Data-driven predictive maintenance and time-series applications

Kefalas, M.

Chapter 1

Introduction

1.1 Background

"Prediction is very difficult, especially if it's about the future!" With these words, Nobel laureate in physics Niels Bohr managed to describe in a humorous way the inherent difficulty of knowing what will happen in the future. Mankind has always been intrigued and challenged to know what to anticipate. Having this knowledge, one can better prepare against what is about to occur, thus avoiding any - usually unwanted - surprises, such as a disaster. Driven by the power that such knowledge can have, humans have managed to devise various methods to foresee what *could* happen and when. In antiquity, for example, the Oracle of Delphi in Greece would consult kings and emperors on military campaigns by foretelling (in a rather cryptic way) who, for instance, would prevail in battle through the use of fumes and vapors [29, 72, 1]. The Babylonians would predict the weather from cloud patterns and astrology [137, 172] and people in the Middle Ages through what came to be known as weather lore [33]. In Christianity and other religions, prophets would try to read the signs of the times to prepare people for the ominous days that were approaching. Even though, a lot of time has passed since those days, human nature has remained the same. Nowadays, of course, human interests and needs have shifted, but the requirement on what to expect is still necessary and pivotal to the existence of our present-day living.

Modern societies rely heavily on energy, transportation, and (heavy) industry to function smoothly. As a result, maintenance of these systems is of paramount importance, in order to maximize safety, usage, and to minimize (unnecessary) downtime, as this can save superfluous costs and nuisance. Inevitably, the need to foresee/predict when a system will require maintenance has gained popularity since the 1990's, with a slow shift from the classic corrective or reactive maintenance (RM) and scheduled or preventive maintenance (PM) policies towards what is known as predictive maintenance (PdM) [103, 184, 160]. The former two - still predominant - types of maintenance either perform corrective actions when they are needed (e.g., when a system failure has already occurred) or in regular time intervals in order to prevent failures (based on e.g., maintenance history of similar assets), respectively. At a closer look, however, these maintenance strategies naturally introduce significant extra costs due to machine downtime, component replacement or even unnecessary maintenance interventions. In contrast to these maintenance polices, PdM through the careful monitoring of the system in question (condition-based maintenance or CBM), determines its state/condition, in order to predict when maintenance should be performed, overcoming the inefficiencies of the aforementioned methods. This way, maintenance can be planned accordingly and in advance by activating needed protocols, preparing the necessary logistics (e.g., personnel, spare parts) and by optimizing the overall maintenance schedule. PdM achieves this by estimating what is known as the remaining useful life (RUL) of an asset. The RUL determines the time remaining until the system in question exceeds its normal operational envelope and is no longer deemed useful.

The available technologies and applications that enable and are associated with PdM are typically called prognostics and health management (PHM). RUL as the means to an end, lies in the heart of PHM and has thus received a lot of attention from the academic and industrial communities. In general, there are three major classes of approaches which deal with the RUL estimation. Namely, model-based, data-driven, and hybrid methods [59, 227, 21, 184, 156]. Model-based methods (or physics-based methods) rely on established mathematical/physical models of the asset and degradation process and consequently require a thorough understanding of the asset's and degradation process's physics. This can sometimes prove to be prohibitively costly, in terms of time and money, especially for very complex systems. Data-driven methods are, on the other hand, relatively easier to develop as they do not call for (a lot of) expert or domain knowledge. However, they do require large amounts of (maintenance) data, a constraint that is not always possible due to the nature of certain applications. Machine learning (ML) and deep learning (DL) commonly belong to this category. Lastly, hybrid (or fusion) methods attempt to leverage the advantages of the two previous methods while minimizing their limitations by combining (or fusing) model-based approaches and augmenting them with data-driven methods and vice versa.

Despite their differences, all methods, in general, make some use of the sensor data generated by the assets for machine monitoring and/or it's maintenance history. This has been made possible due to the progress of (big) data, machine learning, deep learning and the related fields and applications (e.g., internet-of-things), the wide availability of sensors [177], the increase of computational power and storage (e.g., cloud computing), and the reduction in the prices of these. Typically, the generated data are recorded over time, as a sequence of measurements, known as time-series. All methods use the generated data in order to either develop the equations of the underlying physical process in the model-based approaches or to extract knowledge in order to make inferences in the data-driven approaches.

