Engineering Applications of Artificial Intelligence 41 (2015) 139–150

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

Predicting the need for vehicle compressor repairs using maintenance records and logged vehicle data

Rune Prytz [a,*,1], Sławomir Nowaczyk [b], Thorsteinn Rögnvaldsson [b], Stefan Byttner [b]

[a] Volvo Group Trucks Technology, Advanced Technology & Research, Göteborg, Sweden
[b] Center for Applied Intelligent Systems Research, Halmstad University, Sweden

ARTICLE INFO

Article history:
Received 11 June 2014
Received in revised form
17 January 2015
Accepted 11 February 2015

Keywords:
Machine Learning
Diagnostics
Fault Detection
Automotive Industry
Air Compressor

ABSTRACT

Methods and results are presented for applying supervised machine learning techniques to the task of predicting the need for repairs of air compressors in commercial trucks and buses. Prediction models are derived from logged on-board data that are downloaded during workshop visits and have been collected over three years on a large number of vehicles. A number of issues are identified with the data sources, many of which originate from the fact that the data sources were not designed for data mining. Nevertheless, exploiting this available data is very important for the automotive industry as means to quickly introduce predictive maintenance solutions. It is shown on a large data set from heavy duty trucks in normal operation how this can be done and generate a profit.

Random forest is used as the classifier algorithm, together with two methods for feature selection whose results are compared to a human expert. The machine learning based features outperform the human expert features, which supports the idea to use data mining to improve maintenance operations in this domain.

1. Introduction

Today, Original Equipment Manufacturers (OEMs) of commercial transport vehicles typically design maintenance plans based on simple parameters such as calendar time or mileage. However, this is no longer sufficient in the market and there is a need for more advanced approaches that provide predictions of future maintenance needs of individual trucks. Instead of selling just vehicles, the sector is heading towards selling complete transport services; for example, a fleet of trucks, including maintenance, with a guaranteed level of availability. This moves some of the operational risk from the customer to the OEM but should lower the overall cost of ownership. The OEM has the benefit of scale and can exploit similarities in usage and wear between different vehicle operators.

Predicting future maintenance needs of equipment can be approached in many different ways. One approach is to monitor the equipment and detect patterns that signal an emerging fault, which is reviewed by Hines and Seibert (2006), Hines et al. (2008a,b), and Ma and Jiang (2011). A more challenging one is to predict the Remaining Useful Life (RUL) for key systems, which is reviewed by Peng et al. (2010), Si et al. (2011), Sikorska et al. (2011) and Liao and Köttig (2014). For each of these approaches there are several options on how to do it: use physical models, expert rules, data-driven models, or hybrid combinations of these. The models can look for parameter changes that are linked to actual degradation of components, or they can look at vehicle usage patterns and indirectly infer the wear on the components. Data-driven solutions can be based on real-time data streamed during operation or collected historical data.

We present a data-driven approach that combines pattern recognition with the RUL estimation, by classifying if the RUL is shorter or longer than the time to the next planned service visit. The model is based on combining collected (i.e. not real-time) data from two sources: data collected on-board the vehicles and service records collected from OEM certified maintenance workshops. This presents a number of challenges, since the data sources have been designed for purposes such as warranty analysis, variant handling and financial follow-up on workshops, not for data mining. The data come from a huge set of real vehicles in normal operation, with different operators. The challenges include, among others, highly unbalanced data sets, noisy class labels, uncertainty in the dates, irregular readouts and unpredictable number of readouts from individual vehicles. In

* Corresponding author.
E-mail addresses: rune.prytz@volvo.com (R. Prytz), slawomir.nowaczyk@hh.se (S. Nowaczyk), thorsteinn.rognvaldsson@hh.se (T. Rögnvaldsson), stefan.byttner@hh.se (S. Byttner).

[1] Rune Prytz holds a licentiate in Signals and Systems engineering. He is active as an research engineer at AB Volvo and holds a position as Uptime systems & services Specialist.

http://dx.doi.org/10.1016/j.engappai.2015.02.009
0952-1976/

addition, multiple readouts from the same truck are highly correlated, which puts constraints on how data for testing and training are selected. We specifically study air compressors on heavy duty trucks and the fault complexity is also a challenge; air compressors face many possible types of failures, but we need to consider them all as one since they are not differentiated in the data sources.

Predictive maintenance in the automotive domain is more challenging than in many other domains, since vehicles are moving machines, often operating in areas with low network coverage or travelling between countries. This means few opportunities for continuous monitoring, due to the cost of wireless communication, bandwidth limitations, etc. In addition, both the sensors and computational units need to fulfil rigorous safety standards, which makes them expensive and not worth adding purely for diagnostic purposes. Those problems are amplified due to a large variety of available truck configurations. Finally, heavy duty vehicles usually operate in diverse and often harsh environments.

