

Text Mining in Insurance: From Unstructured Data to Meaning

Diego Zappa¹✉, Mattia Borrelli², Gian Paolo Clemente³, Nino Savelli³

¹ Department of Statistical Sciences, Università Cattolica del Sacro Cuore, Milan, ² Reporting Actuary, Swiss Re, ³ Department of Mathematics for Economic, Financial and Actuarial Sciences, Università Cattolica del Sacro Cuore, Milan

Keywords: n-grams, natural language processing, policy premiums, text mining

Variance

Vol. 14, Issue 1, 2021

Every day, insurance companies collect an enormous quantity of text data from multiple sources. We present a strategy to make beneficial use of the large amount of information available in documents by exploiting natural language processing. After a brief review of the basics of text mining, we describe a case study in which, by analyzing the accident narratives written by the researchers of the National Highway Traffic Safety Administration of the U.S. Department of Transportation, we aim to extract latent information that can be used to fine-tune policy premiums. The process involves two steps. First, we classify the reports according to the relevance of their content to determine the risk profiles of the people involved. Next, we use these profiles to create new latent risk covariates for a company's ratemaking process.

Address for Correspondence: diego.zappa@unicatt.it

1. INTRODUCTION

Text mining is a knowledge-intensive process in which the dataset is represented by documents (Bash 2015). Originally introduced for descriptive purposes only, it has recently evolved to include methods able to classify documents according to their latent topic or to infer information about the "sentiment" of customers or the users of social networks. The approaches have been boosted by the evolution of both the computational efficiency of the algorithms necessary to analyze textual data and the technology needed to store information.

This work does not aim to explain in detail what text mining is. Relevant and comprehensive references are provided by Bash (2015) and Feldman and Sanger (2006). What we wish to show is how insurance companies can exploit this methodology to extract valuable information from un-

structured data. In insurance, as in many other contexts, there is a strong need for text classifiers because companies collect an enormous quantity of text data every day from multiple sources, such as customer feedback, claims adjuster notes, underwriter notes, police reports on accidents, medical records, surveys, e-mail messages, web documents, social media posts, etc. Text analysis could help companies to refine the following:

- Marketing campaigns
- Brand management
- Fraud detection
- Claims management and compensation
- Subrogation
- Relationships between the help center and clients
- Analysis of contract clauses

a Associate Professor of Statistics at faculty of Banking, Finance and Insurance. He teaches courses on statistics for insurance for both graduate and master students.

Correspondence: diego.zappa@unicatt.it

b Reporting Actuary at Swiss Re. He is responsible for the valuation of Life & Health business in Southern and Western Europe. Full Qualified Actuary, member of the National Order of Actuaries (ONA). Experience as consultant for Life and Health insurance companies.

c Associate Professor of Non-Life Insurance and Social Insurance at faculty of Banking, Finance and Insurance. He teaches courses on non-life topics for both under-graduate and graduate students. Full Qualified Actuary, member of the National Order of Actuaries (ONA). Experience as consultant for general insurance companies. Speaker at several international conference and author of several publications on international journals.

d Full Professor of Non-Life Insurance and Life Insurance at faculty of Banking, Finance and Insurance. He teaches courses on risk theory topics for both under-graduate and graduate students. Fully Qualified Actuary, member of the National Order of Actuaries (ONA). Large experience as consultant for general insurance companies and national/international insurance regulation. Speaker at several international conference and author of several publications on international journals.

To assess the possibilities offered by text mining, a UK insurer¹ recently introduced a motor insurance policy that granted discounts to first-time drivers who agreed to give the company access to their Facebook profiles² (recent changes to data privacy regulations do not allow the experiment to be repeated). Given access to the profiles, the company developed an algorithm that analyzed all user posts, likes, and GPS locations in order to perform a personality test. By analyzing each user's style of writing, the algorithm could uncover positive and negative traits. In this way, personality traits could be used as predictors of a customer's life and driving behavior and to assess criteria for determining eligibility for an insurance discount.

As suggested by this experiment, one of the potential competitive advantages of text exploration in insurance is the possibility of enriching customer risk profiles based on the standard structured "Company DB" customer database. Ideally, we are referring to a framework in which actuaries gather information from an unstructured "Cloud DB," fed by an external source (e.g., documents, web sources, etc.). The Cloud DB shares information with the Company DB through appropriate link variables that are able to connect the Cloud DB profiles and the company's customers, thereby allowing the aggregation of the two apparently disjointed data sources.

Text mining is a challenging research field. Issues include the need to analyze very large quantities of data, the unstructured nature of text data, and the complexity of finding keys to standardize language for inferential purposes. For example, in the case of insurance companies, the language in text data varies from colloquial to formal. In police reports or claims adjuster notes, we might find that the terminology used is often repetitive, with lexical structures that are sometimes fixed and predictable, but this does not apply to social media, which is always changing and therefore cannot be analyzed using standard methods.

This contribution fits into the big data paradigm (Bühlmann et al. 2016). Generally speaking, big data may be depicted as an unstructured, large, heterogeneous, and unstable data set that often hides latent relevant information not measurable through a standard sampling process. Big data may include documents; the tweets on the web; any social network; sentiment about the health of the economy, the status of a country, or a company; or the flow of documents produced during daily work (e.g., reports, recipes, phone calls, e-mails). Using an extensive description of a case study, we aim to show the possibilities offered by text mining to extract latent information that might be used by insurers to fine-tune policy pricing or to better assess customer risk profiles.

The paper is organized as follows. Section 2 briefly describes how a document can be analyzed using text mining. Section 3 shows how natural language processing (NLP) (Clark, Fox, and Lappin 2010) algorithms may be used to classify a document. Section 4 tests the efficacy of text mining in extracting latent information useful for insurance pricing and presents results of NLP applied to a collection of reports produced by the National Highway Traffic Safety Administration (NHTSA 2008) about accidents that occurred in the United States between 2005 and 2007.³ Conclusions are provided in the final section.

2. N-GRAMS AND THE PREDICTION OF WORDS

Documents are the focus of text mining. A collection of documents is known as a corpus. Once we have chosen the document unit, we need to set the granularity level of the analysis, after which we can analyze single characters, words, phrases, or even groups of phrases.

One of the most common tools used to analyze a document is the n -gram, a sequence of n words (or even of characters) obtained from a document. This is a sort of rolling window of size n : by moving this window by one position at a time, we obtain a list of new n -grams. For example, let us consider the sentence "The police stopped a vehicle without insurance." We can build different n -grams by varying the length n :

- $n = 1$ (unigram) returns {"The," "police," "stopped," "a," "vehicle," "without," "insurance"}
- $n = 2$ (bigram) returns {"The police," "police stopped," "stopped a," "a vehicle," "vehicle without," "without insurance"}
- $n = 3$ (trigram) returns {"The police stopped," "police stopped a," "stopped a vehicle," "a vehicle without," "vehicle without insurance"}

The choice of the size of the n -grams is linked to the complexity of the problem and to the scope of the analysis. Using a high value of n means incorporating more context into the units of the document, while a low n value means that the basic unit of data will be more granular.

An n -gram model can also be seen as a Markov chain of $m - 1$ elements, where m is the length of the sentence (e.g., number of words). Therefore, the outcome of the n -grams might be seen as a stochastic process that predicts words given a certain history or context. Since the model is a Markov chain, it does not take into account all the history of the previous words but considers only the most recent word/history. This is a simplification of how documents are

¹ See <https://www.telegraph.co.uk/insurance/car/insurer-trawls-your-facebook-profile-to-see-how-well-you-drive/>.

² See <https://www.telegraph.co.uk/technology/2016/11/02/facebook-blocks-admiral-app-that-priced-car-insurance-based-on-s/>.

