Machine learning for monitoring the condition of critical systems

Applied Machine Learning (ML) is currently undergoing a boom, spurred on by availability of both data and computational power. Industrial equipment control and monitoring systems generate huge volumes of data. Through capture and analysis of this data, engineers are combining expert knowledge with ML to unlock the potential of this data. Decision support can be bolstered by improving the depth and breadth of information available to decision makers. Through a careful combination of expert elicitation and automated analysis of Condition Monitoring (CM) data one can streamline maintenance practices, moving closer to a fully Condition-Based Maintenance (CBM) approach. This presentation discusses types of ML analysis that are applicable to CM data along with presenting the challenges and benefits.

1 Introduction

Methods in Machine Learning (ML) in recent years have become more accessible accessible and prolific. Engineers are increasingly holding the ML lens to problems and assessing its suitability as a solution. The discipline of Prognostic Health Management (PHM) assesses applications of ML alongside traditional reliability modelling to provide deeper insight into causes and effects of equipment fault and degradation. Condition Monitoring (CM) data can be leveraged by ML methods to extract useful *information*. By application of suitable algorithms, predictions for the future evolution of the system can be made. This is primarily aimed at aiding decision support for Condition-Based Maintenance (CBM) and mission planning, ultimately improving reliability and availability of assets.

Applications in PHM are often focussed on high value asset domains where maintenance is expensive. Environments such as subsea, aerospace and outer space are areas where PHM can make significant contributions to decision support [1] and aid in reduction of risk.

2 Machine Learning

Machine learning methods generally fall into two categories: (1) Classification and (2) Regression. Numerous specific algorithms exist to carry out these general processes, each with nuanced selection criteria for specific problems. A reason these methods are becoming so widely applied across different sectors is that there exists a wealth of readily deployable software libraries openly available for performing this type of analysis.

Anomaly Detection is discussed as a third category though this is considered a subset of the former categories.

The selection of the methodologies for each application is a complex process which involves assessing the volume, resolution and types of historical data and data streams readily available for a system. A *"data poor"* situation will have quite different requirements to a *"data rich"* situation.

2.1 Classification

Classification deals with *labelling* data, this type of problem is synonymous with commonly encountered software features such as facial recognition and Automatic Number Plate Recognition (ANPR) systems where the input is a general *class* (i.e. photo) and the output is a *label* (i.e. number plate value).

The basic concept is taking a *training* dataset containing *labelled* examples of a specific category or *class*, a hypothesis function is formed which can then compare unseen examples and provide an output class, accompanied by a prediction accuracy. This methodology is useful in the recognition of common fault modes through use of a set of *known* fault data and compares features to infer the predicted fault mode.

2.2 Regression

Regression in ML, similarly to classical statistics, takes a set of datapoints and makes an inference about trends and projects them within some domain, commonly *time series* but this can be any domain i.e. frequency domain. In terms of condition monitoring this is used to learn from historical datapoints and predict features such as Time to Failure (TTF) and provide Remaining Useful Life (RuL) predictions for assets/components.

2.3 Anomaly Detection

A common problem encountered, particularly with *high-value* assets, is obtaining sufficiently many instances of faults in datasets. Robust designs coupled with well-designed maintenance regimes often preclude generation of such datasets [2]. Additionally expensive assets are often too costly to sacrifice for lifecycle testing to generate the volumes of data required. Anomaly detection provides a method allowing ML algorithms to be trained on readily available *steady-state* CM data.

Anomaly detection gives an indication that there is a problem with the asset, with further analysis being required to understand the nature of the fault. By combining Anomaly Detection with regression and classification these algorithms can be integrated into a decision support system.

3 Applications

The focus of our work is on automation of vibration analysis. Vibration analysis is a commonly used practice in assessing the proper functioning of mechanical machinery, particularly rotational equipment because it is an information rich data source and can be easily understood through basic signal processing methods such as Fast Fourier Transform (FFT). ML algorithms can be trained to observe how the frequency spectrum evolves over time and assess changes that may be indicative of system degradation. If good training data is available for this — through either lifecycle testing or historical profiles of operational assets — then the algorithm can use pattern recognition to compare the ongoing stream of operational data with the historical data to detect incipient fault and provide a fault mode estimation.

