

Enhancing Regulatory Compliance by Using Artificial Intelligence Text Mining to Identify Penalty Clauses in Legislation

Nachshon Goltz and Michael Mayo*

This article provides the theoretical basis of machine learning for text classification and presents a two-stage experiment of (1) training multiple models and selecting the best one, and (2) employing a sliding window detection in order to identify penalty clauses in regulation.

As regulatory compliance (or compliance governance) becomes ever more challenging, attempts to engage IT solutions and especially artificial intelligence (“AI”) are on the rise. This article suggests that regulatory compliance can be enhanced by employing an AI model trained to identify penalty clauses in the regulations. The article provides the theoretical basis of machine learning for text classification and presents a two-stage experiment of (1) training multiple models and selecting the best one, and (2) employing a sliding window detection in order to identify penalty clauses in regulation. Results benchmarked using an algorithm-based penalties API suggests further development is needed.

Introduction

Challenges in regulatory compliance are closely associated with the increased business opportunities resulting from globalization and Information Communication Technology.¹ There is a drastic increase in the regulatory requirements with which businesses must comply, not only in sheer number but also in complexity, confronting businesses with the need to adapt to a complex and evolving regulatory environment.² In such an environment, organizations are finding compliance with the numerous regulations to be very expensive,³ with an estimated \$1 trillion spent worldwide on regulatory compliance.

Compliance is defined as “adhering to the requirements of legal, industry, and organizational standards and codes, to principles of good governance, and to accepted community and ethical standards.”⁴ Regulatory compliance is defined as the set of activities and policies in place in an enterprise to ensure that the business activities required to achieve the business goals of the company comply with the relevant normative requirements.⁵ Sadiq and Governatori define compliance as ensuring that business processes, operations, and practice are in accordance with a prescribed and/or agreed set of norms.⁶

According to El Kharbili et al.,⁷ regulatory compliance consists of measures and directives, which are implemented by policies, internal controls, and procedures and which are modeled for business processes. They identified three classes of compliance rules: regulations, information technology security standards, and quality standards.

When businesses are confronted with evolving and diverse regulatory requirements from multiple sources, it is often suggested that an integrated approach to governance, risk, and compliance (“GRC”) is useful.⁸ At the heart of integrated GRC is the adoption of a risk-based approach to compliance, where resources are allocated to the areas they are most needed based on risk levels.

Mahler⁹ defines legal risk as a risk resulting from a set of facts that are assessed under a set of legal norms. By comparison, compliance risk is considered a risk resulting from a failure to comply with laws, regulations, rules, related self-regulatory organization standards, and applicable codes of conduct. A common denominator for both risks is legal compliance risk, a risk resulting from a failure to adhere to the requirements of the law.

At a strategic level, compliance is naturally related to the concept of risk. Noncompliant situations expose a company to risks that can often be mitigated. Risk mitigation is the actual driver for internal compliance auditing. The risk entity represents the risks a company wants to monitor; risks are associated with compliance requirements.¹⁰

KPMG, one of the four big global accounting companies, claims that “[t]he top risk perceived by senior executives is the growing regulatory pressure from governments around the world. C-level executives in almost all industries say this, not just those in Financial Services, where companies are facing arguably the greatest regulatory challenge in their history.”¹¹

The market value for compliance-related software and services was estimated as over \$32 billion in 2008,¹² whereas only the eGRC market size is expected to grow from \$19.42 billion in 2016 to \$38 billion by 2021.¹³ This boost in business investment is primarily a consequence of regulatory mandates that emerged as a result of events that led to some of the largest scandals in corporate history such as Enron,¹⁴ WorldCom (United States),¹⁵ HIH (Australia),¹⁶ Societe Generale (France), and, above all, the global economic crisis of 2008.

