

IMPROVING AIRCRAFT MAINTENANCE, REPAIR AND OVERHAUL: A NOVEL TEXT MINING APPROACH

Jun Yu

Knowledge Transfer Centre
University of Reading
Reading, United Kingdom
yujunuk@gmail.com

Stephen Gulliver

Informatics Research Centre
University of Reading
Reading, United Kingdom
s.r.gulliver@henley.reading.ac.uk

Abstract — Aircraft Maintenance, Repair and Overhaul (MRO) feedback commonly includes the engineer’s complex text-based inspection report. Capturing and normalizing the content of these textual descriptions is vital to cost and quality benchmarking, and provides information to facilitate continuous improvement of MRO process and analytics. As data analysis and mining tools requires highly normalized data, raw textual data are inadequate. This paper offers a textual-mining solution to efficiently analyse bulk textual feedback data.

Despite replacement of the same parts and/or sub-parts, the actual service cost for the same repair is often distinctly different from similar previously jobs. Regular Expression algorithm was incorporated with an aircraft MRO glossary dictionary in order to help providing additional information concerning the reason for cost variation. Professional terms and conventions were included within the dictionary to avoid ambiguity and improve the outcome of the result. Testing results show that most descriptive inspection reports can be appropriately interpreted, allowing extraction of highly normalized data. This additional normalized data strongly supports data analysis and data mining, whilst also increasing the accuracy of quotation costing of future work. This solution has been effectively used by a large aircraft MRO agency with positive results.

Keywords- Aircraft, MRO, Text Mining, Quotation

I. INTRODUCTION

The global economic crisis, compounded by the rising cost of fuel, has led to most airlines strategically targeting cost as a top priority. When considering Aircraft Maintenance, Repair and Overhaul (MRO), airlines demand the very best cost efficiency, whilst ensuing no compromise in the areas of quality and safety. A common practice when costing a new service job is to compare the current job with multiple examples of similar previous work. To achieve this, however, comparative consistency has to be found between the current job and existing cases in a historical MRO database. As each fault is potential unique, it is impossible however to capture detailed feedback information in formalized forms. Use of a textual description supports the detailed capture of unique fault information, yet feedback quantity, style and layout of reports varies considerably between service companies and/or engineers. Formalizing

structure of feedback, although ideal for analytics, limits the ability of engineers to document detailed comments or suggestions for process improvement.

Aircraft MRO typically involves many complicated details. For example: A Boeing 747-200 vertical stabilizer repair might relate to one of thousands of small parts; pricing of regular maintenance on an engine can vary considerably according to its warranty status; even a passenger seat belt repair can vary hugely based on every individual case. In order to allow consistent comparison, Business Intelligence (BI) is vital to analyse, sort and capture meaningful and consistent semantics from repair feedback. Such individual differences are easily ignored when jobs are summarized and/or data normalized. Information such as defect part-number, repair type, supplier contact number, final quote and completed date are commonly normalized in structure, so are easy to formalize and store for future reference. The more descriptive information in the engineer’s inspection report, however, is often ignored or left unprocessed.

BI systems offer great potential to help MRO companies. Increased feedback supports MRO companies to better understanding the reasons for deviation in service cost, and therefore supports improvement in the accuracy of future cost estimation. Moreover, descriptive inspection reports regularly contain detailed comments about the job and/or suggestions for process improvement, which offer the MRO suppliers process feedback that currently goes to waste. Current tools for analysing descriptive data (i.e. written reports or text) are not directly useful to MRO descriptive feedback, however the following sections described the problems and how adaption of current tools facilitated measurable improvements in the accuracy of job costing.

II. PROBLEMS

A conventional MRO process involves: a customer (airline) delegating an MRO job to an MRO agency. The agency references an MRO database of completed jobs and quotes both i) a ‘cheapest’ price, and ii) the ‘best’ supplier to fulfil that MRO job. Once the quote is granted by the customer, the part is sent, checked, repaired, and returned, the customer is billed and then payment is cleared.

A key issue during the above process is the creation of the quotation. How can we ensure that the ‘cheapest’ price is quoted, as all jobs are uniquely different? Moreover, how

can we promise the ‘best’ supplier is recommended? These issues depend on the historical database containing enough relevant information to answer each question. Normally quotes are based solely on normalized data, which we have mentioned is unable to provide information about certain factors. For example, an engine repair cost varies from a few hundred USD to more than several hundred thousand USD depending on the sub-parts involved or warranty status. An average price, without more information, will clearly be very inaccurate. Capture of more precise historical data is clearly important to ensure the best possible quotation for customers. This paper offers a text mining approach to improve both the data quality and quantity.