1.2 Objectives

The estimation of the RUL is, in general, an inherently challenging problem. The remaining life is not simply a target variable that can be predicted from the time-series sensor measurements. It is a variable that needs to be inferred from a longer trend of degradation patterns and when those appear. In the data-driven domain specifically, there exist various challenges, such as the correct pre-processing of the data, the definition of the RUL variable in supervised settings, as well as the selection of the appropriate learning algorithm to be used. Generally, the design choices of the algorithm and its respective hyperparameters, next to the choices that need to be made during pre-processing, make this task challenging for end-users. Apart, however, from the technical difficulties of the matter, further valuable objectives arise.

A point of key importance is to be able to quantify the confidence of the estimation. Determining the RUL simply as a point-wise estimation gives no information about the uncertainty of our prediction. For example, how sure are we that our car will need a new alternator in 6 weeks? After all, when dealing with safety-critical and operations-critical questions we would like to be sure (as sure as one can be) about the future. Furthermore, it is also desirable to be able to predict the chance that a machine operates without a fault or failure up to some future time. This way, maintenance personnel can determine whether the inspection interval is appropriate or not.

This work is devoted to the usage of data-driven methods for the estimation of the RUL in the context of PdM. Therefore, an extensive study of data-driven methods in PdM is conducted to point out the shortcomings, the advancements in the field, as well as opportunities for future work. Based on this, we developed pipelines with more or less complex methods, such as Automated Machine Learning (AutoML), Deep Neural Networks (DNNs), hyperparameter optimization (HPO) and combinations of these to tackle said questions.

In addition, we investigate a method to perform multi-objective scheduling optimization of a job-shop, which is an appropriate next step of the RUL estimation. After all, apart from safety, schedule optimization is another major reason for the existence of PHM and for the notion of RUL specifically. Beyond that, this work investigates an application of interpretable artificial intelligence (AI) through the use of genetic programming (GP) on a real-world problem from the aviation industry. It further proposes extensions and ideas, in order to incorporate similar approaches in the field of PHM. Finally, given the fact that this work deals mainly with time-series data, we examine a real-world problem from the medical domain concerning classification. For this task, we make use and extend an earlier implementation of a data-driven pipeline originally developed for the automotive industry, exhibiting the significance of knowledge transfer and how it can impact research between different scientific fields.

The relevant research questions (RQ) of this work are stated in the following:

RQ1 What are the advancements, drawbacks, and opportunities in PHM and specifically of data-driven PdM in the aerospace industry?

RQ2 Can automated machine learning (AutoML) methods be applied for the estimation of the RUL in data-driven PdM?

RQ3 Can we propose an automated framework for configuring RUL prediction models which are highly accurate and have less estimation uncertainty?

RQ4 Can explainable AI facilitate the understanding of the data generating process of industrial processes?

RQ5 Can tabu search (TS) support the solution of the multi-objective flexible job-shop scheduling problem (FJSSP)?

RQ6 Can time-series techniques from industry lend themselves to applications in the medical domain?

The work presented in this thesis has been carried out in the context of the NWO (Nederlandse Organisatie voor Wetenschappelijk Onderzoek) project, CIMPLO (Cross-Industry Predictive Maintenance Optimization Platform). The CIMPLO-project aims at developing a cross-industry predictive maintenance optimization platform, which addresses the real-world requirements for dynamic, scalable multiple-criteria maintenance scheduling. Amongst the industry partners involved in the CIMPLO-project, KLM (Koninklijke Luchtvaart Maatschappij) and Honda Research Institute Europe have been involved in the work presented in this dissertation.

1.3 Outline of the Dissertation

In this Section, the outline of this dissertation is presented. Furthermore, we briefly describe the motivation and contents of each following chapter.

- Chapter 2 provides a brief overview of the notions used in the following chapters. Its aim is to acclimatize the reader with the terminology used and serve as a quick reference guide.

- Chapter 3 provides a general introduction to the reader of the fields of PHM and PdM through definitions and examples, and defines the notion of the RUL. It further emphasizes on data-driven applications of PdM in the aerospace industry due to the criticality that PHM has in it and its future. Drawbacks and future research suggestions are also presented. **RQ1** is addressed in this chapter. Contents of this chapter are (partly) based on [230]; Duc van Nguyen, Marios Kefalas, Kaifeng Yang, Asteris Apostolidis, Markus

Olhofer, Steffen Limmer, and Thomas Bäck. A Review: Prognostics and Health Management in Automotive and Aerospace. International Journal of Prognostics and Health Management, 10(2):35, 2019.