Since the global financial crisis of 2008, spending on compliance has broken new records: over \$1 trillion is spent worldwide on regulatory compliance.¹⁷ Over one million people are employed around the world in the regulatory compliance industry.¹⁸

Several recent works have addressed the issue of (regulatory) compliance in the context of business process management, service computing, and cloud computing domain.¹⁹ The general idea is to determine whether the constraints (i.e., norms) imposed by some regulatory framework (ranging from acts, to regulations, to industry standards, to best practices and internal policies) are met by some IT systems.²⁰

In order to enable an automated compliance checking, it is necessary that, besides a formal representation of process models, compliance requirements exist in an analyzable format. A formalized compliance requirement is a structural pattern,²¹ which defines what the structure of a subsection of a process model has to look like in order to comply with the underlying rule.

Assessing compliance demands for an interpretation and translation of the requirements provided in natural language in an actionable rule description (especially in the case of principle-based regulations).²² Therefore, AI may be suitable for this task.

AI is the theory and development of computer systems able to perform tasks that would otherwise require human intelligence. Example tasks include visual perception, speech recognition, decision making under uncertainty, learning, and translation between languages.²³

Venture capital investments in companies developing and commercializing AI-related products and technology exceeded \$2 billion since 2011.²⁴ Leading players such as IBM,²⁵ Google,²⁶ and Facebook²⁷ have invested heavily in developing their AI capabilities.

No doubt the alleged AI legal application that has received the most public attention is ROSS, a system supported by IBM's

Watson division, claiming to be a “junior associate.”²⁸ One of the co-founders of the ROSS team describes it as “[b]asically, what we built is a [sic] the best legal researcher available.”²⁹ Even without being made available to the public or presented in any public demonstration, ROSS has become a symbol of legal AI technology.

Other fields in which AI has been used within the legal profession are e-discovery (Recommind, Equivio—now part of Microsoft), forecasting outcomes of IP litigation (Lex Machina—now part of Lexis), providing fact and context-specific answers to legal, compliance, and policy questions (Neota Logic) and contract life-cycle software, including discovery, analysis, and due diligence (Kira Systems and KM Standards).

Given the challenging and complex regulatory structure, the cost of compliance and the presumed abilities of AI, the authors’ theoretical assumptions are:

1. it is possible to train an AI model to identify penalty clauses in regulations, and
2. identifying penalty clauses in regulations can enhance compliance.

This article will briefly describe the field of machine learning for text classification, and will discuss the proposed AI system, initial exploratory experiments, further sliding window detection experiments, followed by a summary and conclusions.

Machine Learning for Text Classification

In this research we take a machine learning approach to detecting compliance clauses in legislation. Machine learning is one of the dominant paradigms within AI and is concerned with learning from data. The particular style of machine learning we focus on is classification, i.e., learning rules for assigning class labels to unlabeled examples.

In the classification approach to machine learning the user provides the system with a set of labeled examples $\{(X_1, Y_1), (X_2, Y_2), \dots (X_n, Y_n)\}$ called the training data where each X_i is a fixed length numeric feature vector of dimensionality D . Each Y_i is an element of the set of class labels $C = \{c_1, c_2, \dots, c_K\}$. The aim of a classification-based algorithm is to train a model that can accurately

predict class labels for unlabeled examples X_{n+1} , X_{n+2} , ... etc., which are not present in the training data. The unlabeled examples are called test data. If the true labels for the test data are known (but withheld from the learning algorithm), then the accuracy of any given learning algorithm can be assessed and different algorithms for classification can be compared.

Hundreds of machine learning algorithms for classification have been developed. In the experiments described here, we select a representative sample of algorithms that are commonly used in research and practice, specifically the following:

- Naïve Bayes,³⁰ an efficient algorithm for learning a conditional probability distribution $P(C | X_i)$ over the class labels. The algorithm's efficiency and simplicity arises from its assumption of the conditional independence of the features, which is often an incorrect assumption but can still yield high classification accuracy.
- The C4.5 decision tree learning algorithm.³¹ Decision trees have the advantage that they are often succinct and therefore interpretable by humans. A decision tree consists of a set of tests of feature vector values arranged into a tree structure, with predicted class labels positioned at the leaves of the tree.
- Sequential Minimal Optimization ("SMO"),³² an algorithm for learning support vector machines. Support vector machines are state-of-the-art classifiers that learn the optimal hyperplane separating examples with different class labels. When unlabeled examples are to be classified, the example's position with respect to the learned hyperplane is used to predict its class label.
- Random forests,³³ an ensemble method in which multiple decision trees are learned, each tree being learned from different random projection of the training data. When new examples need to be classified, the decision trees vote on the class label, and the class label with the majority vote is predicted.