III. TEXT MINING

Textual data mining relates to the process of extracting high-quality information from large quantities of textual content. By defining structural patterns within the text, normalised forms of information can be derived, which can be used to add value to existing formalised sources of information. Common textual data mining techniques include: textual categorization [1], textual clustering [2], subject / object extraction and sentiment analysis [3]. Text data mining is used commonly in online media form processing, especially in the domain of CRM (Customer Relationship Management) [4]. Sentiment textual analysis is used to capture user satisfaction, often taken as feedback after media viewing or interface testing [5]. To facilitate general usage, text mining is now being incorporated, at some level of complexity, within the majority of Business Intelligence (BI) software solutions, however no one has considered the best form of textual data mining when analysing MRO feedback text.

IV. PROPOSED METHODS

A. Incorporating Regular Expression

Regular expressions are patterns of characters that match or fail to match, sequences of characters in text, allowing users to identify the presence of desired pieces of text [6]. In short, a regular expression is a specific kind of textual pattern. The Microsoft (MS) .Net Framework library contains classes that implement regular expression (Regex), which provides a powerful, flexible and efficient method for processing text [7]. Regular expression’s extensive pattern-matching notation allows users to quickly parse large amounts of text to find specific character patterns, to extract, edit, replace or delete textual substrings [8] thus splitting descriptive text into analysable individual words. This process, however, needs to be done very strategically. A four stage process was implemented in our work.

Firstly, from the complete engineering report, we narrowed down the scope of feedback text since extraneous information such as the engineer’s signature, company name, reply/forward contents, only done the text longer and the outcome more confused.




Figure 1. An example screen dump demonstrating the removal of extraneous textual content.

Secondly, scoped descriptive text was split at the sentence level in order to remove irrelevant information. A simple English sentence starts with a capital letter or a punctuation mark, and ends with a full-stop. In reality a few exceptions are tolerated, e.g. starting with a lowercase letter or ending with a question/exclamation mark etc., yet by finding the end of sentences, and by adding a newline character, we were able to accurately split text paragraphs into distinct sentences (see Fig.1 – relevant information). Finding sentence features relies entirely on development of appropriate Regular Expression syntax. A Regular Expression Tester was used to identify the syntax of sentence structures before implementing it into programmable code (see Fig.2).

Both matching a capital letter and finding a special symbol, the basic syntax of Regular Expression has to be adapted. MS .Net provides a Regex class that is a powerful search engine for string searches. Most MS .Net programming languages, such as C# or VB, they can easily implement a string search by declaring a use of System.Text.RegularExpressions at the program beginning.

Once the scoped text was split into sentences, the third stage is to split these sentences into individual words with all punctuations, irrelevant marks (e.g. brackets) and linking words being removed. Information could be lost at this stage however key information should hopefully remain.

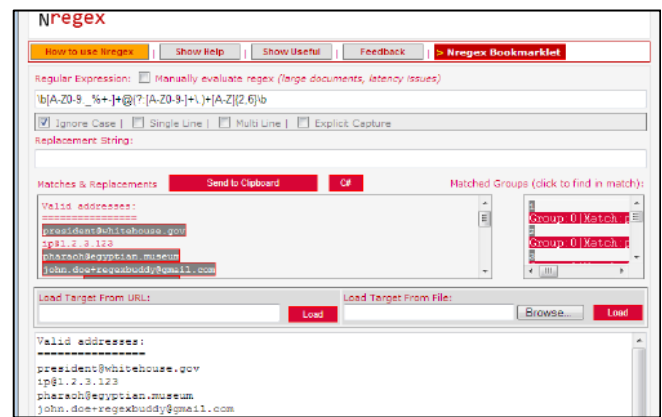


Figure 2. A Regular Expression Tester (<http://www.nregex.com/nregex>)

Finally, with reference to the aircraft MRO glossary dictionary, we classified, prioritized and counted the frequency of each word. The following section describes the construction and use of the MRO dictionary.

