Machine Learning and Knowledge Discovery for Engineering Systems Health Management

Introductory Chapter

Ashok N. Srivastava and Jiawei Han

As the complexity of engineered systems grows it is becoming more imperative than ever to develop tools and technologies that can manage the complex interactions within and between these systems to maintain a high degree of reliability. In many cases these systems are instrumented with sophisticated sensors that are used to monitor and in some cases control the system. This book explores the development of state-of-the-art tools and techniques that can be used to automatically detect, diagnose, and in some cases, predict the effects of adverse events in an engineered system on its ultimate performance. This gives rise to the field Systems Health Management, in which methods are developed with the express purpose of monitoring the condition, or 'state of health' of a complex system and through automated means detecting faults, failures, and abnormalities in the system condition. If such an anomaly is detected, a diagnosis, or discovery of the root cause of the anomaly is usually necessary in order to enact an appropriate mitigation strategy. In some cases, it may become essential to assess, again in an automated fashion, the impact of the anomaly on overall system performance. This relies on an estimate of the remaining useful life of the system under study.

For the purposes of this discussion, there are fundamentally two approaches that can be taken to address systems health management. The first approach has been taken and explored in the engineering community for many decades, and that is to build high-fidelity, physics-based models of the system under study. These physics-based models are generally built on sets of coupled, potentially nonlinear partial differential equations that describe the evolution of the system through time and space given an initial condition and a set of forcing functions. The models are based on first principles and thus, by definition, obey the laws of physics from the outset. For example, in an engine simulation, the laws of thermodynamics, Newton's laws, and other fundamental physical understandings are implicitly and explicitly expressed. This degree of fidelity to physical laws gives these methods enduring power and predictive capability.

At the other extreme, one can take a purely statistical or data-driven approach, in which one takes the data obtained from sensors on a real system (or the output of simulated sensor readings from an engineering simulation) and builds a statistical model that relates the inputs to the system with the measured outputs. This model could be based on a potentially complex, nonlinear mapping of inputs to outputs that is estimated through the minimization of a loss function. This approach, based on machine learning, is extremely powerful in that it uses the observed data to make an internal model of the system under study. In many cases, this model does not lend itself to physical interpretability, but it can have predictive power to estimate the evolution of the system in both time and space.

Certainly the view presented above is simplified: many machine learning systems encode knowledge of the system under study, and similarly, many physics-based systems can adapt to data as it is observed. Most Bayesian approaches to machine learning fundamentally encode a statistical model that describes a system based on information prior to the arrival of data and then adapts the model as new data is observed. The degree to which they encode physical laws varies based on the nature of the model and its flexibility. However, the bridge between physics-based approaches and data-driven approaches needs to be explored further in the context of systems health management. This book attempts to

show fundamental algorithms and their operation in a wide-variety of domains that can further help bridge this divide.

Many engineering systems differ fundamentally from other complex systems in that they obey physical principles. We note, however, that not all engineered systems must obey physical laws. For example, a complex software system may have no relation to physical laws and thus may not require the same type of modeling as a hardware-based system. If a software system is simulating a physical system, those physical laws must be encoded in the software with sufficient fidelity.

Another aspect of engineering systems is that there are often maintenance or other types of reports written about them that can describe off-nominal behavior and the events that lead up to such behavior. These documents can be in the form of maintenance reports or usage reports that are written by different authors. For example, in the case of aviation and space systems, there can be extensive documentation that describes fault and failure modes that are written when anomalies are detected, diagnosed, and then when a mitigation strategy is employed.

The subject matter of this book therefore spans many dimensions. Part I begins with five chapters describing data-driven methods for anomaly detection, diagnosis, and prognostics for analyzing massive data streams and associated performance metrics. It also describes the analysis of text reports using novel approaches in machine learning to help detect and discriminate between failure modes. Part II covers physics-based methods for diagnostics and prognostics, in which these methods adapt to observed data. The chapters in this section cover physics-based, data-driven, and hybrid approaches to studying damage propagation and prognostics with applications in composite materials and solid rocket motors. Part III discusses application domains such as distributed data centers, aircraft engines, and embedded real-time software systems. Given the interdisciplinary nature of the field, there is considerable overlap in these dimensions as reflected in the chapters in each section.