In this work, each training example represents a fragment of text and the set of class labels is {yes, no}, indicating whether or not the fragment contains a compliance clause. For example, one of the clauses labeled yes in our dataset³⁴ is:

“Section 8(1.2): Fines

A person who contravenes section 7.5 is guilty of an offence and liable

(a) for a first offence, to a fine of not more than \$10 000, and

(b) for a 2nd or subsequent offence, to a fine of not more than \$100 000.”

Machine learning algorithms cannot be applied directly to text fragments like this because the classifier expects examples to be fixed length vectors numbers. We therefore must convert the text fragments into vectors before they can be further processed, and two common sets of approaches exist for achieving this: bag-of-words (“BOW”)-based approach and word2vec (“W2V”)-based approach.

The BOW approach is a classical text processing approach in which the order of words in a document is ignored and only the presence/absence of a word is important. For example, the clause above can be converted into the vector shown in Table 1 using the BOW approach. Typically stop words such as “is” and “a” are removed before forming the vector.


One drawback of the BOW approach is that the dimensionality D of the training example vectors is equal to the number of distinct words in the set of training examples. Hence, for large and varied text datasets, the dimensionality can be quite large (in the hundreds or even thousands of words).

A more modern alternative to BOW is the recently proposed W2V approach.³⁵ W2V is inspired by the distributional hypothesis in linguistics,³⁶ which proposes that words that occur in similar contexts have similar meaning. One W2V method uses a neural network to predict the context of a word (i.e., the words preceding and following the word) given the word. Because the neural network contains a hidden layer, the activation values on the hidden layer are used to predict the context whenever a word is presented to the network. Following the distributional hypothesis, words with similar meaning should therefore have similar activation values on the hidden layer of the network. These activation values can therefore be used to “represent” the words as points in a space whose size corresponds to the dimensionality of the neural network’s hidden layer. Such a representation is called an embedding.

Figure 1 illustrates word embedding for some of the words in the example text fragment shown earlier. Directly visualizing points in a very high (i.e., 300) dimensional space is difficult, so shown

Table 1. Example of a BOW Representation of a Text Document	
contravenes	1
convicted	0
fine	1
guilty	1
liable	1
day	0
continues	0
subsequent	0
... etc.

Figure 1. Words Vectors from the Previous Example Visualized by Plotting Their First Two Principle Components



in the figure are the first two principle components of the word embedding. The first principle component is measured along the x axis. It can clearly be observed that words related to compliance such as “Fines,” “guilty,” “liable,” and “contravenes” lie close together along this component but are separated along the second principle component. Other pairs of words (e.g., “more” and “than”) that are related are also close together in embedding space.

Table 2. Algorithm for Computing a 300-Dimensional Representation of a Document				
Word	x_1	x_2	...	x_{300}
A	0.044241	0.089298	...	0.014781
person	0.120826	-0.108401	...	0.040490
who	0.050564	0.061801	...	0.031993
contravenes	0.062897	-0.111759	...	0.048083
... etc.
Average	0.016457	-0.033592	...	-0.002725

Because training a neural network to compute the word embedding for millions of words is computationally expensive (typically requiring billions of examples for robustness and therefore significant resources) we make use of a pre-trained set of three million word embedding provided by Google. Prior research has shown that this set of embedding is surprisingly robust even when applied to text classification problems that are not news related.³⁷

The dimensionality of each word embedding is 300 (significantly lower than that arising from the BOW method), and given all the word vectors in a document, a single “document vector” for an entire text fragment can be produced by a simple averaging procedure. This process is outlined in Table 2, and has been shown in several prior works³⁸ to produce robust and accurate feature vectors for classification.