³ At the time of this paper submission, the NHTSA database was available at <https://crashviewer.nhtsa.dot.gov/LegacyNMVCCS/Search>. It is beyond the scope of this paper to give details about it. We simply highlight some relevant goals of the NHTSA research: understanding the pre-crash environment; investigating the causes of the rollover problem; and describing the traffic environment in which the crash occurred, the behavior of the individuals and vehicle involved, the specific outcomes, any drug impairment, etc.

written because it does not focus on grammar rules to estimate words but only on their context. According to this framework, we can build the joint probability of words in a sentence using the chain rule:

$$P(w_1 w_2 \dots w_m) = \prod_{i=1}^m P(w_i | w_1 w_2 \dots w_{i-1}) \quad (2.1)$$

where $w_i, i = 1, \dots, m$ is the i -th word of a sentence of length m and equals $P(w_1)$ when $i = 1$. It is worth noting that we obtain the joint probability of a document by multiplying the probability of each word conditioned over all the previous words. For the sentence reported above, we obtain:

$$\begin{aligned} &P(\text{The police stopped a vehicle without insurance}) \\ &= P(\text{The}) \times P(\text{police}|\text{the}) \times P(\text{stopped}|\text{The police}) \\ &\quad \times P(\text{a}|\text{The police stopped}) \\ &\quad \times P(\text{vehicle}|\text{The police stopped a}) \\ &\quad \times P(\text{without}|\text{The police stopped a vehicle}) \\ &\quad \times P(\text{insurance}|\text{The police stopped a vehicle without}) \end{aligned}$$

To compute the probability, this approach requires counting and dividing the occurrences of each word for all the possible sentences. An example is:

$$\begin{aligned} &P(\text{vehicle}|\text{The police stopped a}) \\ &= \frac{\#(\text{The police stopped a vehicle})}{\#(\text{The police stopped a})} \end{aligned}$$

Unfortunately, this will lead nowhere because we will rarely have enough data to compute these probabilities consistently. To avoid this issue, an n -gram language model is used. The aim is to limit the context in which a word is used. Instead of using all the previous word history, we take only a subset of it. Therefore, equation (2.1) can be approximated using an n -gram model of order n according to:

$$\begin{aligned} P(w_1 w_2 \dots w_m) &= \prod_{i=1}^m P(w_i | w_1 w_2 \dots w_{i-1}) \\ &\cong \prod_{i=1}^m P(w_i | w_{i-(n-1)} \dots w_{i-1}) \quad (2.2) \\ &= \frac{\#(w_{i-(n-1)} \dots w_{i-1}, w_i)}{\#(w_{i-(n-1)} \dots w_{i-1})} \end{aligned}$$

where the number of words used to condition the probabilities is $n - 1$.

For instance, in the case of a bigram language model, each word probability will be conditioned only by the previous word. In our example we have:

$$\begin{aligned} &P(\text{The police stopped a vehicle without insurance}) \\ &= P(\text{The}) \times P(\text{police}|\text{The}) \times P(\text{stopped}|\text{police}) \\ &\quad \times P(\text{a}|\text{stopped}) \times P(\text{vehicle}|\text{a}) \times P(\text{without}|\text{vehicle}) \\ &\quad \times P(\text{insurance}|\text{without}) \end{aligned}$$

The use of an n -gram model allows us to reduce the number of cases in which occurrences are counted. Intuitively, as the order of the n -gram model increases, the n -gram frequency decreases. This is because every character sequence obtained from the model needs to be matched over the entire corpus (for example, “stopped a vehicle”).

The above description is an oversimplified approach to document analysis. In every text mining application, a key point is to understand how words and punctuation are joined together to express concepts. The definition of grammar rules is certainly crucial, but to assure a proper representation of data we need to simplify these issues.

	Terms												
	Are	Autonomous	Car	Cars	Claims	Change	Driving	Insurance	Look	Motor	No	Self	
Document 1	0	1	1	0	1	0	0	1	1	0	1	0	
Document 2	1	0	0	1	0	1	1	1	0	1	0	1	

Table 1. Numerical representation of documents

The “bag-of-words” model is an approach used to deal with the complexity of text data. It assumes that all the terms have the same importance in a document; that is, there is no distinction between different parts of speech (verbs, nouns, adjectives, etc.). We are not interested in their position in the text, and so they can be seen as a set of strings without meaning. The term *bag of words* is self-explanatory because it refers to a document as a bag of words that can be extracted without considering the order.

2.1. REPRESENTING TEXT DATA

To show how a document can be numerically represented, let us consider, for instance, the following two simple phrases (which will be generically called *documents*):

1. “Autonomous car insurance. Look, no claims!”
2. “Self-driving cars are set to change motor insurance radically.”

Next, let us consider a matrix on whose columns we place every unique term in the corpus and on whose rows we place the document IDs. When terms are present / not present in documents, the cells will take the value 1/0, thereby giving the matrix a non-negativity connotation.

This representation allows us to move from unstructured data (i.e., a collection of text documents) to structured data that can be analyzed by applying data mining methods. This numerical structure also has three important properties: sparsity, non-negativity, and side information. Punctuation also plays an important role in giving context to documents, especially when we have corpora composed of documents retrieved from the Internet, e-mails, and other sources of data in which informal language is typically used. We can add the frequency with which signs are present among classified documents. This will give us an indication of how frequently some punctuation signs are in documents.

To obtain a simpler numerical representation of the corpus, we can opt to reduce the number of features, e.g., prepositions, “generic” verbs (*be*, *have*), etc. In [Table 1](#) we have eliminated duplicated terms, punctuation, derived words, and so on. These are a few of the dimensionality reduction techniques applied in the following paragraphs to reduce the complexity of document representation.

2.2. TOKENIZATION

Tokenization is a process that consists of breaking up documents (classified as words, n-grams, or phrases) into elements called *tokens*, which are used as an input to text mining procedures.

Tokens might be similar but stylistically very different from one another: for example, some have capital letters while others have numbers or punctuation signs in them. In addition, some tokens have no meaning and thus would add no information to the documents' representation. To resolve these issues, we apply the most commonly used normalization techniques to the tokens to obtain the most meaningful subset of them. Examples of nonrelevant tokens are case, punctuation, numbers, hyperlinks, and spelling.

A relevant example is given by American and British English spelling. Despite the fact that both fall under the English language category, there are cases in which the spelling of certain words differs:

- *Analyze* – *Analyse*
- *Program* – *Programme*
- *Center* – *Centre*

In an information extraction problem, this causes difficulties: even if the words have the same meaning, the fact that they are spelled differently prevents us from obtaining matches during search operations. Obviously, the same holds for words that are used differently in the two dialects: *holiday* or *vacation*, *hire purchase* or *installment plan*, and so on.

Once the tokenization process is completed, we reduce the number of features to select the more significant ones by removing the so-called stop words. We can think of stop words as words that are useful for the syntactic construction of phrases but whose semantic value is important relative to the context in which they are inserted. For example, the use of the preposition *to* or *for* can change the interpretation of phrases (“I’ll take it to him” and “I’ll take it for him” have different meanings). Thus, in this case, removing the prepositions might affect the interpretation of the phrases and undermine the document’s meaning. To deal with stop words we use a stop words dictionary, which is a collection of all the terms that we consider unnecessary to our specific problem and that can thus be removed. Since stop words dictionaries are already available in all of the most common languages, a dictionary can be downloaded and then tweaked by adding or removing specific terms.