This information gives some indication of the overall *'health'* of an asset, and it is common for PHM practitioners to use fault datasets to construct a *'health index'* [3]. A health index is an indicator of the health of an asset often chosen to be 100% to 0% with 0% meaning failure. By imposing this health index on the asset life, algorithms can be trained to recognised what features in the data correspond to each health state and use this as a predictor.

In vibration terms it is common to set thresholds as a health indicator, though the absolute vibration may be exceeded during shock load and similar conditions. ML offers more intelligent methods to carry out this analysis. Learning the baseline operational frequency spectrum across a range of operational conditions mitigates misclassification of faults. Additionally feeding the system positional context about the frequency spectrum allows fault mode classification to take place.

4 Conclusions & Future Work

4.1 Benefits

Machine Learning (ML) carries potential to *automatically* perform in-depth *online* analysis that was previously not possible due to the volumes of data. This can be performed using embedded systems and/or centralised computing resources. High resolution measurements can provide valuable insight into the inner workings and minutiae of mechanical and electrical systems with *'interesting'* data being surfaced by automatic ML algorithms which can provide actionable information for asset operators whilst also providing valuable datasets for *offline* fault analysis and system operation optimisation.

The application of ML for PHM still requires specialist knowledge in understanding the fundamentals of the physical fault modes and mechanisms but this knowledge is imparted to the system when designing the algorithm allowing much wider application.

With the growth of PHM systems there is a real potential for changes, not only technical but contractual. The potential for use of Contracts for Availability (CfA) is vastly increased by the increased understanding and knowledge permitted by the asset state estimation available from these advanced monitoring algorithms. This is well known in the commercial aviation sector where major manufacturers have transformed how they sell jet engines and ultimately their business strategy [4].

4.2 Future Work

CM data is a source of the *outputs* of a system but there is also potential to make use of the control system *inputs* to create a holistic PHM system which also takes into account the control data and learns relationships between systems throughout the life of an asset, a concept known as *Lifelong Machine Learning* [5]. Future work will look into the *fusion* of the control data and CM data for further improvement in fault predictions and to investigate how different usage profiles affect the asset condition.

Currently the bulk of the work in this field focusses on individual assets or component parts though the potential is much greater than this. Further works will investigate the integration of this information into Integrated Platform Management System (IPMS) [6] and similar systems.

Integration of fault predictions into autonomous systems is a related area of research undergoing development. Autonomous agents such as automated drones often carry out some level of mission planning on board by themselves. The fault predictions provided by PHM can allow the state of the system and any impending faults to be taken into account in mission planning.

References

[1] N.-H. Kim, D. An, and J.-H. Choi, "Prognostics and health management of engineering systems," *Switzerland: Springer International Publishing*, 2017.

[2] H. L. Resnikoff, "Mathematical Aspects of Reliability-Centred Maintenance," 1978.

[3] A. Saxena, J. Celaya, B. Saha, S. Saha, and K. Goebel, "Metrics for offline evaluation of prognostic performance," *International Journal of Prognostics and Health Management*, vol. 1, no. 1, pp. 4–23, 2010.

[4] D. J. Smith, "Power-by-the-hour: the role of technology in reshaping business strategy at rolls-royce," *Technology Analysis & Strategic Management*, vol. 25, no. 8, pp. 987–1007, 2013.

[5] Z. Chen, B. Liu, R. Brachman, P. Stone, and F. Rossi, *Lifelong Machine Learning: Second Edition*, ser. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2018.

[6] M. J. Roemer, C. S. Byington, G. J. Kacprzynski, and G. Vachtsevanos, "An overview of selected prognostic technologies with reference to an integrated phm architecture," in *In Proceedings of the First International Forum on Integrated System Health Engineering and Management in Aerospace, Big Sky*, 2005, pp. 3941–3947.

Speaker Biography

Ross Dickie graduated from the University of Edinburgh with a BEng(Hons) in Electrical & Mechanical Engineering in 2014 and received an MSc in Renewable Energy & Distributed Generation in 2016 from Heriot Watt University. He is currently a PhD researcher at Heriot Watt University as part of the Smart Systems Group and co-sponsored by MacTaggart Scott in Edinburgh, Scotland. His current research investigates the applications of advanced data analysis within industrial systems aiming to identify actionable information and gain fault insights.