A disadvantage of both approaches to generating feature vectors from text is that they typically ignore any parts of the text that are non-alphabetical, such as punctuation symbols, formatting, and numbers. This may impact the performance of the system when the presence or absence of these features (e.g., fine amounts) in the text is important.

Initial Exploratory Experiments

Our initial exploratory experiments tested the theoretical assumptions brought in the introduction, according to which the use of machine learning will enable to identify compliance clauses in the regulations and, hence, enhance compliance. Given the initial exploratory nature of this experiment, we avoided using analysis

based on Formal Concept Analysis and Distributional Semantic Techniques (such as Latent Semantic Analysis and Topic Modeling).

Curated Canadian Dataset

In order to test the hypothesis, we identified 40 negative and 40 positive penalty clauses within Canadian federal and provincial legislation.³⁹ The negative examples are clauses from laws or regulations that does not contain penalties. The positive examples are clauses taken from Canadian federal and provincial legislation and regulation that contains penalties. Both clauses contain dollar amounts.

For example, the following clause taken from the Canada Deposit Insurance Corporation Act,⁴⁰ is a positive penalty clause:

50.1 Every member institution or other person who commits an offence under this Act is liable on summary conviction

- (a) in the case of a natural person, to a fine not exceeding \$100,000 or to imprisonment for a term not exceeding twelve months, or to both; or
- (b) in any other case, to a fine not exceeding \$500,000.

And a negative example would be the following clause taken from Alberta: Election Finances and Contribution Disclosure Act Section 13(2) Exemptions:⁴¹

Money or goods provided by any person that do not exceed \$50 in the aggregate are not a contribution for the purposes of this Act but shall be recorded as to the gross amount by the chief financial officer of the recipient unless the donor specifically requests that the amount be considered a contribution.

Results

Given the dataset of examples described in the previous section, we next performed a set of experiments to determine which machine learning algorithm would classify the examples with the highest accuracy. As described above, we focused on four different machine learning algorithms (naïve Bayes, C4.5, SMO, and random forests) along with two different methods for converting documents into feature vectors (BOW and W2V).

Table 3. Leave-One-Out-Cross-Validation Percent Correct Across the Dataset of 80 Examples, by Feature Vector Type and Machine Learning Algorithm

	Naïve Bayes	C4.5	SMO	Random Forest
BOW	97.50	92.50	96.25	97.50
W2V	91.25	78.75	100.00	92.50

The approach taken to evaluate the machine learning in this initial set of experiments is leave-one-out-cross-validation, a methodology in which a given machine learning algorithm is trained on all but one of the labeled examples in the dataset (79 in this case) and tested on the single remaining held-out example. This train and test procedure is then repeated for each of the remaining 79 examples until each example has been tested, so that an overall accuracy metric for algorithm can be computed.


Note that since the single test example in each round is not present in the training data for that round, then the algorithm is always being tested on an example that it has not “seen” in its training dataset. Therefore, the outcome of the experiment should give an indication of how well each algorithm performs on unlabeled data.

To carry out our experiments, we used implementations of the algorithms available in Weka 3.8.0⁴² with default settings.⁴³ As mentioned previously, the word vector data is obtained from Google⁴⁴ after being trained on a corpus of news items.

The results of the experiment are summarized in Table 3. As can be clearly observed, the SMO algorithm in conjunction with W2V achieves the greatest performance with 100 percent accuracy. Unfortunately, all of the other W2V approaches underperform the BOW approach, especially the C4.5/W2V combination, which achieves only 79 percent. The average accuracy, however, for all approaches is above 90 percent.

Curious as to the reason why the accuracy for most methods was so high, we next explored the accuracy of a very simple classification method that classifies an example as positive if it contains the word “fine” and negative otherwise. The accuracy of this simple approach is 96.25 percent on our dataset of 80 examples. This fact happens to explain why the BOW approach performs well compared to the W2V approach: the BOW feature vector has an explicit feature for the presence of the word “fine” (see Figure 1),

Figure 2. Leave-One-Out-Cross-Validation Percent Correct Across the Dataset of 80 Examples, by Feature Vector Type and Machine Learning Algorithm



and a classifier only has to detect this single important feature in order to perform well with high accuracy.

