

New advances in aircraft MRO services: data mining enhancement

Conference or Workshop Item

Accepted Version

Yu, J., Gulliver, S., Tang, Y. and Ke, L. (2011) New advances in aircraft MRO services: data mining enhancement. In: Fourth International Workshop on Advanced Computational Intelligence (IWACI2011), 19-21 October 2011, Wuhan, China. Available at <http://centaur.reading.ac.uk/24806/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <https://doi.org/10.1109/IWACI.2011.6160002>

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

New Advances in Aircraft MRO Services: Data Mining Enhancement

J. Yu, S. Gulliver, Y. Tang and L. Ke

Abstract — Aircraft Maintenance, Repair and Overhaul (MRO) agencies rely largely on row-data based quotation systems to select the best suppliers for the customers (airlines). The data quantity and quality becomes a key issue to determining the success of an MRO job, since we need to ensure we achieve cost and quality benchmarks. This paper introduces a data mining approach to create an MRO quotation system that enhances the data quantity and data quality, and enables significantly more precise MRO job quotations.

Regular Expression was utilized to analyse descriptive textual feedback (i.e. engineer's reports) in order to extract more referable highly normalised data for job quotation. A text mining based key influencer analysis function enables the user to proactively select sub-parts, defects and possible solutions to make queries more accurate. Implementation results show that system data would improve cost quotation in 40% of MRO jobs, would reduce service cost without causing a drop in service quality.

I. INTRODUCTION

Maintenance, Repair and Overhaul in the airline industry involves the fixing of out of order or broken mechanical or electrical devices. MRO activity also includes routine scheduled maintenance to minimize the risk of a future fault. A standard MRO process involves: the customer (i.e. an airline) delegating an MRO job to an MRO agency. The agency references a MRO database of completed jobs and quotes two critical pieces of information. Firstly the MRO agency provides information about the 'cheapest' price. Secondly the MRO agency provides information

about the 'best' supplier to fulfil that MRO job. The customer then chooses and grants a quote. Once a quote is granted, the part is sent, checked, repaired, and returned to the MRO agency. The customer is billed for the work and payment is cleared.

The key issue during this MRO process is the creation of accurate and appropriate quotes. Incorrect quotation significantly impacts the chance of either a quote being granted (i.e. because the quote is too high), or may lead to a reduction in customer trust (i.e. the quote is too low and the bill is therefore perceived to be too high). How can MRO agencies ensure that the 'cheapest' price is quoted, as all jobs are uniquely different? Moreover, how can we promise that the 'best' supplier is recommended? These issues depend on the historical MRO database containing enough relevant information to answer these questions. Traditionally MRO quotes were based solely on normalized feedback data. Normalized feedback, however, is unable to provide information about the many more descriptive factors; that are important when comparing complex engineering jobs. An engine repair cost may vary from a few hundred USD to more than several hundred thousand USD depending on the sub-parts involved or warranty status. An average price therefore 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 solution to improve capture of both data quality and quantity within the maintenance and service domain.

II. A CONCEPTUAL FRAMEWORK

A. Text Mining

Textual data mining relates to the process of extracting high-quality information from large quantities of textual content [1]. The purpose of Text Mining is to automate the processing of unstructured (textual) information [2], so that meaningful semantic value can be obtained, i.e. by defining structural patterns within the text, normalised forms of information can be derived. This normalized semantically rich data can be used by common data mining (statistical and machine learning) algorithms to support the identification of clusters, associations and sequences in MRO activity.

In its most basic form, text mining allows words in the text to be catalogued in a dictionary and counted. This allows us to determine a matrix of how many

Manuscript received July 1, 2011. This work was supported in parts by the Informatics Research Centre and the Knowledge Transfer Centre from University of Reading (UK).

J. Yu is with University of Reading, United Kingdom. His current role is KTP (Knowledge Transfer Partnership) associate, interest includes Business Intelligence, Data Mining and .Net programming (e-mail: yujunuk@gmail.com).

