

A Comparison of LSA and LDA for the Analysis of Railroad Accident Text

Trefor Williams^{a*}, John Betak^{b, c}

^aRutgers University, New Brunswick, USA, 08854

^bCollaborative Solutions LLC, Albuquerque, USA, 87122

^cThe University of Texas at Austin

Abstract

Latent Semantic Analysis (LSA) and Latent Dirichlet Allocation (LDA) were used to identify themes in a database of text about railroad equipment accidents maintained by the Federal Railroad Administration in the United States. These text mining techniques use different mechanisms to identify topics. LDA and LSA identified switching accidents, hump yard accidents and grade crossing accidents as major accident type topics. LSA identified accidents with track maintenance equipment as a topic. Both text mining models identified accidents with tractor-trailer highway trucks as a particular problem at grade crossings. It was found that the use of the two techniques was complementary, with more accident topics identified than with the use of a single method.

Keywords: railroad, text mining, accidents

1. Introduction

There are various causes of railroad accidents. In the United States, railroads are required to report accidents with damage costing greater than US \$9,200 in 2010 to \$10,500 in 2015 to the Federal Railroad Administration. The accident form includes a field for a textual description of the accident. Text mining is defined in the context of discovering previously unknown information that is implicit in the text but not immediately obvious¹. There are various methods of text mining with distinct methods of identifying underlying topics in the text. This paper compares the results of applying Latent Dirichlet Allocation (LDA), a type of probabilistic topic modeling, and Latent Semantic Analysis (LSA), a natural language processing technique, to the text field in a database of Federal Railroad Administration Equipment Accident Reports.

There are few existing examples of the application of text mining to the railroad industry. Williams and Betak^{1 2 3} have used LDA analysis, to study grade crossing accidents, equipment accidents and accident investigation reports. Brown⁴ has used LDA models to mine text from railroad accident reports. The output from these LDA models are used to enhance data analytic models that predict the cost of extreme accidents. Previous railroad studies have only used LDA. The goal of this paper is to compare and contrast the output of LSA models with LDA to determine if use of the two models leads to additional insights when applied to railroad data.

The data used for this study are from the Federal Railroad Administration's railroad equipment accident database available on line at <http://safetydata.fra.dot.gov/officeofSafety/publicsite/Query/AccidentByStateRailroad.aspx>. The database includes text fields

that describe each accident. The length of the text field varied from a few words to paragraphs with several sentences. These text data were mined to reveal additional knowledge about railroad accidents. Data for the railroad equipment accidents were collected from January 2010-February 2015. There were

Table 1. Examples of Accident Text

Accident 1	ENGINEER STARTED TO PULL TRAIN AHEAD WHEN HE WAS RADIOED TO STOP, CARS ON GROUND. INVESTIGATION FOUND THREE CARS ON GROUND. ALSO FOUND PREVIOUS CREW POSSIBLY LINED SWITCH WRONG WHICH CAUSED THE DERAILMENT.
Accident 2	Y-SDG2321-13 DERAILED 3 ARTICULATED RAILCARS WHILE PULLING OUT OF YARD TRACK 9802 DUE TO TRACK WIDE GAGE. NO HAZARDOUS MATERIALS WERE RELEASED.

* Corresponding author. Tel.: +18324452880

Fax: +17324452232; E-mail: tpw@soe.rutgers.edu

© 2019 International Association for Sharing Knowledge and Sustainability.

DOI: 10.5383/JUSPN.11.01.002

12,447 accidents reported during this period. Table. 1 shows some example accident descriptions.

2. Text Mining Algorithms

The details of the LSA and LDA algorithms are presented in this section. This section illustrates how the two algorithms use different mechanisms to automatically generate the topics in the text corpus. A topic is a grouping of related words.

2.2. Latent Dirichlet Allocation

Topic modeling algorithms are statistical methods that analyze the words of unstructured original texts to automatically discover the themes that run through them. Topic models automatically organize a text collection into its major themes. A frequently used topic-modeling algorithm is Latent Dirichlet Allocation (LDA). Details of the LDA Algorithm are given by