However, W2V-based algorithms do not have this shortcut—instead the word vector for “fine” is averaged into one overall document vector, as illustrated in Table 2. The fact that the SMO/W2V algorithm came out as the top performer therefore makes this approach all the more impressive.

Next, we decided to perform further tests of two of the algorithms on more challenging and realistic test data. The results of that experiment are described in the next section.

Sliding Window Detection Experiments

In the next set of experiments, we took two of the classifiers used in the previous experiment and tested them on some more challenging and realistic test documents. The first classifier we tested was SMO/W2V, which was selected because it had the highest accuracy in the previous experiment. The second classifier we selected was random forest/W2V, which, while not the most accurate method, is interesting to examine further because it gives reasonable probability estimates when it makes a prediction (e.g., 0.63 probability of a yes), whereas SMO typically does not produce proper probabilities.

The methodology we used for this experiment was as follows. One of the machine learning algorithms (either SMO/W2V or random forest/W2V) was trained on all 80 examples in our labeled dataset. Next, we obtained a selection of further test documents—four Canadian federal and provincial Acts and one U.S. document. These documents had not been labeled. These documents were

much longer than our training examples since they are complete acts rather than small fragments.

To account for the size disparity between the training examples and the longer test documents, we adopted a sliding window approach and considered only 50 words of each test document at a time. For each possible position of the sliding window, a prediction was made (yes or no) using both of our models. We then advanced the sliding window by 25 words (to ensure that each window would overlap) and repeated the process.

This meant that the number of predictions made per test document varied depending on its length, from 38 predictions (for the shortest document) to 109 predictions (for the longest document).⁴⁵

Results

*The Ferries Act*⁴⁶—This Act contains one penalty clause that was identified positively by both methods.

*Motor Vehicle Fuel Consumption Standards Act*⁴⁷—Both methods missed the first penalty clause and identified it as negative. However, the second penalty clause was clearly identified by both methods as positive. Few additional clauses discussing “inspector,” “order,” and other enforcement-related subject were identified positive by the SMO/W2V and scored relatively high (≥ 40) in the random forest/W2V.

*The Partnership Act*⁴⁸—This Act does not contain any penalties. SMO/W2V yielded 10 positive results and the random forest/W2V yielded 14 results with yes probability equal or higher than 50 percent. Two positive results in the SMO/W2V yielded less than 50 percent in the random forest/W2V and two of the positive results (≥ 50 percent) in the random forest/W2V yielded negative results in the SMO/W2V. It is interesting to note that although this legislation did not contain any penalties, the majority of the clauses identified by both methods as positive dealt with related subject of liability.

*The Pay Equity Act*⁴⁹—This Act does not contain any penalties. The SMO/W2V indicated three cases of positive penalty clauses whereas only the middle one was positive in the random forest/W2V as well. These clauses discussed related subject of court order.

In order to benchmark the results, we ran the same legislation through publicly available algorithm-based penalties API.

Table 4. Sliding Window Detection Experiments					
Test Documents	Penalty Clauses	SMO/W2V	Random Forest/W2V	False Positive Content	Benchmarking
The Ferries Act	2	V/X	V/X		V/V
Motor Vehicle Fuel Consumption Standards Act	3	X/X/X 5 false positive	V/X/X 5 false positive	“inspector,” “order,” and other enforcement-related subjects	V/V/V
The Partnership Act	0	10 false positive	14 false positive	subject of liability	0
The Pay Equity Act	0	3 positive	1 positive	subject of court order	0

With regard to the first Act (the Ferries Act), the algorithm-based API reported: “Found penalty clauses but not dollar figures.” The penalty clause in this Act (positively identified by the sliding window method) contain the words “dollar,” “two,” and “eight,” but not the signs “\$,” “2,” and “8.” This could be the reason for the algorithm-based API to miss this clause while the sliding window correctly identified it. This was also the case with the second Act, the Motor Vehicle Fuel Consumption Standards Act.