S. Gulliver is a senior lecturer from the University of Reading. His areas of interest includes: multimedia perception, usability, pervasive and mobile devices, personalisation, attentive displays, virtual environments, Eye-tracking and data analysis. (e-mail: s.r.gulliver@henley.reading.ac.uk).

Y. Tang is the deputy director of the Informatics Research Centre, University of Reading, his expertise and research interests include bioinformatics; corporate strategy planning; project management; benchmarking; performance measurement and large scale databases. (e-mail: y.tang@reading.ac.uk)

L. Ke is a head of a professional IT development team. He has more than 15 years' experience working on the development of improving aircraft MRO service using IT technologies. (e-mail: lisheng_ke@yahoo.co.uk)

times each word has been used throughout the textual description. This basic process can be extended to exclude common stop words, for example, non-subject words that flow of readable sentences, yet have limited value to content analysis, e.g. "because" and "therefore". Functioned lists of stop words are publically available to help stop words without meaning entering the database. In addition stem checking and synonym comparison can be included to restrict multiple words being added to the dictionary with the same meaning, e.g. 'jump, jumped and jumping' or 'replace, swap and exchange'. Inclusion of multiple very similar meaning words dilutes the importance of meaning across words, often causing important meaning to be lost during data mining processes.

Common textual data mining techniques include: textual categorization [3], textual clustering [4], subject / object extraction and sentiment analysis [1]. Text data mining is used commonly in online media form processing, especially in the domain of CRM (Customer Relationship Management) [6]. Sentiment textual analysis is used to capture user satisfaction, often taken as feedback after media viewing or interface testing [7]. To the best of our knowledge no one has previously considered textual data mining when analysing MRO feedback text.

B. 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 [8]. In short, a regular expression is a specific kind of textual pattern. Regular expression patterns are defined using a formalized language, which is then interpreted by a regular expression processor. This processor parses text, and outputs substrings that meet the defined specification.

The formalized pattern-matching notation allows users to quickly parse large amounts of text to find specific character patterns. Textual substring are extracted, edited, replaced or deleted [10]; thus splitting descriptive text into analysable individual words. The MS .Net Framework library contains classes that implement regular expression (regex). Most MS .Net programming languages, such as C# or VB, can easily implement a string search by declaring use of System.Text.RegularExpressions [9]. This process, however, needs to be done strategically.

III. METHODS

A. Incorporating Regular Expression

In our work we suggest the use of four phases to successfully facilitate the best use of regular expressions.

Firstly, extraneous information such as the engineer's signature, company name, reply/forward contents, etc. was removed from the completed report text. This additional information is normally available

via normalized feedback and only acts to add confusion to the text mining process. By removing such information, we are trying to limit the inclusion of irrelevant information from entering the MRO dictionary; which would therefore subsequently require more regular data audit and cleaning processes to be undertaken. Figure 1 shows the removal of irrelevant information.

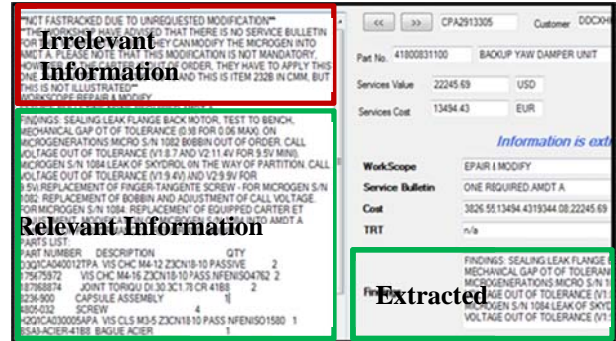


Fig. 1. An example demonstrating the removal of extraneous textual content.

Secondly, remaining descriptive text was split at the sentence level in order to remove additional irrelevant text. English sentences commonly start with a capital letter or a punctuation mark, and end with a full-stop. By identifying sentences, and by adding a newline character as required, we were able to accurately split text paragraphs into distinct sentences (see figure 1 – relevant information). Finding sentence features relies entirely on development of appropriate Regular Expression syntax, e.g. `[A-Z].*?[.!?]\s*((?=\s[A-Z])\s|$)`. A Regular Expression Tester was used to identify the syntax of sentence structures before implementing it into programmable code (e.g. <http://www.nregex.com/nregex>). To increase process success, inclusion of syntactic exceptions was subsequently considered, i.e. starting with a lowercase letter or ending with a question/exclamation mark etc.