- Chapter 4 serves to discuss the difficulties and challenges that accompany the RUL prediction task in the data-driven domain. In this view, it introduces to the reader the notion and necessity of automated machine learning (AutoML). The presented method is validated on a simulated dataset for turbofan engines. **RQ2** is addressed in this chapter. Contents of this chapter are (partly) based [112]; Marios Kefalas, Mitra Baratchi, Asteris Apostolidis, Dirk van den Herik, and Thomas Bäck. Automated Machine Learning for Remaining Useful Life Estimation of Aircraft Engines. In 2021 IEEE International Conference on Prognostics and Health Management (ICPHM), pages 1–9, Detroit (Romulus), MI, USA, June 2021. IEEE.

- Chapter 5 deals with a topic of high significance in PHM and specifically in data-driven PdM, namely uncertainty quantification. The motivation behind this chapter lies in the fact that in addition to the RUL prediction, one needs to assess also the confidence of that prediction. This is especially crucial in operations-critical and safety-critical applications, where an indication of the remaining time until failure should be as confident as possible. **RQ3** is addressed in this chapter. Contents of this chapter are (partly) based on [116]; Marios Kefalas, Bas van Stein, Mitra Baratci, Asteris Apostolidis, and Thomas Bäck. An End-to-End Pipeline for Uncertainty Quantification and Remaining Useful Life Estimation: An Application on Aircraft Engines. In 2022 7$^{\text{th}}$ European Conference of the Prognostics and Health Management Society, Turin, Italy, July 2022. PHM Society.

- Chapter 6 discusses the importance of explainability in data-driven methods in PHM. Through a case study with real-world data from the aerospace industry we motivate the criticality of explainability, the advantages that accompany such methods, and future research directions. Additionally, the notion of symbolic regression (SR) is presented and how it can be used in prognostics. **RQ4** is addressed in this chapter. Contents of this chapter are (partly) based on [113]; Marios Kefalas, Juan de Santiago Rojo, Asteris Apostolidis, Dirk van den Herik, Bas van Stein, and Thomas Bäck. Explainable Artificial Intelligence for Exhaust Gas Temperature of Turbofan Engines. Journal of Aerospace Information Systems, pages 1–8, 2022. Publisher: American Institute of Aeronautics and Astronautics eprint: https://doi.org/10.2514/1.I011058.

- Chapter 7 deals with the next step that arises in PHM/PdM applications. Scheduling. The idea presented here, is how to optimally schedule the maintenance of assets. Specifically, in this chapter we introduce and deal with the flexible job-shop scheduling problem

(FJSSP), an NP-hard problem that closely resembles real maintenance settings. We discuss its difficulties and propose a method to solve its multi-objective variant. **RQ5** is addressed in this chapter. Contents of this chapter are (partly) based on [115]; Marios Kefalas, Steffen Limmer, Asteris Apostolidis, Markus Olhofer, Michael Emmerich, and Thomas Bäck. A tabu search-based memetic algorithm for the multi-objective flexible job shop scheduling problem. In Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO '19, page 1254–1262, New York, NY, USA, 2019. Association for Computing Machinery.

- Chapter 8 exhibits how time-series techniques from industry can lend themselves to applications in the medical domain. Specifically, we show that a time-series pipeline that originated in the automotive industry can facilitate the diagnostic process in the field of Neurology. **RQ6** is addressed in this chapter. Contents of this chapter are (partly) based on [114]; Marios Kefalas, Milan Koch, Victor Geraedts, Hao Wang, Martijn Tannemaat, and Thomas Bäck, Automated Machine Learning for the Classification of Normal and Abnormal Electromyography Data, 2020 IEEE International Conference on Big Data (Big Data), 2020, pp. 1176-1185. IEEE.

- Chapter 9 concludes the work of this thesis by summarizing the answers to the research questions and providing suggestions for future work.

1.4 Author's Contributions

Below we show, in chronological order, a list of publications that the author has contributed to.

[1] Asep Maulana, **Marios Kefalas**, and Michael Emmerich,"Immunization of networks using genetic algorithms and multiobjective metaheuristics", 2017 IEEE Symposium Series on Computational Intelligence (SSCI), 2017, pp. 1-8. IEEE. DOI: `https://doi.org/10.1109/SSCI.2017.8285368`.

[2] **Marios Kefalas**, Steffen Limmer, Asteris Apostolidis, Markus Olhofer, Michael Emmerich, and Thomas H. W. Bäck. 2019. "A tabu search-based memetic algorithm for the multi-objective flexible job shop scheduling problem". In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '19). Association for Computing Machinery, New York, NY, USA, 1254–1262. DOI:`https://doi.org/10.1145/3319619.3326817`.