2.3. STEMMING AND LEMMATIZATION

Once the removal of numbers and punctuation has been completed, we can further simplify the features using a stemming or a lemmatization process. In English, as well as in other languages, verbs have different inflectional forms or suffixes to which they are related. For example:

- *see, saw, seen* → *see*
- *insurance, insurer, insure, insured* → *insur*

Stemming is merely a heuristic process that truncates the end of every word to reduce it to its common base, also called the *root*. In English a process that truncates the last characters of every term may have success in removing excess characters, leaving only the common base for the matching algorithms. The most common stemming algorithm for the English language is the one proposed by Porter (1980).

To improve the accuracy of the truncation process, we can use another technique called *lemmatization*. Basically, instead of defining rules to truncate words, we use a lemmatizer, which carries out a full morphological analysis to accurately identify the so-called lemma for each word. Lemmatization removes inflection suffixes and compresses words into the lemma, which we define as the canonical form of a term, stripped of most conjugation suffixes and transformations. Intuitively, we can think of a lemma as the word we look up in a dictionary when searching for a specific term.

To show the impact of these processes, we refer to a corpus of documents collected from Twitter. We have retrieved tweets regarding the insurance sector in Italy. The tweet-gathering process was spread across three months, making it possible to collect 15,909 tweets that included the word “*assicurazione*” (Italian word for “insurance”). The number of documents in this corpus is far greater than the number of documents used to explain how to prepare a text before the representation. Table 2 emphasizes how the processes outlined above are able to reduce the document complexity.

The original corpus of 15,909 documents is composed of 237,091 tokens. This number indicates the total number of words from the tokenization process. Note that after the stop words have been eliminated, the quantity of the tokens was reduced to 6% of the original data set.

2.4. VECTOR SPACE MODELS

Once the list of tokens is ready, we need to find a way to describe the whole document. One choice is to use vector space models (Lappin and Fox 2015, Ch. 16).

A document can be seen as a vector whose dimensions are given by the number of features (Turney and Pantel 2010). This representation is called a *document-term matrix*. By switching columns and rows, we obtain the so-called term-document matrix.

Placing documents into a multidimensional space requires an accurate coordinate system. Ideally, documents that are similar are also close to each other, while documents that are semantically different need to be distant. Using only term frequency is not enough to capture similarities if there are no informative words with a high occurrence rate. Hence, the tf-idf (term frequency-inverse document frequency) technique is typically used; this is a statistic that reflects the importance of terms in a corpus.

In general, a term is more important than others when it occurs multiple times in a document and when it is also rarer than other terms in the corpus. If we are analyzing a collection of documents from a division of an insurance company, we may find that terms such as *claims*, *underwriter*, *premium*, and *reserve* appear quite frequently. How-

Table 2. The impact of normalization and tokenization on a corpus of tweets

Process	Tokens
Original data	237,091
Tokenization	34,770
Token normalization	15,051
Stop words	14,847
Stemming	10,878

ever, their presence is common to every document in the collection, and so the value of the tf-idf statistic will not be as high as we would have expected while using only term frequency. Terms such as *windstorm*, *subrogation*, and *tsunami* may be less frequent than the previous ones, but they are more peculiar and rarer. Peculiarity and rarity are two qualities we look for when representing documents. To this end, the td-idf statistic combines two quantities:

1. Term frequency
2. Document frequency

2.4.1. TERM FREQUENCY

Term frequency is the frequency of the term t in the document d . It is the result of a simple tabulation process of the document text:

$$tf(t, d) = f_{t,d} = \frac{d T_t}{T} \quad (2.3)$$

where $d T_t$ is the number of occurrences of t in d , while T is the total number of terms in the document. One problem with term frequency regards long documents. Longer texts have a high probability that some words will be repeated, thereby leading us to conclude that the term frequency of these words is higher than what we would expect in shorter documents. To mitigate this effect, we can use a normalized term frequency that includes the use of a smoothing parameter based on a combination of a weighting parameter and a damping function: $ntf(t, d) = a + (1 - a) \frac{tf(t, d)}{tf_{\max}(t, d)}$, where $tf_{\max}(t, d)$ is the maximum term frequency of all the terms in the document d and a is the smoothing parameter that can assume values between 0 and 1 (typically around 0.4–0.5).

2.4.2. DOCUMENT FREQUENCY

Document frequency indicates the *inverse document frequency* of the term t in the collection of N documents. It is equal to the log of the ratio between the number of documents in the collection, N , and the number of documents N_t in which the term t appears:

$$idf(t, N) = \log \left(\frac{N}{N_t} \right) \quad (2.4)$$

This quantity describes how many rare terms there are in the corpus. Thus, in the presence of a rare term, equation (2.4) assumes a high value; otherwise, when many documents share the word, equation (2.4) takes on a small value.

2.4.3. TERM FREQUENCY–INVERSE DOCUMENT FREQUENCY

Ultimately, we can compute the tf-idf statistic for each document by simply multiplying the term frequency and inverse document frequency as follows:

$$tf - idf(t, N) = tf(t, d) \times idf(t, N) \quad (2.5)$$

tf-idf assumes a high value in the case of a high term frequency and a low frequency in the collection. To better explain how the tf-idf works, we use the following five documents:

1. “Autonomous car insurance. Look, no claims!”
2. “Self-driving cars are set to change motor insurance radically.”
3. “Insurers agree that, reputationally, their brand image is often made or broken during the claims process.”
4. “Industry executives broadly agree that advanced analytics can be used to drive value in insurance.”
5. “Progress has been slower in other lines of business, such as general liability, most specialty lines, and other elements of life insurance.”

We proceed by building a term-document matrix using the tf-idf statistic. [Table 3](#) shows how common stems such as *insur*, *claim*, and *agre* have low tf-idf values because of their presence in multiple documents. The statistic enables us to describe documents by considering peculiar terms and avoiding a representation based on common terminology. In fact, the stem *insur* has value 0—despite its high frequency, the highest among all the stems—meaning that it will not affect the representation of the corpus at all. Words that appear only once in a document will have a high tf-idf value and influence the position of the document in the multidimensional space.⁴

⁴ To further compress the table, for each document we may compute the average tf-idf statistic associated with every stem. A document with a high value means there are words in it that are on average more peculiar compared to others.

Table 3. Term-document matrix (tf-idf)

Terms	Documents				
advanc	0.000	0.000	0.000	0.211	0.000
agre	0.000	0.000	0.132	0.120	0.000
analyt	0.000	0.000	0.000	0.211	0.000
autonom	0.464	0.000	0.000	0.000	0.000
brand	0.000	0.000	0.232	0.000	0.000
broad	0.000	0.000	0.000	0.211	0.000
broken	0.000	0.000	0.232	0.000	0.000
busi	0.000	0.000	0.000	0.000	0.211
can	0.000	0.000	0.000	0.211	0.000
car	0.264	0.189	0.000	0.000	0.000
chang	0.000	0.332	0.000	0.000	0.000
claim	0.264	0.000	0.132	0.000	0.000
drive	0.000	0.000	0.000	0.211	0.000
element	0.000	0.000	0.000	0.000	0.211
execut	0.000	0.000	0.000	0.211	0.000
general	0.000	0.000	0.000	0.000	0.211
imag	0.000	0.000	0.232	0.000	0.000
industri	0.000	0.000	0.000	0.211	0.000
insur	0.000	0.000	0.000	0.000	0.000
liabil	0.000	0.000	0.000	0.000	0.211
life	0.000	0.000	0.000	0.000	0.211
line	0.000	0.000	0.000	0.000	0.422
look	0.464	0.000	0.000	0.000	0.000
made	0.000	0.000	0.232	0.000	0.000
motor	0.000	0.332	0.000	0.000	0.000
often	0.000	0.000	0.232	0.000	0.000
process	0.000	0.000	0.232	0.000	0.000
progress	0.000	0.000	0.000	0.000	0.211
radic	0.000	0.332	0.000	0.000	0.000
reput	0.000	0.000	0.232	0.000	0.000
selfdriv	0.000	0.332	0.000	0.000	0.000
set	0.000	0.332	0.000	0.000	0.000
slower	0.000	0.000	0.000	0.000	0.211
specialti	0.000	0.000	0.000	0.000	0.211
use	0.000	0.000	0.000	0.211	0.000
valu	0.000	0.000	0.000	0.211	0.000

2.5. PLACING DOCUMENTS IN A MULTIDIMENSIONAL SPACE

Having represented the documents in a space, it is also interesting to compare pairs of documents to understand document similarities.