The paper is structured as follows. A survey of related works introduces the area of data mining of warranty data. This is followed by an overview of the data sets and then a methodology section where the problem is introduced and the employed methods are described. This is finally followed by a results section and a conclusion section.

1.1. Related work

There are few publications where service records and logged data are used for predicting maintenance needs of equipment, especially in the automotive industry, where wear prediction is almost universally done using models that are constructed before production.

In a survey of artificial intelligence solutions in the automotive industry, Gusikhin et al. (2007) discuss fault prognostics, after-sales service and warranty claims. Two representative examples of work in this area are Buddhakulsomsiri and Zakarian (2009) and Rajpathak (2013). Buddhakulsomsiri and Zakarian (2009) present a data mining algorithm that extracts associative and sequential patterns from a large automotive warranty database, capturing relationships among occurrences of warranty claims over time. Employing a simple IF–THEN rule representation, the algorithm filters out insignificant patterns using a number of rule strength parameters. In their work, however, no information about vehicle usage is available, and the discovered knowledge is of a statistical nature concerning relations between common faults. Rajpathak (2013) presents an ontology based text mining system that clusters repairs with the purpose of identifying best-practice repairs and, perhaps more importantly, automatically identifying when claimed labour codes are inconsistent with the repairs. Related to the latter, but more advanced, is the work by Medina-Oliva et al. (2014) on ship equipment diagnosis. They use an ontology approach applied to mining fleet data bases and convincingly show how to use this to find the causes for observed sensor deviations.

Thus, data mining of maintenance data and logged data has mainly focused on finding relations between repairs and operations and to extract most likely root causes for faults. Few have used them for estimating RUL or to warn for upcoming faults. We presented preliminary results for the work in this paper in an earlier study (Prytz et al., 2013). Furthermore, Frisk et al. (2007) recently published a study where logged on-board vehicle data were used to model RUL for lead-acid batteries. Their approach is similar to ours in the way that they also use random forests and estimate the likelihood that the component survives a certain time after the last data download. Our work is different from theirs in two aspects. First, a compressor failure is more intricate than a battery failure; a compressor can fail in many ways and there are many possible causes. Second, they also attempt to model the full RUL curve whereas we only consider the probability for survival until the next service stop.

Recently Choudhary et al. (2009) presented a survey of 150 papers related to the use of data mining in manufacturing. While their scope was broader than only diagnostics and fault prediction, they covered a large portion of literature related to the topic of this paper. Their general conclusion is that the specifics of the automotive domain make fault prediction and condition based maintenance a more challenging problem than in other domains; almost all research considers the case where continuous monitoring of devices is possible.

Jardine et al. (2006) present an overview of condition-based maintenance (CBM) solutions for mechanical systems, with special focus on models, algorithms and technologies for data processing and maintenance decision-making. They emphasise the need for correct, accurate, information (especially event information) and working tools for extracting knowledge from maintenance databases. Peng et al. (2010) also review methods for prognostics in CBM and conclude that methods tend to require extensive historical records that include many failures, even “catastrophic” failures that destroy the equipment, and that few methods have been demonstrated in practical applications. Schwabacher (2005) surveys recent work in data-driven prognostics, fault detection and diagnostics. Si et al. (2011) and Sikorska et al. (2011) present overviews of methods for prognostic modelling of RUL and note that available on-board data are seldom tailored to the needs of making prognosis and that few case studies exist where algorithms are applied to real world problems in realistic operating environments.

When it comes to diagnostics specifically for compressors, it is common to use sensors that continuously monitor the health state, e.g. accelerometers for vibration statistics, see Ahmed et al. (2012), or temperature sensors to measure the compressor working temperature, see Jayanth (2010). The standard off-board tests for checking the health status of compressors require first discharging the compressor and then measuring the time it takes to reach certain pressure limits in a charging test, as described e.g. in a compressor trouble shooting manual Bendix (2004). All these are essentially model-based diagnostic approaches where the normal performance of a compressor has been defined and then compared to the field case. Similarly, there are patents that describe methods for on-board fault detection for air brake systems (compressors, air dryers, wet tanks, etc.) that build on setting reference values at installment or after repair, see e.g. Fogelstrom (2007).

In summary, there exist very few published examples where equipment maintenance needs are estimated from logged vehicle data and maintenance data bases. Yet, given how common these data sources are and how central transportation vehicles are to the society, we claim that it is a very important research field.

2. Presentation of data

Companies that produce high value products necessarily have well-defined processes for product quality follow-up, which usually rely on large quantities of data stored in databases. Although these databases were designed for other purposes, e.g. analysing warranty issues, variant handling and workshop follow-up, it is possible to use them also to model and predict component wear. In this work we use two such databases: the *Logged Vehicle Data* (LVD) and the *Volvo Service Records* (VSR). In this work we have used data from approximately 65,000 European Volvo trucks, models FH13 and FM13, produced between 2010 and 2013.

2.1. LVD

The LVD database contains aggregated information about vehicle usage patterns. The values are downloaded each time a