The Partnership Act does not contain penalty clauses. However, both our experiment and the algorithm-based API identified positive penalty clauses. The API was not specific and just reported: “Found penalty clauses but not dollar figures.”

Finally, with the Pay Equity Act, the algorithm-based API identified correctly that there are no penalty clauses while our experiment identified three cases of false positive penalty clauses (dealing with compliance-related matters).

The results are summarized in Table 4.

Conclusions

We conducted two experiments with the purpose of using AI text mining in order to identify penalty clauses in legislation.

From a practical perspective, the second experiment results are more important since the testing was done using “real data”

(i.e., legislation), whereas in the first experiment the samples were tested against each other.

While the results of the first experiment (BOW/W2V) were impressive with above 90 percent accuracy (*see* Table 2), the results of the second experiment (Sliding Window Detection) showed less accuracy. Nonetheless, the use of an algorithm-based penalties API as a benchmark shows potential in further development of the AI-based method.

In light of the lessons learned from this experiment, there is a need to consider replacing two basic elements used in the current method:

1. the basis for the current transformation system from words to vector was based on Google's news dataset. Further development of the AI-based method would benefit from creating a specific legal dataset that may generate more accurate transformation and thus improve the system overall accuracy;
2. the current system used Google's W2V tool to transform the text into a bag of word vectors. This tool as we have used it ignores punctuation marks and numbers, since it is only concerned with vectors for English words. This may not be adequate for the purpose of identifying penalty clauses where symbols such as dollar signs and figures may be important. Further development of the AI-based method would benefit from using Facebook's "fastText" system,⁵⁰ which builds word vectors indirectly from n-gram character vectors and can therefore produce vectors for punctuation marks and numbers properly.

Employing these elements along with a bigger dataset of examples may produce better results. Moreover, using the two methods, an algorithm-based system and an AI-based system, in tandem could produce the most accurate results and hence enhance regulatory compliance.

Notes

* Nachshon (Sean) Goltz is a lecturer at Waikato University School of Law and co-founder of Global-Regulation.com. Michael Mayo is a senior lecturer in the computer science department at the University of Waikato. The authors thank

Addison Cameron-Huff and Antonino Rotolo for their illuminating comments and Liam Whittaker for his research assistance. This article was first presented at MIREL 2017—Workshop on “MIning and REasoning with Legal texts” as part of ICAIL 2017—The 16th International Conference on Artificial Intelligence and Law, June 16, 2017, London (UK).

1. Samson Esayas & Tobias Mahler, Modeling compliance risk: a structured approach, *Artificial Intelligence and Law* (2015) 23:271–300.

2. Giblin C., Liu AY., Muller S., Pfitzmann B., and Zhou X. (2005) Regulations expressed as logical models (REALM). In: Proceedings on legal knowledge and information systems: the eighteenth annual conference, IOS Press, pp. 37–48.

3. Ponemon Institute (2011) The role of governance, risk management and compliance in organizations study of GRC practitioners. Research Report, Sponsored by RSA, Security Division of EMC, Bedford.

4. Standards Australia (2006) Compliance programs AS 3806-2006, 2nd ed. Australia, Sydney.

5. Governatori, Guido, Business Process Compliance: An Abstract Normative Framework, *Information Technology* 55(6) (2013) 231-238 at 231.

6. Shazia Wasim Sadiq, Guido Governatori: Managing Regulatory Compliance in Business Processes, *Handbook on Business Process Management* (2) 2015: 265-288.

7. El Kharbili, Marwane, Sebastian Stein, Ivan Markovic, and Elke Pulvermueller (2008): Towards a Framework for Semantic Business Process Compliance Management, *CEUR Workshop Proceedings*, 339: 1-15.

8. Bonazzi R., Hussami L., and Pigneur Y. (2010) Compliance management is becoming a major issue in IS design. In: D’Atri, A., Saccà, D. (eds) *Information systems: people, organizations, institutions, and technologies*. Springer, Berlin, pp. 391–398.

9. Mahler T. (2010) Legal risk management: developing and evaluating elements of a method for proactive legal analyses, with a particular focus on contracts. Dissertation, University of Oslo.