- **Part I: Data-Driven methods for Systems Health Management.** This part covers data-driven methods for anomaly detection, diagnosis, and prediction and methods to analyze massive data streams arising from complex systems with associated performance metrics
    - Mining Data Streams: Systems and Algorithms, by Aggarwal and Turaga
    - A Tutorial on Bayesian Networks for Systems Health Management: Choi, Darwiche, Zheng, and Mengshoel
    - Heterogeneous Fleetwide Anomaly Detection: Oza and Das
    - Discriminative Topic Models: by Shan, Agovic, and Banerjee
    - Prognostic Performance Metrics: Goebel, Saxena, Celaya, Saha, and Saha
- **Part II: Physics-based methods for Systems Health Management.** This part discusses physics-based diagnostic and prognostic methods which incorporate and adapt to observed data
    - Gaussian Process Damage Prognosis under Random and Flight Profile Loading, Chattopadhyay and Mohanty
    - Fatigue Damage Prognosis Updating under Uncertainties: Guan and Liu
    - Physics-based methods of failure analysis and diagnostics in human space flight, Smelyanskiy, Luchinsky, Hafiychuk, Osipov, Patterson-Hine, and Hanson
    - Model-based Tools and Techniques for Real-Time System and Software Health Management, Abdelwahed, Dubey, Karsai, and Mahadevan
- **Part III: Applications.** Application domains such as distributed computer data centers, aircraft engines, and embedded real-time software systems are discussed.

- Real-time identification of performance problems in large distributed systems, Goldszmidt, Woodard, and Bodik
- A Combined Model-Based and Data-Driven Prognostic Approach for Aircraft System Life Management, Orchard, Vachtsevanos, and Goebel.
- Hybrid Models for Propulsion Health Management, Volponi and Rajamani
- Extracting Critical Information from Free Text Data for Systems Health Management, Kao, Poteet, and Augustine

Part I: Data-Driven Methods for Systems Health Management

Chapter 1: Aggarwal and Turaga begin this portion of the book by giving a comprehensive treatment on the subject of monitoring, managing, and extracting real-time information from massive data streams. This is central to the theme of the book since many complex engineering systems are instrumented with high frequency sensors producing enormous amounts of information. The authors discuss the architecture of a distributed stream mining system and also overview algorithms for classic data mining and machine learning tasks such as classification, clustering, pattern extraction, and time series analysis. The authors also describe methods to deal with non-stationary systems which are prevalent in application domains.

Chapter 2: The second chapter of this book gives a tutorial on Bayesian networks, which comprise the backbone of many probabilistic reasoning systems for diagnosing faults in complex systems. The chapter discusses this modeling technique in light of a set of examples meant to show the use of Bayesian networks for modeling and learning in the context of systems health management. The authors describe how these networks naturally model uncertainty due to measurement and modeling and discuss reasoning and learning in complex systems. Choi, et al. conclude this chapter with a discussion of a complex real-world application to electrical power systems.

Chapter 3: A key feature of many, but not all, engineering systems is that there may be numerous copies of engineering systems operating in different environments. In this chapter, Oza and Das consider the problem of mining data from fleets of systems (such as aircraft) where data from continuous sensors, discrete switches, and text reports must be combined to assess the overall health of the system. The chapter features an in-depth discussion of previous attempts to mine data from fleet-wide sources and then discusses novel statistical and kernel-based algorithms for anomaly detection of the data types that arise from fleets. The chapter concludes with a discussion of the impact of these methods on safety critical systems.

Chapter 4: The majority of this book primarily concerns itself with analysis and interpretation of numeric data coming from hardware or software systems. This chapter differs in that it addresses the analysis of text documents. Although text mining is a vast area of research by itself, we included two chapters on the topic area because they address key problems that have applicability both to systems health management and machine learning. Shan et al. discuss discriminative topic models, which are a family of models that simultaneously address the problem of grouping documents with similar topics together in an automated fashion and then classifying these documents into different categories. These models are shown on the real-world data from the Aviation Safety Reporting System to be competitive with existing methods with the additional benefit of extracting interesting topics.

Chapter 5: For many machine learning algorithms, the performance metric is specified and the learning system then optimizes the performance metric based on the observed data. This chapter discusses performance metrics for prognostics, where one estimates the remaining useful life of a component in an engineering system. The chapter gives an overview of the prognostics sub-discipline of systems health management and then discusses the development of performance metrics that are tailored to the needs of the prognostics community. Issues such as end-user requirements, non-stationarity, and different time scales are also discussed. The chapter features a set of guidelines on choosing different prognostic metrics based on both user requirements and the characteristics of the observed data.

Part II: Physics-based methods for Systems Health Management

Chapter 6: In contrast to the topics in Part I, the chapters in Part II feature discussions of methods that are physics-based but also incorporate a data-driven element. Chattophadyay and Mohanty discuss a novel application of Gaussian Process Regression to the problem of estimating fatigue damage propagation using these methods but also give a theoretical understanding of the physics underlying fatigue damage prognosis. The chapter gives a clear introductory treatment of Gaussian Processes and then shows the performance of these methods on data from simulations as well as real-world experiments. A key issue that the chapter addresses is the fact that these materials are subjected to different loading conditions. The authors demonstrate how these methods can accommodate loadings that are observed in flight as well as random loading.