When using vector space models, the closer the document vectors are, the more similar the documents are (see, for example, Pennington, Socher, and Manning 2014). To evaluate the position of the documents, we use a measure called *cosine similarity*, a statistic that also implies a length normalization process. Given only two documents, let $\mathbf{V}_1$ and $\mathbf{V}_2$ be the two column vectors containing the tf-idf val-

ues for the M terms in the collection. The cosine similarity is defined as

$$\cos(\theta) = \frac{\mathbf{V}_1 \cdot \mathbf{V}_2}{\|\mathbf{V}_1\| \|\mathbf{V}_2\|}$$

We report below the matrix computed by using the corpus of the five documents described above. Values show that this collection includes very dissimilar documents. The small size of the collection and the limited number of stems justify the results. In this case, the use of tf-idf has over-weighted the search for rare terms with high discrimination abilities.

$$\begin{bmatrix} 1.000 & 0.086 & 0.071 & 0.000 & 0.000 \\ 0.086 & 1.000 & 0.000 & 0.000 & 0.000 \\ 0.071 & 0.000 & 1.000 & 0.038 & 0.000 \\ 0.000 & 0.000 & 0.038 & 1.000 & 0.000 \\ 0.000 & 0.000 & 0.000 & 0.000 & 1.000 \end{bmatrix}$$

3. NATURAL LANGUAGE PROCESSING

After the very demanding process of document normalization, we are ready to find statistical methods that can “naturally” process language to extract information. We are looking for efficient statistical methods that “simulate” the process of our brain when we are reading or listening to people speaking.

The n-gram approach may fail in this task. It does not take into account the overall context of a phrase unless the size of the n-gram windows is very wide; however, when the size of the n-gram increases, the accuracy of estimated word probabilities decreases. Supervised machine learning methods are not able to deal with this issue without a time-consuming and sometimes painful calibration of the training phase.

NLP uses a mix of artificial intelligence, computer science, and computational linguistics to extract meaning from documents; recognize text; and, ultimately, model and shape text in order to compose original content (Bowman et al. 2015).

In insurance, NLP may be useful in making inferences about fraud (Kolyshkina and van Rooyen 2006; Stout 1998) or customer sentiment (Liu 2015; Ceron, Curini, and Iacus 2016), or (see our case study) in extracting latent information that cannot be measured by standard methodologies.

A revolutionary method to process language is the so-called word2vec. Word2vec is not just an algorithm but a class of algorithms introduced by Mikolov et al. (2013a, 2013b) that contains two different models: continuous bag of words (CBOW) and skip-gram. These algorithms are an example of deep learning (Wiley 2016). Deep learning is a branch of machine learning based on multilayered neural networks that involve both linear and nonlinear transformations. The algorithms in Mikolov, Chen, Corrado, et al. can be defined as two instances of a shallow neural network (with only one hidden layer). The importance of word2vec methods does not lie only in their predictive abilities. We show that these methods also enable us to provide a meaningful representation of words. Instead of trying to represent words in relation to their use in a document, the representation of words obtained with word2vec models concentrates on the meanings of the words themselves. For example, using the concept of a “window” in the n-gram models, we define the context as the words that surround the central word window itself. For example, in the string “deep learning methods,” the words *deep* and *methods* compose the context of the central word *learning*.

3.1. CONTINUOUS BAG OF WORDS

The continuous bag of words (CBOW) model was introduced by Mikolov et al. in 2013. We start by describing the




Figure 1. A simple CBOW model with only a one-word context

simplest version based on a one-word context. In this version, the model predicts one word, given only one context word, to form a bigram. For example, given the word *insurance*, the model will try to predict the immediately following word.

To achieve this aim, we train the neural network based on the structure reported in Figure 1 (see also Rong 2014).

For each word we are interested in, we use a vocabulary of size V from the corpus of documents we are analyzing. Each word will be represented by one-hot vector $\mathbf{x}$ of dimension V . Next, we shrink the V dimensions into a smaller space of size $N < V$ (hidden layer). The transition is made by using the $V \times N$ weight matrix. Each node h of the neural network is defined by the simple linear activation transformation $\mathbf{h} = \mathbf{x}^T \mathbf{W}$. The transition from the hidden layer to the output layer is performed by using the $N \times V$ matrix $\tilde{\mathbf{W}}$.

These matrices provide the score $\mathbf{u} = \tilde{\mathbf{W}} (\mathbf{x}^T \mathbf{W})^T$ computed for every word in the vocabulary.

Let $\mathbf{x}_I$ be the input vector (for instance, the word *insurance*), and suppose we want to maximize the probability $P(\mathbf{x}_O | \mathbf{x}_I)$, where $\mathbf{x}_O$ is the word we expect to find next to $\mathbf{x}_I$ (say, *policy*). Once we have computed an initial score $\mathbf{u}$ by using the stochastic gradient descent (Guelman 2012), it is possible to update the weights of $\mathbf{W}$ and $\tilde{\mathbf{W}}$ until convergence. The scores $\mathbf{u}$ reveal what words within V are closer to $\mathbf{x}$.

One of the advantages of the word2vec algorithms is that $\mathbf{W}$ is also used to obtain a denser and more meaningful representation of the words included in the input layer. From a practical point of view, the N dimensions of $\mathbf{W}$ represent the word similarity, i.e., the N words that are most likely to be close to the words used in the training phase. Using the notion of cosine similarities described in Section 2.5, we can compute the similarities between words and obtain additional information about terms and their semantics.

Another property of word2vec models regards “clustering themes,” or the ability of these models not only to find similar words but also to discover a class of items relating to similar topics. These models do not typically perform topic modeling, but practical evidence shows that a topic can be assigned to similar word vectors. Therefore, terms that ap-

pear in the same context can be interpreted as observations of a latent topic distribution over words.

Hence, word2vec models can be useful in discovering semantic relationships between words. This is one of the most important features of this class of algorithms, allowing us to overcome the simple similarity between words. Similarities between groups of words, known as *linguistic regularities* (Mikolov et al. 2013b), can indeed be detected.

4. INSURANCE CASE STUDY

We analyzed 6,949 police reports (written by National Motor Vehicle Crash Causation Survey researchers sent to crash scenes) on accidents in the United States between 2005 and 2007. The researchers listened to police scanners; interviewed everyone on the scene; and collected structured data such as the date and time of the accident, weather and road conditions, driver use of medications, driver use of a cell phone, and so on. Then the researchers wrote a brief report, describing the accident in a maximum of 1,200 words.

In the literature, an example of applying text mining to this database was proposed by Borba (2013, 2015), who extracted dichotomous variables (called *flags*) from documents to study the impact on accidents of weather conditions, the status of the cars, the use of mobile phones while driving, the dynamics and the locations of accidents, and the driver's condition (presence of driver fatigue, use of a cell phone, medications, drugs, alcohol, prescriptions).