B. The Aircraft MRO Glossary Dictionary

Aircraft maintenance, repair, and overhaul (MRO) is a critical process with strict requirements to safety, security and quality, therefore its execution is highly standardized by international bodies; including the Air Transport Association [9].

The aircraft MRO glossary dictionary is an alphabetical list of industrial terms extracted from engineering reports, which is based on half a million MRO jobs. The dictionary, however, is not only a list of terms, but acts as a network of interlinked dimensions allowing reference using Terms, Part Number, Repair Type, Suppliers, Customers etc. For instance, a specific term may appear a large number of times on the report that relates to a specific part or a specific MRO job type.

The Aircraft MRO glossary dictionary offers a reference base for analysis relating to previously processed job reports. According to the frequency of the word appearing on the report, and by referring to the aircraft MRO glossary dictionary additional normalized information can be inferred, e.g. part number, repair type, service price etc., yet this normalized data may have already been provided in the engineer's feedback. Words describing the reasons for failure, symptoms of a broken part, and the ultimate solutions are of considerable value during text mining. Accordingly, the following section discusses identification of the key influencers and prioritization, which is used to determine importance.

C. Key Influencers Analysis

Split words are treated differently, depending on use of contextual use of the word and the sentence structure, i.e. where the word belongs. To determine this, a three stage process was used.

English sentence structures are categorized for proper extraction of the key influencers in a sentence. Krohn [11] noted that English sentence structures include the following forms: Subject-Verb-Object, Subject-Verb, Subject-Verb-Indirect Object-Object, Subject-Verb-object-indirect Object etc. The first stage is to define the structure of the given sentence. Based on the defined sentence structure, the second stage is to analyse and determine the subjects, adjectives, objects and verbs in the given sentence. A key influencer (word) is then classified as its attribute (see Fig.3).

The third stage of the process involves the analysis of sentence content, i.e. word attributes and sentence structure. A key influencer diagram is used to develop an easy impact reference (see Fig.3). During stage 3, the word frequency in each sentence was also counted, as an important parameter to sentence meaning. A real example of key influencers is shown in fig.3. This example relates to an 'Oil Cooler and Relief Bypass Valve Assembly', and clearly shows influencers, analysed from past engineer reports, which relate to problems, processes or actions involving this part.

Influencers	Attribute	Part	Relative Impact
Flange	noun	341D8028-501	
Nicks	verb	341D8028-501	
Sealing	noun	341D8028-501	
SPM	noun	341D8028-501	
Scrath	noun	341D8028-501	
Replacement	noun	341D8028-501	
Clean	verb	341D8028-501	
Blend	noun	341D8028-501	

Figure 3. Key influencers relating to 'Part Number' 341D8028-501; an Oil Cooler and Relief Bypass Valve Assembly.

The three influencers with the most impact were defined as: seal, clean, replacement. Despite appropriate and relevant engineering experience, we are able to infer from these influencers that the fault may be either a blockage or a broken seal. If cleaning the part fails, then a replacement part may be needed. Estimation of 'activity based costing' is also therefore possible with subsequent analysis.

V. TESTING

A. Training and Testing Data

An MRO history database was created as training data from 470,449 existing MRO jobs. Each job includes both normalised and unnormalised data. Normalised data are processed, properly classified and stored within a relational database (MS SQL Server) which relates to information such as job reference number, part number, repair type, service cost, customer, supplier and completion date etc. Unnormalised data, i.e. the engineer inspecting reports, is contained within the descriptive paragraphs.

Testing data was taken from the MRO live database, which at the time in question included 97,857 MRO jobs. Fig.4 illustrates the relative proportions of training and testing data sets assuming continued growth in the testing data set of about 50-80 new MRO jobs per day.

B. Testing Procedure

The reference number of any suspected broken part was firstly fed into the test system. First tier query results retrieve normalized information including the lowest/highest service cost, repair types involved, suppliers who have completed