10. Silveira, Patricia, et al., On the Design of Compliance Governance Dashboards for Effective Compliance and Audit Management, January 2009, Conference: Service-Oriented Computing. ICSOC/ServiceWave 2009 Workshops—International Workshops, ICSOC/ServiceWave 2009, Stockholm, Sweden, November 23-27, 2009, Revised Selected Papers.

11. <https://assets.kpmg.com/content/dam/kpmg/pdf/2013/08/expectations-risk-management-survey-v3.pdf>.

12. Hagerty, J., J. Hackbush, D. Gaughan, and S. Jacobson (2008). The governance, risk management, and compliance spending report, 2008–2009: Inside the \$32B GRC Market. AMR Research, Boston, 25 Mar. 2008.

13. Enterprise Governance, Risk, and Compliance (eGRC) Market by Component (Software (Type & Usage), Service), Deployment Model (Cloud & On-Premise), Organization Size, Business Function (Finance, IT, Legal, & Operations), & Vertical—Global Forecast to 2021, <http://www.marketsandmarkets.com/PressReleases/enterprise-governance-risk-compliance.asp>.

14. Enron’s Collapse: The Overview; Enron Collapses as Suitor Cancels Plans for Merger, <http://www.nytimes.com/2001/11/29/business/enron-s-collapse-the-overview-enron-collapses-as-suitor-cancels-plans-for-merger.html>.

15. Worldcom's Collapse: The Overview; Worldcom Files for Bankruptcy; Largest U.S. Case, <http://www.nytimes.com/2002/07/22/us/worldcom-s-collapse-the-overview-worldcom-files-for-bankruptcy-largest-us-case.html>.

16. HIH: The Inside Story of Australia's Biggest Corporate Collapse, <http://www.smh.com.au/articles/2003/03/14/1047583693489.html>.

17. http://www.blplaw.com/media/know-how/speed-of-business/BLP_The-Speed-of-Business-impact-of-regulation.pdf.

18. <https://www.linkedin.com/vsearch/p?keywords=regulatory+compliance&queryRewriteParams=1>.

19. Becker J., Delfmann P., Eggert M., and Schwittay S. (2012) Generalizability and Applicability of Model-Based Business Process Compliance-Checking Approaches—A State-of-the-Art Analysis and Research Roadmap. *BuR—Business Research Journal* 5(2):221–247; Fellmann M., and Zasada A. (2014) State-of-the-Art of Business Process Compliance Approaches. In: 22st European Conference on Information Systems, ECIS 2014, Tel Aviv, Israel, June 9-11, 2014.

20. Mustafa Hashmi, Guido Governatori, Moe Thandar Wynn: Normative requirements for regulatory compliance: An abstract formal framework. *Information Systems Frontiers* 18(3): 429-455 (2016).

21. Ghose, Aditya and George Koliadis (2007): Auditing Business Process Compliance, in: Bernd J. Krämer, Kwei-Jay Lin, and Priya Narasimhan (eds.), *Service-Oriented Computing (ICSOC 2007)*, Springer: Berlin et al., 169-180.

22. C. Giblin, S. Müller, and B. Pfitzmann. From Regulatory Policies to Event Monitoring Rules: Towards Model-Driven Compliance Automation. IBM Research Report RZ 3662, Zurich, October, 2006. 10. K. Namiri, N. Stojanovic. A Semantic-Based Approach for Compliance Management of Internal Controls in Business Processes. *CAISE'07*, pp. 61-64, 2007.

23. Oxford Dictionaries, "Definition of artificial intelligence," https://en.oxforddictionaries.com/definition/us/artificial_intelligence.

24. Demystifying artificial intelligence—What business leaders need to know about cognitive technologies, Deloitte Uni. Press.

25. IBM, "IBM Watson," <https://www-03.ibm.com/press/us/en/presskit/27297.wss> (IBM has committed \$1 billion to commercializing Watson, its cognitive computing platform).