Chapter 7: Guan and Liu provide a comprehensive treatment to the problem of fatigue damage prognosis using probabilistic methods to appropriately model uncertainties. A probabilistic prognosis framework is given which shows how prior probability distributions, simulation methods and response measures can be combined to help estimate remaining useful life of structures. They show a hierarchical Bayesian method for modeling numerous sources of uncertainty including uncertainties due to modeling, variability, and measurement error. The authors show how a Bayesian approach can be used to model prior information and can be updated as new data is observed. This approach can lead to computational burdens which are also addressed through the development of a decoupled Markov Chain Monte Carlo sampler. The authors demonstrate their technique on experimental data.

Chapter 8: The dynamics of an engineering system guide its evolution and can thus play a central role in diagnosing an observed anomaly. In Chapter 7, Smeliyanskiy et al. show a novel method to learn a dynamical inference of stochastic nonlinear models. The analytical approach is based on path integrals and a method is shown to infer model parameters in a dynamical noise environment. The research results are presented on the familiar Lorenz chaotic system and then extended to a discussion of an application of these methods to fault diagnosis of solid rocket motors. The authors discuss further research for using dynamical inference algorithms to predict fault dynamics in the solid rocket motors.

Chapter 9: Although many examples given in this book are related to physical engineering systems, we also treat the issues of detecting faults in software systems. Software systems differ in numerous ways from hardware systems. A key distinction in the context of this book is that most software systems do not obey physical laws and thus require a new approach to modeling. This chapter discusses an approach known as Time Failure Propagation Graphs which are traditionally applied to hardware systems but are adapted to address software systems. The formal structure of the TFPG is discussed and a reasoning algorithm to perform diagnosis on software systems is presented. The authors also address the issue of prognostics of impending faults.

Part III: Applications

Chapter 10: The Applications part of the book focuses on the use of machine learning and physics-based approaches to analyze complex systems. The first chapter in this part discusses an approach to real-time identification of performance problems in large distributed data centers. These complex systems are becoming commonplace in their implementation at major internet companies and have extremely high reliability requirements. These requirements pose major technical challenges, one of which is to rapidly determine whether an observed problem has been seen before or whether this is a new problem. The authors describe this so-called crisis identification problem and provide a method to extract the signature of these problems and then discuss a full Bayesian approach to clustering these signatures. The results of applying these methods on a real-world data center are also discussed.

Chapter 11: Orchard et al. discuss the particle filter algorithm and its application to estimating the probability density function of a state vector in a system with known dynamics. This probability density function is updated as new data is observed. The technique is described in the context of prognostics which is the approach to estimating the remaining useful life of a component. A key issue that arises in prognostics is modeling the effects of uncertainty due to model misspecification, measurement noise, and other sources of uncertainty in the estimate of remaining useful life. The authors present the algorithms and discuss the results of their methods on the problem of fatigue crack growth.

Chapter 12: Volponi and Rajamani describe hybrid models for analyzing aircraft engines to detect, diagnose, and predict the future trends of system faults. The paper takes an approach of combining physics-based and data-driven techniques to garner the benefit of both approaches. The combined approach employs physics-based models to learn operating set points rapidly, and then data-driven techniques to adapt to the change in operating point. The chapter features a description of a typical onboard engine model architecture and then a contrasting hybrid architecture, and accentuates the differences between these two approaches. A key contribution of this chapter is the inclusion of a discussion on verification and validation of hybrid models in the engine community and the requirements that are given on such models from major regulators.

Chapter 13: A key application area of text mining in the context of systems health management is in understanding the linkage between maintenance reports and free narratives written by the users of engineering systems. This chapter discusses an approach to linking information from maintenance reports regarding aircraft and narratives from the Aviation Safety Reporting System. The paper discusses an approach that combines natural language processing and text mining techniques to address this goal.

This book is the culmination of the hard work of numerous individuals including, first, and foremost the chapter authors. Their dedication and willingness to share their insights is invaluable. We also thank the reviewers for key insights and contributions and Randi Cohen at Taylor and Francis for her enduring help. A. N. Srivastava wishes to thank the NASA Aviation Safety Program System-Wide Safety and Assurance Technology Project for supporting this work.

Dedication

ANS dedicates this book to his late father, Dr. Jagdish N. Srivastava, whose guidance, broad vision, and fundamental insights helped shape him.