Figure 4. The proportion of training data and testing data

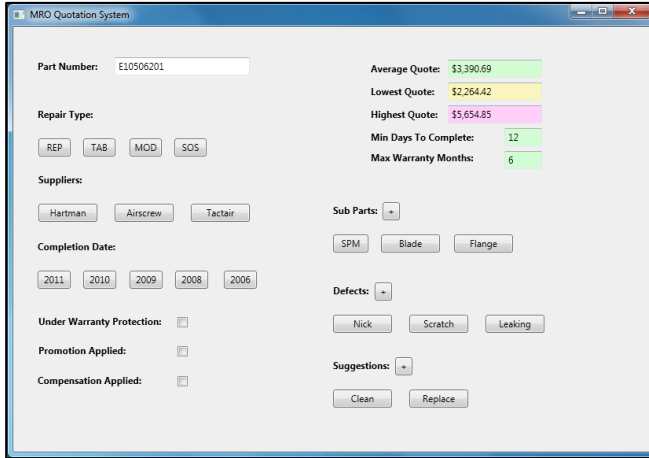


Figure 5. A screen shot for the test environment.

similar jobs before, etc.

Information extracted from the engineer's inspection report (i.e. the key influencers) was classified as being one of three categories: sub-parts, defects or suggestions for easy interface reference (Fig.5). As the user selects more influencers, the displayed information changes.

To validate the accuracy of the system, query results, i.e. an average quote price was calculated by the testing system, and compared with the actual service cost for each job. It has been agreed that a deviation of 20% from the quotation price is acceptable according to the industrial convention. However, as the quotation system improves in the future, the allowed range could be reduced to 10% or even less.

If extreme deviation is identified then reports will be manually checked to ensure that the reasons for differences in pricing are revealed, such as the introduction of a new sub-part, or high inflation that will clearly make current prices higher than that of historical quotes. In this case, new sub-parts will be added to the MRO glossary dictionary and an inflation multiplier will be added to the historical cost during quotation calculations. In our experience, however, there are many cases in which there is no clear reason to support a rise in the service cost above that of the system estimated cost. Replacement of the same parts and/or sub-parts, implementation of the same repair type, in the same duration, often lead to a cost that is distinctly different from previously completed jobs. In the following section, we look at the success and failure resulting from our work, and provide detailed figures to validate whether the mining of engineering inspection reports supports the accuracy of service quotation.

VI. RESULTS

A. Successes

Overall comparison figures showed that in 97,857 test cases, the system could successfully estimate job costs (within a 20% deviation range) 76.3% of the time. With a 25% deviation range, the figure rose to 84.9%. Jobs were successfully estimated 88.7% of the time for if a cost

deviation of 30% was deemed acceptable. Only 65.5% of cases fall within the 10% cost deviation range (see Fig.6).

Our proposed approach, i.e. incorporating mined data from engineer service reports, increased estimation accuracy by approximately 18% over the existing system where results were solely based on normalised data (see Fig.6).

A high proportion of jobs, however, were still found to have a service cost greater than 10% above the estimated system estimated quote. Apart from a small portion of cases, which related to inflation, most additional costs related to over spending caused by incorrect selection of a supplier. We believe that this issue could be better managed via provision of additional supporting information and use of appropriate Business Processes. Selecting the right supplier is the key to achieving cost efficiency and quality benchmarking, so we aimed to extend our current work on cost estimation to help support supplier selection.

Another positive outcome of implementing this system is the on-going and continuous expansion of the aircraft MRO historical database. The accuracy of quotes depends on the accuracy of the reference dictionary. It is important that we have both i) enough (i.e. quantity) or data, relating to past jobs; but also ii) that the quality of that data are worth referencing. This system adds considerable extra value by identifying key influences from previously unused engineering reports. This information increases the quantity of data, but also provides key influencer analysis, which should ensure data quality. Furthermore, by presenting the key words regarding each specific job (part), customers can proactively specify the job nature, making the quote more accurate. Using our system action based costing is also possible, which allows customers to see both i) the range of possible costs, but also ii) the justification for why variation might occur.

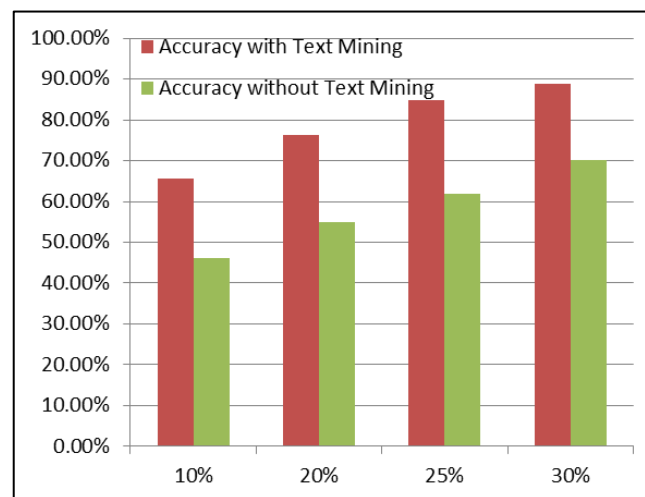


Figure 6. Accuracy based on deviation ranges (with/without text mining)