Our purpose is to classify documents not only according to “accident-specific” keywords. In particular, we look for information that can uncover new risk covariates that might be used to fine-tune policy pricing or to improve customer risk profiles. To achieve this aim, a straightforward application of n-grams may be misleading. For example, a key issue is the profiling of drivers on the basis of the substances reported in the police reports, if available. The difficulty here is that, on the one hand, the words *medication*, *drug*, *prescriptions*, and *alcohol* are not always explicit in police reports, and on the other hand, some substances could belong to more than one category. In addition, some verbal expressions such as “he was not taking medication” or “the driver was aggressive but the BAC results were negative” could easily lead to the wrong classification. Therefore, without a supervised check by an expert, the n-gram approach could fail. To avoid these pitfalls, we have improved the analysis of documents by training an NLP system, based on a word2vec algorithm, to “automatically” classify the substances. In what follows we will focus mainly on this case, but the same procedure can be applied to other contexts, such as driver-related factors (distracted, using a cell phone, etc.).

The procedure we used to analyze the data set is based on the following seven steps:

1. Obtain the data set
2. Apply a text normalization procedure
3. Select substances (word2vec)
4. Tag parts of speech
5. Filter off-topic cases

6. Create the n-gram structure and flags
7. Apply the prediction model

4.1. OBTAINING THE DATA SET

Unfortunately, at least at the time of the present paper's submission, the data set containing the accident narratives cannot be directly downloaded in a user-friendly format. The data set is published online, where it can be queried using the search form created by the NHTSA. Each case is stored in an XML file that can be downloaded. The peculiar trait of XML is that it can encode any piece of a document. Indeed, as it is a markup language, it allows tags to be created to host new and different types of information. An example is provided in [Figure 2](#), which shows how information regarding driver fatigue during an accident can be encoded.

Given the structure of XML, each case is stored differently. The quality of the information collected depends on how detailed the report is. Each request made from the interface is passed to the server containing the database using the HTTP GET method. The result of the GET request is a webpage that reads the underlying XML code of the selected case and displays it clearly.

Once we have downloaded the XML file of all the cases, we need to parse them to understand the information encoded in the XML language. The objective of the procedure is to extract the information encoded in sections or tags and store it in a structured form. [Figure 3](#) is an example of an accident narrative we have extracted.

4.2. PERFORMING TEXT NORMALIZATION

To simplify the text mining approach, we divided each document into phrases, using the dot as a string separator. The corpus size increases from the initial 6,949 document cases to 198,480 phrases. The next step in the normalization process consists of polishing up the phrases by removing punctuation signs, symbols, and numbers. We also removed the stop words during the normalization process.

For example, the next-to-last phrase (in bold in [Figure 3](#)) will be changed to “driver was crash three weeks prior crash was taking medications ibuprofen vicodin ultram.”

4.3. SELECTING SUBSTANCES

Once we have normalized the text, we focus on the identification of substances (over-the-counter and prescription medications, narcotics, and alcohol) that will be used to filter the on-topic cases that will be analyzed later. Instead of creating a list of possible substances, we opted to use word2vec methods. We were indeed able to train a neural network to predict the context words given a central word as input. This representation allows us to grasp word similarities and to uncover thematic clusters in the text.

In our particular application, we started by finding the terms that had greater cosine similarity to common illegal substances (cocaine, heroin, and so on). We filtered out the corresponding names and collected them in a list. This list was used to calibrate the classification of other drugs, ille-

```

1 <FATIGUE xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
2   <TRIP_LENGTH formCat="2">0.12</TRIP_LENGTH>
3   <TIME_ELAPS_DRIVING formCat="2" sasCode="1" value="1" attrCat="-9999" attrCatStr=
4     "">&lt; half hour</TIME_ELAPS_DRIVING>
5   <TRIP_START_TIME formCat="2">0555</TRIP_START_TIME>
6   <SLEEPSTART_TIME formCat="2">2000</SLEEPSTART_TIME>
7   <SLEEPEND_TIME formCat="2">0430</SLEEPEND_TIME>
8   <LASTSLEEP formCat="2">8.5</LASTSLEEP>
9   <SLEEP24HRS formCat="2">8.5</SLEEP24HRS>
10  <HOURS_AWAKE formCat="2">1.5</HOURS_AWAKE>
11  <AVGSLEEPINT formCat="2">8</AVGSLEEPINT>
12  <SLEEPROTATE formCat="2" sasCode="1" value="1" attrCat="" attrCatStr="">No</
13    SLEEPROTATE>
14  <WORKSHORTEST formCat="2">2</WORKSHORTEST>
15  <WORKLONGEST formCat="2">9.5</WORKLONGEST>
16  <WORKAVG formCat="2">4.5</WORKAVG>
17  <WORKTOTAL formCat="2">18.5</WORKTOTAL>
18  <DRIVER_FATIGUE formCat="2" sasCode="2" value="2" attrCat="-9999" attrCatStr="">
19    Driver not fatigued</DRIVER_FATIGUE>
20 </FATIGUE>

```

Figure 2. Example of an XML file

This single vehicle crash occurred just before dawn during dark hours on a level, dry, left-curve asphalt six lane divided interstate with guardrail's in the median and outside the right shoulder. This interstate runs east and west at this location and has three lanes in each side with no traffic controls. Posted speed limit 105 kmph (65 mph).

Vehicle one (V1), a 1997 Ford Escort sedan driven by a 31 year-old female, was traveling westbound in lane two. V1 drove into the number one lane and off the right side of the road onto a paved shoulder. V1 then impacted a guardrail with its front and rebounded back onto the road re-entered all three lanes and ran off the left side of the road impacting another guardrail with its front. After the second impact, V1 rotated counter clockwise and came to rest on its wheels facing southeast in the median. V1 was towed from the scene.

The critical pre-crash event for V1 was coded: this vehicle traveling, off the edge of the road on the right side. The critical reason for the critical event was coded to V1 as a driver related factor: internal distraction. The driver stated that she was looking to the right and reaching for her cigarettes and lighter on the front seat. **The driver was in a crash three weeks prior to this crash and was taking the medications: ibuprofen, vicodin and ultram.** For this trip she was on an errand going from a coffee shop to her mother's house.

Figure 3. An example of an accident narrative

gal and not. We proceeded to refine the list by searching for the term closer to the term *medication* or *narcotic* to identify substances recorded during the accident investigations. In this way, 290 (legal and illegal) substances were “automatically” detected.

For instance, in Figure 4 we provide a visual representation of the results obtained using the similarity scores between *medication* and *narcotic*. Given a word in the list, the higher of the two scores allows us to define the classification as belonging to either the “narcotics” or “medications” group. We can see how the substances are correctly separated based on their similarity scores. For example, the word *medication* has a medication score of 1, since the two words are equal. The same word has a 0.6 narcotics score; this means that, within the corpora of police reports, there




Figure 4. Drug vectors plotted by their similarities to the terms *narcotic* and *medication*

are instances in which the word *narcotic* is located near the word *medication*, contributing to the high narcotic score for the word *medication*. In general, the illegal drug names are found closer (but not exclusively) to the word *narcotics*, while prescription substances are more likely (but not exclusively) closer to the word *prescription*.

Based on the texts used for the training, if a word is placed above/below the diagonal it means that it is more similar to the word placed on the y/x axis.

It is worth mentioning that the scores computed by the word2vec model are *not* representative of the real similarities between any kinds of drugs. Thus, we should not interpret the scores as a direct representation of the real classification of each substance. Instead, word2vec creates a word embedding for each term of the specific corpus. The word embedding depends on the documents used in the training process, and thus substance discrimination reliability requires accurate and precise documents.