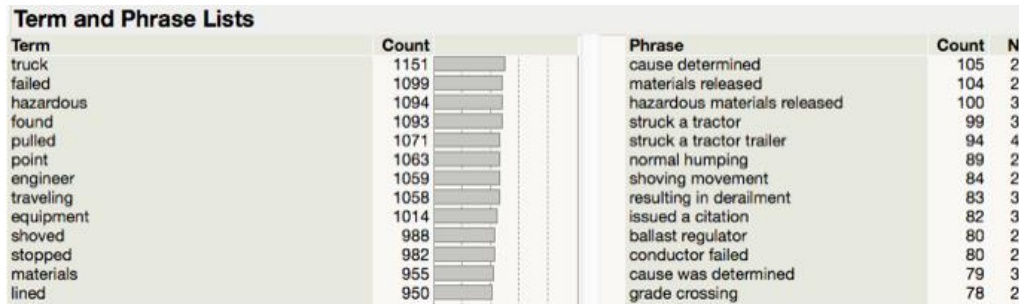


Fig 1. Frequent Words and Phrases

2.1. Latent Semantic Analysis

Text can be characterized by the semantic content it carries. Over the past two decades computational models have been developed to create semantic representations for words encountered in text. One such model is Latent Semantic Analysis (LSA)^{5,6}. LSA is a computational model that works on the notion that words with similar meanings tend to appear in similar contexts. It creates semantic representations for words by analyzing the pattern with which words occur together in documents across thousands of text samples provided to it in a training corpus. Then, from an analysis of the words that do and do not co-occur in the corpus, the model estimates what words should occur in similar documents (i.e., contexts) and are, therefore, close to each other in the semantic space⁷.

Blei⁸. LDA is a generative probabilistic model for collections of discrete data such as text corpora.

The underlying assumption of LDA is that a text document will consist of multiple themes. LDA is a three-level hierarchical Bayesian model where each item of a collection of text is modeled as a finite mixture over an underlying set of topics. Each topic is, in turn, modeled as an infinite mixture over an underlying set of topic probabilities. For text modeling, the topic probabilities provide an explicit representation of a document⁹. Additionally, a topic model generates automatic summaries of topics in terms of a discrete probability distribution over words for each topic, and further infers per-document discrete distributions over topics. In other words, the LDA algorithm automatically identifies words that occur in the accident reports and forms them into ranked topics. In this case, we are using the LDA algorithm's ability to find themes in the FRA accident reports.

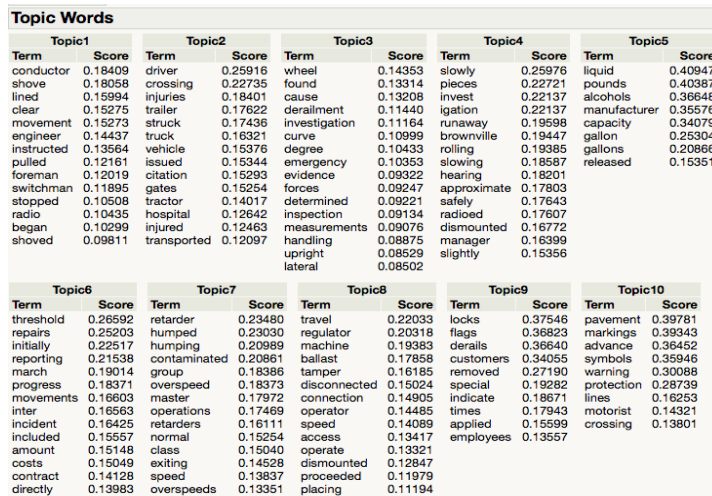


Fig 2. LSA Topics

3. Basic Text Mining

The initial phase of analyzing the railroad accident text was to remove stop words and to tokenize the text into individual words and phrases up to four words in length. Using the JMP text mining software, an initial analysis of the text provides a count of the most frequently occurring words and phrases. The JMP software also provided the facility to add stop words. In processing the text, frequently occurring words like “the” and “and” are automatically removed from the word list. The software uses a standard list of these stop words. It was found that adding the most frequently occurring words in the FRA accident text to the stop word list improved the clarity of the topics. Therefore, words like rail and track were deleted from the analysis.

Figure 1 shows the most frequently occurring words, and phrases after stop words were removed. There are several different frequently occurring phrases that illustrate the nature of many accidents. For example, one of the most frequently occurring phrases is “struck a tractor trailer.” This indicates that accidents at highway grade crossing between a train and a tractor-trailer highway truck. Another interesting term that occurs frequently is “ballast regulator.” This is a type of railroad maintenance equipment that is used to shape and distribute the stone ballast that acts as a foundation for railroad ties. Recently there have been several serious accidents involving maintenance-of-way equipment being struck by trains. The phrase “hazardous materials released” occurs frequently and indicates that in many accidents there are releases of liquid or gaseous chemicals.