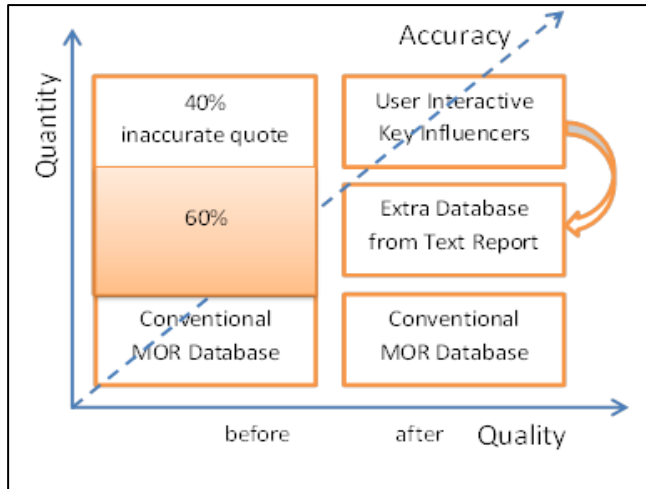


Figure 7. A comparison between the approach adopted before and after

In summary, implementation of the suggested process provided significantly successful improvements in two areas. Firstly, text mining provided more information from the previous unused engineer's report to support company analytics and cost estimation. Secondly, key influencers were defined as the text mining outcome, which that enables user proactively to narrow down the defect scope, supporting more precise data query. Fig.7 visually illustrates the improvement gained by solution adoption.

B. Failures

Despite aircraft MRO processes being highly standardised, i.e. supporting unified operations and terms, testing results showed that most analysis failures were caused because of engineers using highly complex sentence structures, typos or use of irregular English. Further improvement, by adding more comprehensive syntax processing, is needed to consider the flow and individual writing style of individuals.

VII. CONCLUSIONS

The approach introduced in this paper can significantly improve both aircraft MRO job data quantity and quality. This allows undiscovered data, which currently existed unused in engineer's reports, to become available to query. Additional data allows engineers to focus on commonly identified fault areas, hopefully leading to time and cost savings; it supports appropriate selection of supplier, both for parts and services, facilitates activity based costing, which allows similar activities to be quoted independent of part number, and allows costing transparency for customers, letting the customer to see the possible range of cost and the change of this cost will be faced. Therefore a system created based on this approach, which provides a precise quotation mechanism, strongly help the aero industry archive a best cost and quality benchmark.

REFERENCES

- [1] S. Fabrizio. Machine learning in automated text categorization. *ACM Computing Surveys*, 34(1):1–47, 2002
- [2] Andrews, N.O.F and A. Edward. Recent Developments in Document Clustering, October 16, 2007
- [3] T. Peter. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. *Proceedings of the Association for Computational Linguistics*. 2002, pp. 417–424.
- [4] K. Coussemant, and P. Van. "Integrating the Voice of Customers Through Call Center Emails into a Decision Support System for Churn Prediction." *Information and Management*. 45.3 (2008): 164–174. Print.
- [5] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? Sentiment Classification using Machine Learning Techniques". *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2002, pp. 79–86.
- [6] Watt. *Beginning Regular Expressions*. Indianapolis, Ind: Wiley, 2005, pp.2.
- [7] J. Goyvaerts and L. Steven. *Regular Expressions Cookbook*. Beijing: Oreilly, 2009. pp.1.
- [8] T. Nash. *Accelerated C# 2010*. New York, NY: Apress, 2010, pp.232.
- [9] Microsoft MSDN. *NET Framework Regular Expressions*, 2011.
- [10] M. Lutz, S. Alexander and G. Lars. *International Workshop on Topic Maps Research and Applications. Leveraging the Semantics of Topic Maps*. Berlin: Springer, 2007. pp.56.
- [11] R. Krohn. *English Sentence Structure*. Ann Arbor: University of Michigan Press, 1971.