The previous result can be used in the reverse order: in general, even if a word is not mentioned in a document but it occurs often in other documents when other medications/substances/narcotics are mentioned, then it inherits some degree of similarity to the reference group in question. For example, the word *xanax* never appears closer to the term *narcotics*, but it occurs frequently in phrases containing the term *heroin*. Therefore, based on the transitive property, the term *xanax* will be considered more likely to belong to narcotics than to the medications group. The extreme case is when no substance is reported: by using word similarity we can classify a document into one of the categories we have chosen.

The bottom line for using the word2vec model is that we have created an efficient method to discriminate between illegal and legal substances *with little to no effort*. In fact, the ideal word2vec model is one that uses millions or billions of documents (for example, the entire Wikipedia library) to understand word similarities or thematic clustering. However, we are pleased to see how, with just under 7,000 documents, we were still able to reach our goals.

4.4. TAGGING PARTS OF SPEECH

One of the advantages of analyzing text is that we have the possibility to use verbs to detect actions that we are interested in. In order to detect verbs, we use a part-of-speech tagger, a function that assigns a grammar role to each word in a phrase. Since the process used to assign the tag is based on a set of heuristic rules, the procedure can be faulty, especially when new words are presented in input (this happens often with the names of certain drugs). However, for known and common words we were able to extract what we were looking for: the gerund of verbs to understand if an action was in progress or not (“the driver was calling” versus “the driver called”).

4.5. FILTERING OFF-TOPIC CASES

To identify topics, we compute n-grams. Building a database containing all the n-gram information is very expensive from a computational point of view. We estimated that, in our case, this would have led to a structure containing more than 10 million unique expressions. To reduce this number, we needed to find a way to filter some of the phrases created by the normalization procedure. One way to accomplish this task was to search phrases that contain words that are crucial to our application. Using the list of substances created with the word2vec model, we were able to filter the phrases, thereby reducing the database to a more manageable size. The filtering process is quite computationally intensive. However, 97.6% of phrases are labeled as off-topic (see [Table 4](#)).

[Table 4](#) shows that 4,270 accident narratives can be labeled as off-topic since none of their phrases mentioned the presence of drugs or alcoholic beverages. Drugs or drinks are mentioned in 38.5% of the total narratives. We still cannot say if those factors are really present (or not present) because we could have stumbled upon a false pos-

itive—for example, cases that mentioned drugs or drinks only to state that they were *not* present. To overcome this pitfall, we implemented an n-gram strategy to understand the text. Until now, we have focused on a unigram model that considers only whether specific terms are mentioned in the text.

4.6. CREATING THE N-GRAM STRUCTURE AND FLAGS

Building the n-gram structure is only the first step toward extracting information from the accident narratives. To be able to understand the expressions written in the documents, we created a series of flags that exploit the structure we have created. These flags were then used to indicate the presence of drugs, illegal narcotics, and alcohol in the accident narratives. Four different n-grams were investigated.

4.6.1. UNIGRAM

We created four different flags:

- Flag 1 indicates a reference to medications in the text.
- Flag 2 indicates a reference to prescriptions in the text.
- Flag 3 indicates a reference to narcotics in the text.
- Flag 4 indicates a reference to alcohol in the text.

As described above, focusing only on single words does not allow us to obtain the full context. We are not able to capture negations and other more complex text expressions.

4.6.2. BIGRAM

Hence, we expanded our focus to pairs of words. Exploiting a bigram, we tried to detect two particular patterns:

- Negations: A flag is created to indicate the presence of a negation (no, non, not, none) followed by the name of a drug or an alcoholic beverage.
- Actions: A flag is created to indicate the presence of an action (verb with *-ing*) followed by the name of a drug or an alcoholic beverage. We are interested in the gerunds because they are commonly used to express real, fixed, or completed actions.

Transitioning from words to pairs of words has many implications. On the one hand, we could refine our search for substances by accounting for particular combinations of terms. On the other hand, we needed to deal with the ceiling represented by the complexity of the analysis. In the previous step, we had to analyze only 3,261 unigrams, while 16,692 bigrams were generated.

To take an example, using the phrase of the accident narrative highlighted in bold in [Figure 3](#) produces the following bigrams:

1. “driver was”; 2. “was crash”; 3. “crash three”; 4. “three weeks”; 5. “weeks prior”; 6. “prior crash”; 7. “crash was”; 8. “was taking”; 9. “taking medications”; 10. “medications ibuprofen”; 11. “ibuprofen Vicodin”; 12. “vicodin ultram”

Table 4. Recap of the filtering process

	# Cases	# Phrases
Filter	2,679	4,583
Off-topic	4,270	193,870
Total	6,949	198,480

Only a fraction of these bigrams are effectively used in the analysis.

4.6.3. TRIGRAM

Regarding trigrams, there are 24,900 combinations of word triplets—49% more items than are produced with the bigram structure and almost 7 times more than with the unigram structure.

Considering triplets instead of pairs of words further increases the complexity of the analysis, since numerous expressions can be encoded in a larger window. Aside from the classical controls already implemented, we further refined our searching algorithm by taking into account the following:

- **Past continuous tense:** We now track every triplet that contains a reference to this verbal form. Typically, the past continuous is used to express past events that were prolonged in time, or that continued before and after other actions. A classic example is the triplet “was doing drugs.” Notice that in this specific triplet a flag indicating the presence of an action was already implemented in the unigram phase. However, the bigram “doing drugs” does not allow us to exclude the fact there was a negation before the verb *was*. Therefore, to refine the search for enduring past actions we implemented this check, which also has the benefit of removing any doubts regarding negations or other contradictory expressions.
- **Past tense:** We enhanced the check regarding past tenses by identifying triplets that contained the past declination of the auxiliary verbs *to be* and *to have*. This phase allowed us to consider expressions such as “was on medication,” “had taken Percocet,” etc.
- **Negative and positive tests:** We can track the presence of toxicology or blood alcohol content (BAC) tests in the narrative descriptions of the accidents. We are dealing only with phrases that contain at least one occurrence of a substance, so it is fair to assume that any mention of the terms *positive* and *negative* is related to some sort of intoxication test. Exploiting this assumption, we proceed to detect triplets such as “cocaine negative results,” “negative results drug,” “alcohol tested positive,” and “positive opiates cocaine.”

It is clear that triplets offer a broader view on the context of phrases. Just like a child learns to read, the algorithm takes baby steps toward a good comprehension of the meaning of documents.

As emphasized above, we are not declaring a comprehensive list of all the possible expressions that can point toward the correct identification of substances during accidents, which would take too much time. We are only trying to characterize a small set of rules that are at least able to understand the general context in which the crashes happened.

4.6.4. FOURGRAM

The design of the last batch of checks implemented is based on fourgrams and will be used to detect particular samples of false-positive patterns. The degree of complexity increases again, from the previous 24,900 instances with the trigram to 27,642 with the fourgram. We use this structure to detect the following patterns:

- **Prescription medication:** In the accident narratives, it is common to come across this pattern. Reading the narrative, we know that it refers to a prescription; however, in the previous steps of the analysis we may have assigned contradictory flags to it. Because of the presence of the word *medication*, the algorithm may be deceptive and falsely attribute this pattern to the wrong class. By searching for the terms *prescription* and *medication* in a broad context, we were able to contextualize the terms and correctly classify them in the “prescription” category.
- **Prescription sunglasses:** This is another pattern (like the previous one) that can be difficult to deal with. Identifying the presence or absence of prescription sunglasses (lenses or other eyewear items) is not one of our goals; therefore, we need to find a way to eliminate these mentions from the analysis. In order to exclude these patterns, we can build a flag that marks them as “off-topic,” so that they will have little importance in the overall picture.