Table 2. Some LDA Topics

Topic	Top Words
1	bowl utlx hump humped car gatx humping tilx tank operations class retarder fuel cut gallons group foul stalled speed system
2	switch lined point movement ran crossover move run line failed previously reverse switches improperly zone split route running points pulled
3	track cars shoving shoved cut shove pulled job made movement move foreman joint making make coupling clear double couple switchman
4	signal stop damaged fire operator causing pantograph stopped reported bridge control hit failed wire tra machine time ballast contact ck
5	train found emergency mph mp inspection brake speed investigation revealed air traveling curve brakes excessive slack handling grade upright experienced
8	struck crossing truck trailer injuries unit driver vehicle impact tractor road stop lead injured semi fouling gates northbound rig front
9	derailment cars derailed loaded wheels caused empty determined investigation coal curve load inside high grain flat resulted csx hopper center
10	derailed cars head loads ns pulling empties derailling tons units engines st westward shoving locomotives dttx wc eastward sou ft
14	cars released due track hazardous materials yard shoving derailed failure impacted articulated trk rco control contained pulling kicking mt irregular

4. LSA and LDA Model Outputs

The topics generated by the LSA and LDA models are shown and their meaning in a railroad accident context are discussed in this section.

4.1. Latent Semantic Analysis Topics

LSA can be used to automatically generate topics from the words contained in the corpus of text of the railroad equipment accidents. Each topic is a listing of words associated with a




Fig. 3. Topic Frequency from LDA Analysis

particular accident theme. The number of topics to generate is selected by the user. Through experimentation it was found that 10 topics yielded the most useable results. Figure 2 shows the topics generated from the analysis. Figure 2 shows a table of terms in each topic that have the largest scores in absolute value. Each topic is sorted in descending order by the absolute value of the score.

The LSA analysis yielded several interesting topics. They include:

- Grade crossing accidents. Two topics clearly addressed grade crossing accidents. Topic 2 shows there is a grouping of truck and tractor-trailer accidents at highway grade crossings. Topic 10 mentions a text grouping that includes accidents related to highway pavement markings and signs at grade crossings.
- Accidents related to maintenance equipment particularly including tampers and ballast regulators (Topic 8).
- Accidents relating to switching when cars are shoved (Topic 1).
- Accidents in hump yards. Hump yards are yards where freight cars pushed up a small hill and then roll down to be automatically switched to the correct track (Topic 7). Many major classification yards in the United States are hump yards.
- Accidents related to liquid spills (Topic 5).
- Accidents caused by wheel problems (Topic 3).

4.2. Latent Dirichlet Allocation

The LDA software requires the number of topics to be input by the user. Best results were found using 20 topics. Table 2 shows the most interesting topics generated from the LDA modeling. Topic 1 is a topic related to hump yard accidents. Topics 2, and 7 include switching accidents. Topic 3 indicates that there is a grouping of accidents involving cars being shoved. Topic 4 appears to include accidents involving signals. Topic 5 indicates that accidents involving brakes are a major railroad equipment accident theme. Topic 8 is very interesting because it suggests that highway-rail grade crossing accidents involving trucks with

trailers is a common accident type. Fig. 3 shows the number of accidents each LDA topic is associated with. Some accident topics occur much more frequently than other topics. It can be seen from Fig. 3 that many accidents are related to grade crossings (Topic 8). Topic 14 includes the words “hazardous” and “materials” and suggests that accidents involving hazardous materials occur in yards when cars are shoved or kicked. This

Table 3. Comparison of Identified Topics

Topic Meaning	LSA	LDA
Shoving/Switching	X	X
Grade Crossing Accidents	X	X
Wheels	X	X
Liquids Released	X	
Hazardous Materials		X
Track Maintenance Equipment	X	
Braking Accidents		X
Hump Yards	X	X
Derailement/Derailed		X

Table 4. LSA Topic Weights

Accident	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
1	2.45	-2.06	0.60	-5.34	0.74	-0.82	0.17	-0.08	0.85	0.77
2	-2.15	-1.62	-0.77	0.52	-0.44	-1.25	0.38	0.19	-0.26	-0.48

topic also occurs frequently. Words like “derail” and “derailed” are major words in Topics 9, 10, 19 and 20. Topic 20 is the most frequently occurring topic. Its top 5 words are rail, derailed, due, broken and railcars.

