I want to know what is the most standard way of drafting this clause. The baseline is the clause example that contains the most standard terms and the least divergent terms. We can then compare all of the other clauses in a set to the baseline and view their points of similarity and divergence.

The resulting analysis organizes the clauses in a matrix of four quadrants, based on commonality and consistency. We can identify the

clauses that are common and consistent. These are the standard, or "boilerplate" terms. We can find the clauses that are common, but different from one contract to another. These are the critical, highly negotiated terms (and those provisions that are just different because of circumstance or preference). We can find the terms that are uncommon, but when included, they are consistent. These are typically deal-specific markers, such as the overtime clause. Finally, we can identify the uncommon and divergent clauses, representing the deal-specific terms.

This is where we are today. We can find the relevant material. We can identify the relevant elements. And, with this analysis, we can reverse engineer the legal logic from a sample set of documents, and create a term sheet. However, the software cannot tell you how to select the best terms for your client. This is the last stage: the ability to predict the optimal outcome.

What do we need to achieve the third stage of learning? Technically, we have the means to succeed. We have the computer power to crunch the numbers. But we don't yet know if the results will be satisfactory. In order to predict the optimal outcome, we need to capture the key deal terms behind each transaction. These deal terms are data elements such as names, dollar amounts, dates, numbers, percentages, place names and other key business terms. This is the domain of big data.

In the case of contract analysis, we may not have sufficient volume. If we pull examples from EDGAR, we may find a few hundred-thousand employment agreements. But, any one law firm may have just a small handful. In addition, we may not be able to detect statistical correlations. For example, it is commonly known that California courts will not enforce non-compete agreements. However, the analysis of a large sample set shows that most have non-compete terms.

5. CONTRACT STANDARDS

Empirical analysis clearly identifies the contract building blocks. There is a recipe—a standard recipe—for building an employment agreement, for licensing software, or for buying a company. Interestingly, the analysis also shows that the more sophisticated the transaction, then the more standard the document is today. A merger agreement—while much more complex compared to an employment or license agreement—is far more standard in structure. Why is this true? How could merger agreements be more conforming than an employment agreement? First, todays' technology enables us to access and share agreements through EDGAR and other public and private sources. But, this is also true for

many other types of agreement. The principal driver of conformity is the fact that the universe of lawyers that draft merger agreements is far smaller than the universe of lawyers that draft employment agreements, some of whom are not trained as lawyers.

Looking across a wide range of legal agreements, we can detect common patterns. Applying the work of Herbert Simon, who developed the science of complexity, we can use hierarchical deconstruction to further our understanding of contracts. We can deconstruct all agreements into common frameworks, deconstruct agreements into checklists, and deconstruct checklists in elemental clause language.

5.1 FRAMEWORK DECONSTRUCTION

First, we can propose a common structure for all agreements. Common themes appear in exchange agreements, organizational documents, and distribution agreements (for example, wills). We call this the Contract Standards Framework. In the case of exchange agreement—used for purchases, licenses, and performance agreements—we can detect a common three-layer structure shared by a wide range of contracts. The first layer handles the bargain, the mechanics of the exchange, and the term of the obligation. The second set of clauses capture the insurances that seek to guarantee the nature of performance of the exchange. These terms are the representations, statements, conditions, and obligations. The third layer defines the rights, remedies, and general provisions that provide assurances in the event of non-performance or other failure in the bargain. The first layer—comprising the price or core terms—is typically negotiated by business people. The second and third layers are generally negotiated by lawyers. But, in many cases, the lawyers negotiating these non-core terms rarely have the ability to amend the price.

The Contract Standard Framework also serves as a very useful teaching tool to help lawyers learn how to read, review, analyze, and audit agreements. The framework is a hierarchical decision tree of common transaction elements. It can guide lawyers through the contract process, posing a standard set of questions. What is the nature of the consideration given and received? What is included and what is excluded? How is the consideration exchanged?

5.2 AGREEMENT DECONSTRUCTION

Next, we can deconstruct the agreements into checklists, or term sheets. Again we can harness technology to aid our understanding. Contract

analytics can find all the terms found in any particular type of agreement, and organize them into a meaningful table of contents that mirrors the standard organizing structure of the transaction.

Using the statistical measures of commonality and consistency, we can further analyze the terms into a matrix containing four quadrants. In the upper-right sector, we find clauses that are very common and highly consistent, such as a governing law clause. These are the standard (or boilerplate) clauses that can be conformed across all agreement types. In the upper-left quadrant, we find clauses that are very common, but are highly divergent. These are typically the critical, highly negotiated clauses, such as a bonus clause. However, this quadrant also contains clauses that are different due to personal preference, historical circumstances, or because the lawyer selected a particular starting point document.

5.3 LANGUAGE DECONSTRUCTION

Finally, we can decompose the clause language into its elemental building blocks. We train our staff to break the long clause paragraphs into individual sentences, provide captions for each sentence, and redraft each sentence clearly, comprehensively, and concisely. When more clearly presented we can better see the elemental building blocks. They are statements, words defining scope and coverage, standards, and variables.

For example, in an employment agreement, we may find a clause stating: "the Executive hereby accepts employment and agrees faithfully to perform to the best of his abilities, the duties described in section 1." The example contains two main terms: an acceptance of employment and the standard of performance. In the case of the standard, an analysis of a set of agreements finds a range of performance measures. The most common standard is "faithfully and diligently." We also find standards such as "competently" and "to the best of the Executive's ability." With the standards identified, we can now organize them on a continuum of favorability to each party. The balanced standard uses the term "faithfully and diligently." If the parties wish to make the standard more favorable to the employer and measure the performance of the executive to market standards, they should use the word "competently." If the parties wish to measure performance by reference to the capability of the individual, then use the term "best efforts."

6. THE FUTURE

Today, software can rapidly analyze legal agreements and inform lawyers how the agreements are organized, what clauses they contain, and

the full range of standard and deal-specific language. We can see patterns of similarity and difference across all contracts and within any agreement. We can inform practitioners what terms are standard, and which are optional. We can automatically reverse engineer the legal logic behind the agreements. But we cannot instruct lawyers what are the best terms for any particular situation.

We need to deepen our understanding. First, we need to be able to extract the key deal points and match them to contract outcomes. Secondly, we need objective measures of performance to offer some guidance about what makes a "good" contract. Both represent significant programming challenges. But, it is something that we can explore together.

7. GAME THEORY

Can we do this in a way to make teaching fun? Can we turn it into a game? The Contract Standards Framework can serve as a board game. The contract terms can be the playing cards placed on the board. Each side can take turns laying down a set of cards representing their proposed agreement. The other player can accept, reject, or revise the cards until the players come to an agreement. One difficult matter to consider is how do you win? How do you define a "good" contract? What are the appropriate contract performance metrics? It may be that there is no objective contract value, but rather a range of risks and rewards that determine whether the parties will come to an agreement, and thresholds where they will walk away from the deal.

We can use the game to test valuation and behavioral models. We can have a standard set of cards. However, we do not have to give all of the cards to every player. We could hold a few back and see if players will deduce the need for particular terms. We can run scenarios with different players: such as first years, third years, and maybe business school graduates. I hope we can develop the game concept jointly. My company will make the software freely available to all law schools.[3]

3. Note: The Contract Game is now called the Contract Practicum. For more information, please *see* http://lawprofessors.typepad.com/contractsprof_blog/2014/08/northwestern-laws-zev--eigens-new-drafting-exercise.html, [http://perma.cc/L24U-NNQM].