Once the scoped text was split into sentences, the third stage was used to split these sentences into individual words; all punctuation and irrelevant marks (e.g. brackets) were removed. Stop words (e.g. and, but, so) are essential in English prose to ensure that writing is natural and clearly structured, however they themselves hold limited meaning. With reference to the MRO glossary dictionary, stem checking and synonym comparison was also undertaken to limit the sharing of meaning across multiple words. Although significant content will be lost at this stage, key meaning should remain.

Finally, words were classified and prioritized using the frequency of each word as the primary variable. 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 [11].

The aircraft MRO glossary dictionary is an alphabetical list of industrial terms extracted from historical 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 users to reference jobs using Terms, Part Number, Repair Type, Suppliers, Customers etc. The Aircraft MRO glossary dictionary offers a reference base for analysing current jobs against the detailed contain in 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, charge price etc.; however, inferred normalized data may already have been provided in formally structured normalized engineer feedback. Words describing the reasons for failure, symptoms of a broken part and the possible solutions are all of considerable value during the cost quotation process. 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.

The first stage is to define the structure pattern. Krohn [13] noted that English sentences have certain patterns, which relates to how the sentence was constructed, i.e. the order of the elements. In our work we used six sentence structures: i) Subject - Verb, e.g. "The oil cooler has failed"); ii) Verb - Subject, e.g. "There is a strange noise in the conditioner."; iii) Subject - Verb - Object, e.g. "Seals need modification"; iv) Subject - Verb - Indirect Complement, e.g. "The seat-belt bracket looks tired"; v) Subject - Verb - Indirect Object - Object, e.g. "The company gave the engineer a replacement oil-cooler.", Subject - Verb - object - indirect - Object, e.g. "I left the generator running over night". In our solution, English sentence structures were categorized and the key influencers were obtained.

Based on the defined sentence structure, the second stage is to analyse and determine the subjects, adjectives, objects and verbs in the given sentence. New subject influencers, are added to the dictionary. If inappropriate influencers or language is included in engineering reports, then these sentences can be capture before entering the MRO reference database and appropriate action can be taken. Moreover a record can

be taken for use in future supplier selection. All appropriate key influencers (i.e. words) are then classified with an attribute.

The third stage of the process involves the analysis of sentence content, i.e. word attributes, sentence structure and influencer order; a key influencer diagram is used to develop an easy impact reference (see Figure 2). Influencers with an inappropriate order can be highlighted at this point. The Subject-Verb-Object sentence "I smoke crack rocks" is inappropriate in an engineering report, even though each separate influencer describes something that may be realistically included within an engineering report (i.e. engine smoke, a cracked linkage, volcano rock and ash deposits). When attributes are combined in a certain order 'smoke crack' or 'crack rock', the sentence clearly takes on an inappropriate meaning, and action should be taken to remove the sentence from analysis. 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 figure 2. 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. The three influencers with the most impact were defined as: seal, clean, replacement. Despite having no appropriate and relevant engineering experience, we are able to infer from these influencers that the fault might be due to either a blockage or a broken seal. If the part is not dirty or blocked then a replacement part (most likely the 'seal') may need to be ordered. Estimation of 'activity based costing' is also therefore possible with subsequent analysis.

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	

Fig. 2. Key influencers relating to 'Part Number' 341D8028-501; an Oil Cooler and Relief Bypass Valve Assembly.

IV. TESTING

A. Training and Testing Data

A large industrial data set including 470,449 existing MRO jobs from the last ten years (2000-2010) was created as training data. Each job includes both normalised and unnormalised data. Normalised data is 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. The training data size is kept unchanged to provide referable quote benchmarks.