5. Comparison and Analysis of the Topic Modeling Techniques

Both LDA and LSA yield useful information about the nature of the railroad equipment accidents. There is significant overlap in the topics generated by the two methods. Table 3 shows the major accident themes identified by each method. Both techniques clearly show accidents related to hump yards, grade crossings, wheels and switching/shoving. The LSA technique identified accidents related to track maintenance equipment particularly ballast regulators and tampers, while the LDA did not clearly identify this topic. The LDA analysis found braking accidents as a topic, and the LDA analysis also found topics where the highest ranked word was derailment or derailed.

Both LSA and LDA generate topic scores for each accident text. Table 4 show the LSA weightings for the two accidents shown in Table 1. The highest positive weighting number indicates the

topic the accident is assigned to. Examination of Table 4 shows that the highest positive weighting for accident one is given to topic one. This indicates that topic one is the best match for this particular accident. An examination of the text in accident one indicates that it was a switching accident and that the accident has been classified in the correct topic. For the second accident shown in Table 1 the highest topic weighting was for topic 5 although the weighting for this particular accident is much lower indicating it does not fit in the topic as well as the first accident.

6. Future Research: Automatic Classification of Accident Text

The generation of topic scores by both LDA and LSA provide the potential to use the topic scores as a method of categorizing new accidents. The topic scores, coupled with the accidents assigned topic can be used as training data for a classifier, like a neural network, capable of identifying the topic when presented with new data. Additionally, the LDA and LSA output could be linked to data visualizations of numeric data to give a better overall indication of the accident types that are occurring

7. Conclusions

The analysis shows that both techniques find the most frequently occurring accident types. However, the text mining techniques generated several topics that are not as widely known in the railroad industry including the identification of accidents involving ballast maintenance equipment, and the prominence of tractor-trailer highway trucks in grade crossing accidents. This illustrates how the text mining tools can be used to identify

problems that require further investigation. It also illustrates that text mining yields information not observable in the numeric data used in most railroad accident statistical analyses. This further suggests that a very rich field of analysis lies in using these tools on additional railroad databases, such as track inspection reports, railroad incident reports filed for each employee involved in an incident and so on.

The use of the two text mining techniques complements each other. LSA and LDA are in agreement for many of the major accident topics, yet they each generated some topics that the other method didn’t identify. This indicates that using more than one text mining technique that use different mechanisms to identify topics can result in a more meaningful analysis and better identification of accident causes from the text.

References

[1] William T, Betak J, Nelson, C. Applying Topic Modeling to Railroad Grade Crossing Accident Reports. In: Joint Rail Conference 2015, Mar 23-26, San Jose, CA; New York: American Society of Mechanical Engineers, 2015. <https://doi.org/10.1115/JRC2015-5633>

[2] Williams T, Betak, J, Findley, B. Text Mining Analysis of Railroad Accident Investigation Reports. In: Joint Rail

Conference 2016, April 12-15, Columbia, SC; New York:
American Society of Mechanical Engineers, 2016.
<https://doi.org/10.1115/JRC2016-5757>

- [3] Williams, T, Betak J. Identifying Themes in Railroad Equipment Accidents Using Text Mining and Text Visualization, 2016 In: International Conference on Transportation and Development, Jun 26-29, 2016; Washington: American Society of Civil Engineers, 2016.
<https://doi.org/10.1061/9780784479926.049>
- [4] Brown, D. Text Mining the Contributors to Railroad Accidents. *IEEE Transactions on Intelligent Transportation Systems* 2016;**17**(2):346-355.
<https://doi.org/10.1109/TITS.2015.2472580>
- [5] Landauer TK, Dumais ST. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychol Rev* 1997;**104**(2):211-240. <https://doi.org/10.1037/0033-295X.104.2.211>
- [6] Landauer TK, Foltz PW, Laham D. An introduction to latent semantic analysis. *Discourse Processes* 1998 01/01;**25**(2-3):259-284.
<https://doi.org/10.1080/01638539809545028>
- [7] Kwantes PJ, Derbentseva N, Lam Q, Vartanian O, Marmurek HHC. Assessing the Big Five personality traits with latent semantic analysis. *Personality and Individual Differences* 2016 November 2016;**102**(Supplement C):229-233. <https://doi.org/10.1016/j.paid.2016.07.010>
- [8] Blei DM, Ng AY, Jordan MI, Lafferty J. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 2003 05/15;3(4):993-1022.
- [9] Blei DM. Probabilistic Topic Models. *Communications of the ACM* 2012;**55**(4):77-7.
<https://doi.org/10.1145/2133806.2133826>