[3] Duc van Nguyen, **Marios Kefalas**[1], Kaifeng Yang, Asteris Apostolidis, Markus Olhofer,

[1]Duc van Nguyen and Marios Kefalas have contributed equally to this publication (co-first authorship).

Steffen Limmer, and Thomas Bäck. "A Review: Prognostics and Health Management in Automotive and Aerospace", International Journal of Prognostics and Health Management, 10(2):35, 2019. PHM Society. DOI:`https://doi.org/10.36001/ijphm.2019.v10i2.2730`.

[4] **Marios Kefalas**, Milan Koch, Victor Geraedts, Hao Wang, Martijn Tannemaat, and Thomas Bäck, "Automated Machine Learning for the Classification of Normal and Abnormal Electromyography Data", 2020 IEEE International Conference on Big Data (Big Data), 2020, pp. 1176-1185. IEEE. DOI:`https://doi.org/10.1109/BigData50022.2020.9377780`.

[5] Michael Emmerich, Joost Nibbeling, **Marios Kefalas**, and Aske Plaat, "Multiple node immunisation for preventing epidemics on networks by exact multiobjective optimisation of cost and shield-value", arXiv:`https://arxiv.org/abs/2010.06488`, 2020.

[6] **Marios Kefalas**, Mitra Baratchi, Asteris Apostolidis, Dirk van den Herik, and Thomas Bäck, "Automated Machine Learning for Remaining Useful Life Estimation of Aircraft Engines", 2021 IEEE International Conference on Prognostics and Health Management (ICPHM), 2021, pp. 1-9. IEEE. DOI:`https://doi.org/10.1109/ICPHM51084.2021.9486549`.

[7] Victor J. Geraedts, Milan Koch, Roy Kuiper, **Marios Kefalas**, Thomas Bäck, Jacobus J. van Hilten, Hao Wang, Huub A.M. Middelkoop, Niels A. van der Gaag, Maria Fiorella Contarino, and Martijn R. Tannemaat (2021), "Preoperative Electroencephalography-Based Machine Learning Predicts Cognitive Deterioration After Subthalamic Deep Brain Stimulation", Moving Disorders, 36: 2324-2334. Wiley Periodicals LLC. DOI:`https://doi.org/10.1002/mds.28661`.

[8] **Marios Kefalas**, Juan de Santiago Rojo, Asteris Apostolidis, Dirk van den Herik, Bas van Stein, and Thomas H. W. Bäck. "Explainable Artificial Intelligence for Exhaust Gas Temperature of Turbofan Engines", Journal of Aerospace of Information Systems (JAIS), Volume 19, Issue 6, pages 447-454, 2022. American Institute of Aeronautics and Astronautics. DOI: `https://doi.org/10.2514/1.I011058`.

[9] Fan Yang, **Marios Kefalas**, Milan Koch, Anna V. Kononova, Yanan Qiao, and Thomas H. W. Bäck, "Auto-REP: An Automated Regression Pipeline Approach for High-efficiency Earthquake Prediction Using LANL Data", 14th International Conference on Computer and Automation Engineering (ICCAE), 2022, pp. 127-134. IEEE. DOI:`https://doi.org/10.1109/ICCAE55086.2022.9762437`.

[10] **Marios Kefalas**, Bas van Stein, Mitra Baratchi, Asteris Apostolidis, and Thomas H. W. Bäck, "An End-to-End Pipeline for Uncertainty Quantification and Remaining Useful Life Estimation: An Application on Aircraft Engines", 7th European Conference of the Prognostics and Health Management Society (PHME22), 2022. PHM Society. DOI:`https://doi.org/10.36001/phme.2022.v7i1.3317`.

[11] Martijn R. Tannemaat, **Marios Kefalas**, Victor J. Geraedts, Linda Remijn-Nelissen, Anna Verschuuren, Milan Koch, Anna V. Kononova, Hao Wang, and Thomas H.W. Bäck, "Distinguishing normal, neuropathic and myopathic EMG with an automated machine learning approach", Clinical Neurophysiology, pages 1-17, 2022. Elsevier. (*Submitted*)

[12] Michael Emmerich, Yulian Kuryliak, Dmytro Dosyn, Ksenia Pereverdieva, Joost Nibbeling and **Marios Kefalas**, "Multi-objective Targeted Immunization: Non-backtracking vs. Adjacency Matrix", 15th Workshop on Global Optimization (HUGO 2022), pages 1-4, 2022. International Society of Global Optimization.