Testing data was taken from the MRO live database, which at the time in question included 97,857 MRO jobs. The testing result is worked out based on the test data set, its proportion will vary from 16% to nearly 30% within next 12 months as Figure 3 illustrates; since there are about 50-80 new MRO jobs per day added into testing database.

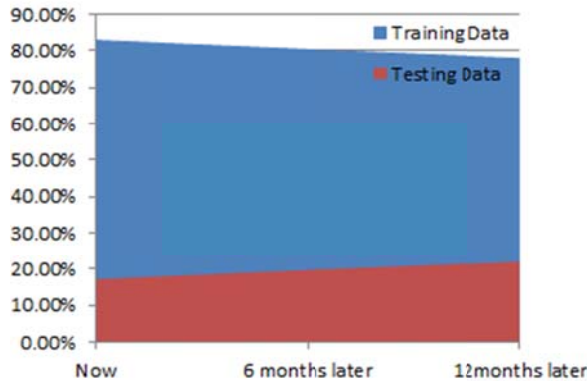


Fig. 3. The proportion of training data and testing data

There are two reasons to separate training and testing data: One is, that was a whole set of historical data for a certain period directly from industry (2000-2010). Another reason is the benchmark as the object to be compared, that needs to be consistent.

Certainly completed testing data will be added to training data as more training data can help improve the accuracy of the benchmarks. However the enlarged training data set will be treated separately for another round of test, i.e. when more fresh data becomes training data. A further comparison will be conducted before and after the training data enlarged.

B. Testing Procedure

The reference number of any suspected broken part was fed into the test system. Initial query results retrieved normalized information including the lowest/highest service cost, repair types involved, and suppliers who have previously completed similar jobs, etc.

Information extracted from the engineer's inspection report (i.e. the key influencers) were classified as being one of three categories: sub-parts, defects or suggestions for easy interface reference (see figure 4). As the user selects more influencers, the displayed information changes.

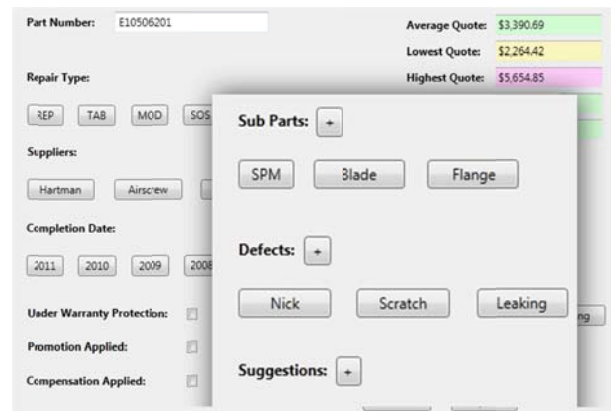


Fig. 4. A screen shot for the test environment.

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. However, as the quotation system improves in the future, the allowed range could be ideally reduced to 10% or less.

If extreme deviation is identified (i.e. over 20 %) then reports were manually checked to ensure that the reasons for differences in pricing are revealed; such as the introduction of a new sub-part, or because of high inflation that will make current prices higher than that of historical quotes. In this case, new sub-parts were added to the MRO glossary dictionary and an inflation multiplier was 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 led 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.

V. 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, if a cost deviation of 30% was reasoned as being acceptable. Only 65.5% of cases fell within the 10% cost deviation range (see Figure 5).

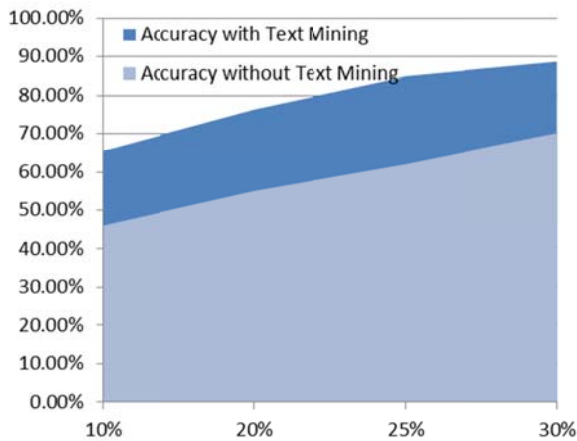


Fig. 5. Accuracy based on deviation ranges (with/without text mining)

Our proposed approach, i.e. independent of deviation range, 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 Figure 5).