Realistically, many patterns may escape our rule-based algorithms, but we are counting on the fact that the flags created in the unigram structure act as a safety net. This concept has its foundation in the filter phase executed upstream. Having retained all the phrases that include at least one occurrence of the presence of the substance, we are 100% sure that in the relative cases there is at least one substance involved. By trying to exclude from the picture all the possible negative expressions, we are only exposing ourselves to the possibility that some patterns for positive detection may be ignored. Nevertheless, by having structured flags that act only on the mere presence of substance names, we are hedging our exposure with regard to missed opportunities.

Finally, we have obtained a list of 24 text-related flags (e.g., “she/he was/was not taking medications/prescriptions/drugs,” “she/he was positive/negative to alcohol test,” “she/he was/was not wearing sunglasses,” “she/he was/was not calling someone,” etc.), where each flag is used to detect a particular pattern in the n -gram structure. Therefore, a flag (0/1 for the absence/presence of a particular pattern) will be associated with every phrase.

4.7. APPLYING THE PREDICTION MODEL

The flags can be used to investigate the effect of a pattern on accidents. For example, the evidence of substances (legal and illegal) and their impact on motor vehicle crashes can be analyzed by studying the percentage of injuries (both minor and more serious, deaths included); we expect the percentage to be higher in accidents where these factors were present than in accidents where they were not present.

By using a logistic regression with an elastic net regularization (Friedman, Hastie, and Tibshirani 2010; Zou and Hastie 2005) to perform regression and variable selection by linearly combining both L1 and L2 penalties, we have estimated the probability of injury in accidents in which at least one substance was present. A justification of the need for this methodology is that text mining problems can be treated as $p > n$ problems, where p represents the number of the predictors and n represents the number of the records. Moreover, some variables may be highly correlated, and methods such as lasso regression will tend to pick only one of them and ignore the others, without preference as to which one has been selected. This is a limit that we do not find in ridge regressions. If the variables are highly correlated when the task at hand is described by an $n > p$ situation, it has been shown that the ridge regression always outperforms the lasso, but it does not allow unnecessary predictors to be filtered. The limits of the previous models do not apply to the elastic net algorithm, a regression analysis method that combines both the lasso and the ridge regressions, overcoming their shortcomings. In our case, briefly consider the logistic model

$$\ln \frac{\Pr(\text{Injury}|\mathbf{x})}{1 - \Pr(\text{Injury}|\mathbf{x})} = \beta_0 + \mathbf{x}^T \beta$$

Let $y_i = 1$ if the i -th report returns at least one injury or $y_i = 0$ in the contrary case. Let $p(\mathbf{x}_i) = \Pr(y_i|\mathbf{x}_i)$. Then the aim is to maximize the penalized log likelihood

$$\max_{\beta_0, \beta} \left[\frac{1}{n} \sum_i (y_i \ln p(\mathbf{x}_i) + (1 - y_i) \ln (1 - p(\mathbf{x}_i))) - \lambda P_\alpha(\beta) \right]$$

where $P_\alpha(\beta) = (1 - \alpha)\|\beta\|_{L_2}^2 + \alpha\|\beta\|_{L_1}$, $0 \leq \alpha \leq 1$ and $\lambda \geq 0$.

Figure 5 shows the probability of injury for all the factors (flags) we considered, ranging from 77% for alcohol to 83.4% for illegal narcotics; the probability is always greater than the benchmark computed using only the off-topic




Figure 5. Probability of injury due to the presence of substances

cases⁵ (see also Borba 2013 for similar results obtained using logistic regression).

All factors present results that comply with our initial prediction. It is more likely that injuries will occur in crashes where at least one of the drivers had taken drugs or was under the effect of narcotics or alcohol. It is interesting to note that the probability of injury is higher in the presence of medications (78.78%) than in the cases in which prescription medicines are involved (77.7%). It is difficult to say whether the two percentages are statistically significantly different. Medications are available over the counter, so they are easier to access and should not affect common activities such as driving. However, it is possible that abuse is higher because of how easily they can be purchased. In addition, medications include antibiotic and other common drugs that are used to treat seasonal illnesses. Therefore, it is possible that in particular seasons more people are effected by allergies, or simply colds, and in order to treat them they use medications. This increases the number of individuals driving under the effect of medications, thus increasing the probability of injury, therefore biasing its real value. In general, this effect is difficult to assess because it is not available in the reports and it cannot be easily measured.

The results regarding alcoholic beverages and illegal narcotics are in line with expectations: it is well known that drug-impaired driving is a crime and can have some serious consequences. The same holds for alcohol-impaired driving, but we need to make a distinction. In the case of illegal narcotics, it may not be easy to detect DWI (driving while intoxicated) cases immediately after a crash. Tests for narcotics are lab tests that may take time. What can be recorded is the possession or possible use of narcotics. We need a list of substances and a series of flags that test their absence/presence at the time of the crash.

⁵ Recall that off-topic cases are those reports in which medications, prescriptions, and alcohol were not mentioned (4,270 cases).

For alcoholic beverages it gets trickier. Determining DUI (driving under the influence) might be difficult because of the different laws on the matter (in the United States, it is possible to drive with a BAC of 0.08%, while for drivers under 21 the zero-tolerance law imposes stricter limits, between 0.01% and 0.05%) and because sometimes the description includes the generic statement that alcoholic beverages have been found in the car. Thus, even if a document mentions that a driver had been drinking, this does not mean that he or she was driving under the influence of alcohol. We do not aim to determine the impact of alcohol infractions (DUI cases) but rather the general influence this factor can have on the way we drive. The results speak clearly: the mere mention of alcoholic beverages in an accident description is sufficient to detect crashes in which the probability of injury was higher than the benchmark.

4.8. FINAL STEP

Insurers may consider this data useful in the process of claim evaluations for the following reasons:

1. Detecting the presence of substances has implications on multiple levels: claim triage, subrogation opportunities, and decisions on whether to renew contracts.
2. Claims adjusters can exploit this information to investigate the circumstances in which crashes occurred, thereby obtaining a more reliable basis on which to begin settlement negotiations.
3. Actuaries can use this information to refine their pricing procedures.

Typically, actuaries base their policy premium calculations on a so-called Company DB, a data set containing details about their customers. Intuitively, a company cannot ignore what is taking place outside its pool of clients, and to correctly price new and existing products it needs to consider phenomena that affect the entire population of non-customers and customers.

By using both structured and unstructured data extracted from the accident narratives, we can build a so-called Cloud DB. This data source should have some classes in common with the Company DB: age, gender, ethnicity, license state, etc. We can interpret these variables as identifying a risk profile that is common to the sources of information. This allows actuaries to use this text-related information in the definition of premiums, opening the door for more extensive exploration and use of documents in order to discover new risk covariates. When writing a policy, we do not aim to detect which individuals are taking narcotics but to explain how this information can be potentially used in order to modify the premium ratings of policyholders.

The rationale for this process is not new in ratemaking. First, we compare the risk of those who have been involved in accidents with those who have not been involved, given a set of covariates. Then we extend this risk to all the customers, mitigating the effect for those who were safe drivers but not reducing their risk to zero. The result is an integration into the premium model algorithm of a latent




Figure 6. Distribution of injury by gender and age

variable that represents the potential risk of driving under the influence of drugs.