26. Amit Chowdhry, "Google to acquire artificial intelligence company DeepMind," *Forbes*, January 27, 2014, <https://www.forbes.com/sites/amitchowdhry/2014/01/27/google-to-acquire-artificial-intelligence-company-deepmind/> (Google has made major investments in AI in recent years, including acquiring eight robotics companies and a machine learning company).

27. Josh Constine, "NYU 'Deep Learning' Professor LeCun will head Facebook's new artificial intelligence lab," *TechCrunch*, December 9, 2013, <http://techcrunch.com/2013/12/09/facebook-artificial-intelligence-lablecun/> (Facebook hired AI luminary Yann LeCun to create an AI laboratory with the goal of bringing major advances in the field).

28. University of Toronto's next lawyer: A computer program named Ross, <http://www.theglobeandmail.com/report-on-business/industry-news/the-law-page/university-of-torontos-next-lawyer-a-computer-program-named-ross/article22054688/>.

29. Addison Cameron-Huff, *Why I'm Skeptical That Watson Will Be Your Next Junior Associate*, March 2, 2015, <https://www.cameronhuff.com/blog/why-i-am-skeptical-of-watson/>.
30. G. John, P. Langley. *Estimating Continuous Distributions in Bayesian Classifiers*. In: *Eleventh Conference on Uncertainty in Artificial Intelligence*, San Mateo, 338-345, 1995.
31. J. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers, 1993.
32. J. Platt: *Fast Training of Support Vector Machines using Sequential Minimal Optimization*. In B. Schoelkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods—Support Vector Learning*, 1998; S. Keerthi, S. Shevade, C. Bhattacharyya, K. Murthy (2001). *Improvements to Platt's SMO Algorithm for SVM Classifier Design*. *Neural Computation*. 13(3):637-649.
33. L. Breiman (2001). *Random Forests*. *Machine Learning*. 45(1):5-32.
34. *Canada Alberta Tobacco and Smoking Reduction Act 2*.
35. T. Mikolov, K. Chen, G. Corrado, and J. Dean. *Efficient estimation of word representations in vector space*. Technical report, arXiv preprint arXiv:1301.3781; T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, *Distributed representations of words and phrases and their compositionality*. In *Advances in Neural Information Processing Systems*, 26 (NIPS 2013), 2013.
36. S. Magnus (2008). *The Distributional Hypothesis*. *Rivista di Linguistica*. 20 (1): 33–53.
37. M. Kusner, Y. Sun, N. Kolkin, and K. Weinberger. *From word embeddings to document distances*. In *32nd International Conference on Machine Learning (ICML), Proc. of*, 2015.
38. *E.g.*, M. Iyyer, P. Enns, J. Boyd-Graber, and P. Resnik. *Political ideology detection using recursive neural networks*. In *52nd Annual Meeting of the Association for Computational Linguistics (ACL), Proc. of*, pages 1113–1122, 2014.
39. *See the dataset at* <https://github.com/mmayo888/MIREL2017Supplementary>.
40. <http://laws.justice.gc.ca/eng/acts/C-3/FullText.html>.
41. http://www.qp.alberta.ca/1266.cfm?page=E02.cfm&leg_type=Acts&isbncln=9780779787241&display=html.
42. <https://www.cs.waikato.ac.nz/~ml/weka/>.
43. E. Frank, M. Hall, and I. Witten (2016). *The WEKA Workbench*. Online Appendix for *Data Mining: Practical Machine Learning Tools and Techniques*, Morgan Kaufmann, Fourth Edition, 2016.
44. <https://code.google.com/archive/p/word2vec/>.
45. *These numbers will change when the Canadian docs are added to the experiment*.
46. <http://nslegislature.ca/legc/statutes/ferries.htm>.
47. <http://laws-lois.justice.gc.ca/eng/acts/M-9/FullText.html>.
48. <https://www.ontario.ca/laws/statute/90p05>.
49. <http://nslegislature.ca/legc/statutes/payequit.htm>.
50. A. Joulin, E. Grave, P. Bojanowski, T. Mikolov, *Bag of Tricks for Efficient Text Classification*, arXiv preprint arXiv:1607.01759, 2016.