A high proportion of jobs, however, were still found to have a service cost greater than 10% above the estimated system 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 aim 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 data (i.e. quantity) relating to past jobs; but also ii) that the quality of that data is worth referencing. This system adds considerable value by identifying key influencers 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 'Activity Based Costing' (ABC) is also possible. ABC is often used in Customer Relationship Management to allow customers to see both i) the range of possible costs, but also ii) the justification for why variation might occur.

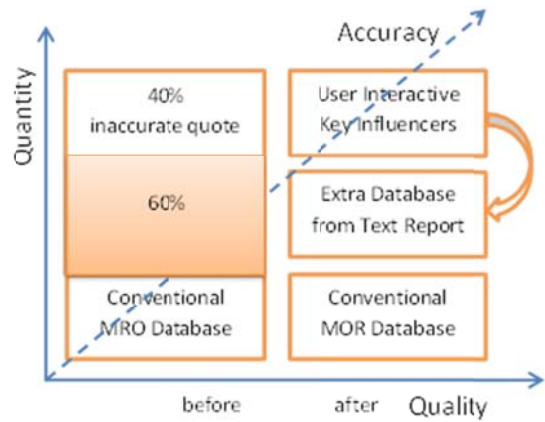


Fig. 6. A comparison of approaches - before and after.

In summary, implementation of the text mining process has significantly improved business function in two areas. Firstly text mining has provided more information from the previously unused engineer's report; to support company analytics and cost estimation. Secondly key influencers were defined as the text-mining outcome, which enables users to proactively narrow down the defect scope, supporting more precise data query. Figure 6 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.

VI. 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 analytic queries. 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; identifies inappropriate feedback; facilitates ABC, which allows similar activities to be quoted independent of part number; and allows costing transparency for customers, letting the customer see the possible range of cost and the chance that this cost will be faced. Incorporating this approach, i.e. facilitating more precise cost quotations, should help the aero industry, as well as many other maintenance and service domains, a better levels of quality at lower levels of cost.

ACKNOWLEDGMENT

This research was funded by a Knowledge Transfer Partnership project connected with an academic body (University of Reading) and an industry company (a leading aircraft MRO Service provider). A particular thank is sent to Mr Paul Taylor (Project adviser) for his support to make the success of this project.

REFERENCES

- [1] X. Zhou, Database Systems for Advanced Applications: 14th International Conference, Dasfaa 2009, Brisbane, Australia, April 21-23, 2009 ; Proceedings. Berlin: Springer, 2009.
- [2] J. D. Hwang and W. J. Wang, "Application of Icp-Aes to Analysis of Solutions." Applied Spectroscopy Reviews. 30.4 (1995): 231-350. Print.
- [3] S. Fabrizio, Machine learning in automated text categorization. ACM Computing Surveys, 34(1):1-47, 2002
- [4] N. O. F. Andrews, and A. Edward, Recent Developments in Document Clustering, October 16, 2007
- [5] 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.
- [6] K. Coussement, and P. den 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.
- [7] B. Pang, B., Lee, L., 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.
- [8] Watt. Beginning Regular Expressions. Indianapolis, Ind: Wiley, 2005, pp.2.
- [9] J. Goyvaerts and L. Steven, Regular Expressions Cookbook. Beijing: O'Reilly, 2009. pp.1.
- [10] T. Nash. Accelerated C# 2010. New York, NY: Apress, 2010, pp.232.
- [11] Microsoft MSDN. NET Framework Regular Expressions, 2011.
- [12] M. Lutz, M. Alexander and G. Lars, International Workshop on Topic Maps Research and Applications. Leveraging the Semantics of Topic Maps. Berlin: Springer, 2007. pp.56.
- [13] R. Krohn, English Sentence Structure. Ann Arbor: University of Michigan Press, 1971.