On a more detailed level, our application could analyze the characteristics of every driver involved in the accidents. Using the fields derived from the XML cases, we were able to build a database including information such as the following for 12,300 drivers:

- License origin, status, endorsements, and restrictions
- Age, gender (see [Figure 6](#)), height, weight, and illness
- External factors: emotions and stress
- External factors: in a hurry, fast following, and traffic
- External factors: distractions and conversation
- Season, day of the week, and time of the accident
- Year of production and make of the vehicle
- Emergency transport in ambulance
- Severity of the accident

We can match these profiles from the police reports with the profiles of the company's own database to add new risk covariates (e.g., probability of using substances or, for some categories, of being under the influence of external factors). To sum up these factors, in practice we built a model that classifies the accident outcome of each driver to predict whether drivers were injured or not during motor vehicle accidents. The variables we used as predictors were those extracted from the XML cases in a structured form, including the four text-related flags referring to medications, prescriptions, illegal narcotics, and alcohol. To test the classification accuracy for each driver, we used an elastic net regularized regression.

[Figure 7](#) reports the receiver operating characteristic (ROC) of our classification procedure, which can be interpreted as the ability of the method to correctly identify the classification categories. In our case, we obtained an area under the ROC of 72.25%, a good measure for a classifier that needs to take into account multiple variables and that uses predictors obtained directly from the text.

5. CONCLUSIONS

We provided a case study that shows how valuable text mining can be for actuarial procedures in pricing models.




Figure 7. Receiver operating characteristic of the method

ROC = receiver operating characteristic; TRP = true positive rate; FPR = false positive rate.

However, the analysis barely scratched the surface of the possible applications for artificial intelligence methods in actuarial and insurance contexts in the near future. The aim was to create text-related risk covariates that may affect ratemaking.

The same process can be reproduced using different contexts. New applications are emerging in terms of analyzing the GPS coordinates of social media posts to obtain proxy measures for the exposure to accidents. Researchers have found a way to use Twitter posts to estimate drinking-while-tweeting patterns to detect particular string patterns (features). This is another example of how it is possible to create new risk covariates linked to data mining of unstructured text documents.

The steps of the algorithm we propose are based on a data-driven process. In each step, the algorithm becomes “more conscious” of the context in which it is operating. It first detects the substances that are present in the doc-

uments and then proceeds to learn about the expressions that are present in the text. Each step allows us to create specific flags that increase the understanding of the text. The algorithm exploits concepts derived from different areas, with the specific aim to maximize interpretability thereby avoiding the use of black boxes.

Future challenges involve the application of these methods to business activities using topic modeling. The insurance industry is well suited to the implementation of these advanced models because of the large quantity of data available to carriers. Text mining represents one example of unconventional methods that could be used to refine insurers’ traditional approach to measuring the risk profiles of their policyholders, thus opening the way for a wave of innovation. Appropriate application of these methods could provide the foundation for extending a data mining approach to big data projects.

ACKNOWLEDGMENTS

This work was sponsored by the Casualty Actuarial Society (CAS). The authors wish to thank the editor and two anonymous referees for their careful reading and the suggestions that helped to improve the quality of the paper. The authors wish to give special thanks to David Core, director of professional education and research at CAS, and Jody Allen, professional education and research coordinator, for their support.

Finally, the opinions expressed in this paper are solely those of the authors. Their employers do not guarantee the accuracy or reliability of the contents provided herein nor do they take a position on them.

Submitted: May 31, 2018 EDT. Accepted: January 09, 2019 EDT. Published: September 21, 2021 EDT.

REFERENCES

- Bash, E. 2015. *The Text Mining Handbook*. Vol. 1. Cambridge, UK: Cambridge University Press.
- Borba, P.S. 2013. "Predictive Analytics, Text Mining, and Drug-Impaired Driving in Automobile Accidents." Seattle, WA: Milliman. <http://us.milliman.com/insight/pc/Predictive-analytics--text-mining--and-drug-impaired-driving-in-automobile-accidents/>.
- . 2015. "Demonstrating the Value of Text Data for Predictive Modeling." Presented at the Southern California Casualty Actuarial Club Spring 2015 Meeting, Pasadena, CA, May 28. <https://www.casact.org/community/affiliates/sccac/0515/Borba.pdf>.
- Bowman, Samuel R., Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2015. "Generating Sentences from a Continuous Space." Preprint: arXiv:1511.06349. <https://doi.org/10.18653/v1/k16-1002>.
- Bühlmann, Peter, Petros Drineas, Michael Kane, and Mark van der Laan, eds. 2016. *Handbook of Big Data*. New York: CRC Press. <https://doi.org/10.1201/b19567>.
- Ceron, Andrea, Luigi Curini, and Stefano Maria Iacus. 2016. "iSA: A Fast, Scalable and Accurate Algorithm for Sentiment Analysis of Social Media Content." *Information Sciences* 367–368 (November):105–24. <https://doi.org/10.1016/j.ins.2016.05.052>.
- Clark, Alexander, Chris Fox, and Shalom Lappin, eds. 2010. *The Handbook of Computational Linguistics and Natural Language Processing*. Malden, MA: Wiley-Blackwell. <https://doi.org/10.1002/9781444324044>.
- Feldman, R., and J. Sanger. 2006. *The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data*. New York: Cambridge University Press.
- Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2010. "Regularization Paths for Generalized Linear Models via Coordinate Descent." *Journal of Statistical Software* 33 (1): 1–22. <https://doi.org/10.18637/jss.v033.i01>.
- Guelman, Leo. 2012. "Gradient Boosting Trees for Auto Insurance Loss Cost Modeling and Prediction." *Expert Systems with Applications* 39 (3): 3659–67. <https://doi.org/10.1016/j.eswa.2011.09.058>.
- Kolyshkina, Inna, and Marcel van Rooyen. 2006. "Text Mining for Insurance Claim Cost Prediction." In *Lecture Notes in Computer Science*, edited by G.J. Williams and S.J. Simoff, 3755:192–202. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/11677437_15.
- Lappin, Shalom, and Chris Fox, eds. 2015. *The Handbook of Contemporary Semantic Theory*. 2nd ed. Malden, MA: Wiley-Blackwell. <https://doi.org/10.1002/9781118882139>.
- Liu, Bing. 2015. *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/cbo9781139084789>.
- Mikolov, T., K. Chen, G. Corrado, and J. Dean. 2013a. "Distributed Representations of Words and Phrases and Their Compositionality." *Nips*, 1–9.
- . 2013b. "Efficient Estimation of Word Representations in Vector Space." *Proceedings of the International Conference on Learning Representations*, 1–12.
- NHTSA. 2008. "National Motor Vehicle Crash Causation Survey: Report to Congress." DOT HS 811 059. <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811059>.
- Pennington, J., R. Socher, and C.D. Manning. 2014. "GloVe: Global Vectors for Word Representation." *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1532–43.
- Porter, M.F. 1980. "An Algorithm for Suffix Stripping." *Program* 14 (3): 130–37. <https://doi.org/10.1108/eb046814>.
- Rong, X. 2014. "Word2vec Parameter Learning Explained." Preprint: arXiv:1411.2738.
- Stout, Nancy. 1998. "Analysis of Narrative Text Fields in Occupational Injury Data." In *Occupational Injury: Risk, Prevention and Intervention*, edited by A.-M. Feyer and A. Williamson, 15–20. London: Taylor & Francis. <https://doi.org/10.1201/9780203212493.ch2>.
- Turney, P. D., and P. Pantel. 2010. "From Frequency to Meaning: Vector Space Models of Semantics." *Journal of Artificial Intelligence Research* 37 (February):141–88. <https://doi.org/10.1613/jair.2934>.
- Wiley, J.F. 2016. *R Deep Learning Essentials*. Birmingham, UK: Packt Publishing Ltd.

Zou, Hui, and Trevor Hastie. 2005. "Regularization and Variable Selection via the Elastic Net." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67 (2): 301–20. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>.