Model Based and Big Data Enabled Predictive Maintenance Capability Development Experience

Mark Mazarek, Changzhou Wang, Darren Macer

The Boeing Company, Seattle, Washington, 98124, USA
mark.a.mazarek@boeing.com changzhou.wang@boeing.com darren.b.macer@boeing.com

ABSTRACT

Airplane health management and predictive maintenance have been in place for well over a decade within the aerospace industry. Predictive maintenance leaders have unique challenges in developing talent pipelines, technology focus areas, balanced with delivery of prognostic insights to customers to improve their operational efficiency. This paper discusses strategies, metrics, lessons learned for building, growing, and sustaining a team of engineers, data scientists, and software developers, with a focus on delivering aerospace prognostic insights.

1. INTRODUCTION

The development of commercial aerospace prognostics is an impactful way to improve airline efficiency and operations. The success of such insights spans multiple complex Original Equipment Manufacturers (OEM), airline technical operations teams and suppliers. This fact demands that the associated partners work together to define, build, validate and deploy prognostics.

For a prognostic capability to be successful, a clear vision along with desired outcomes must be shared with stakeholders, as it will formulate role clarity, working together agreements. The following sections will step through the predictive maintenance capabilities, processes, and metrics, along with skills required to build a predictive maintenance capability.

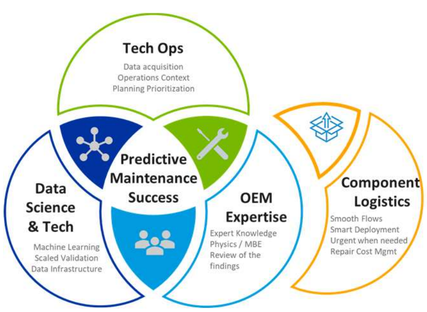

Figure 1: Predictive Maintenance Venn Diagram

2. PREDICTIVE MAINTENANCE CAPABILITY & PROCESS

Expanding on the understanding of the predictive maintenance Venn diagram, we believe it is critical for predictive maintenance success to have a capability consisting of the following personas:

1. System and Hardware Engineers
2. Data Engineers & Data Scientists
3. Software Developers
4. Reliability Engineers

In addition to the personas above, it is also critical for success to have a data platform that meets the requirements of the data science team, and has the flexibility to deploy insights into interfacing systems both internal and customer facing. It is also important that the core predictive maintenance capability interface with organizational stakeholders including but not limited to legal, contracts, information technology teams.

Figure 2 Predictive Maintenance Capability conceptualizes data ingestion, correlation, processing, and deployment interfaces. It is also important to establish data rights agreements, allowing customer understanding of data usage, which will build trust through transparency on data usage and analysis. The capability must include third party airline and component supplier interfaces. A robust data platform includes capabilities to enable data integrity, data quality checks, feedback loops for data generation systems and automated system alerting for when data feeds break.

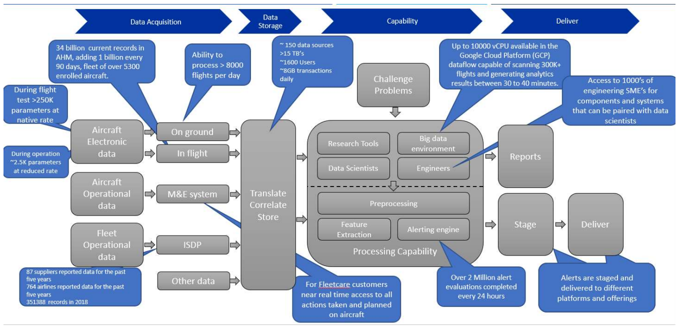

Figure 2: Predictive Maintenance Capability

When leading a predictive maintenance capability and creating of a prognostic research collaboration program from a data science standpoint, we believe it is critical to have a well-defined operator on-boarding process as illustrated in Figure 3, which provides airline customers a high-level understanding of data usage rights, and customer level of effort requirements.

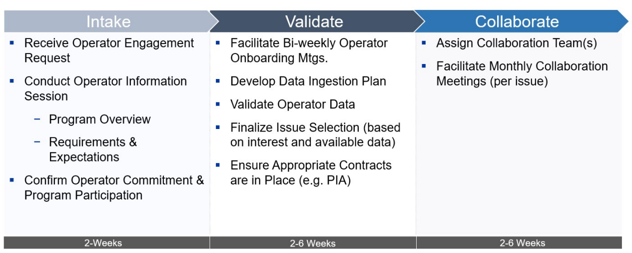

Figure 3: Prognostic Collaboration Onboard Process

In addition to having the right personas and data platform, having a high-level process to share both internally with stakeholders and externally with customers is critical in a few ways. It is important to understand the business requirements of prognostic prior to starting the technical development process. This includes assessing top schedule reliability and maintenance cost drivers experienced by customers. The most important part of the process is that the data science team and engineering teams performing the work can see and understand requirements, expected deliverables including data models, and documentation. It is also important to coordinate with part suppliers from a predictive maintenance leadership perspective. The process shown in Figure 4 outlines the major steps in hypothesizing, developing, generalizing and deploying prognostics.

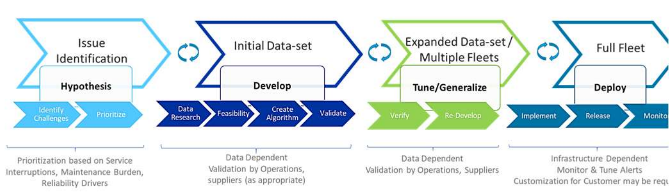

Figure 4: Prognostic Development Process

Beginning with the Hypothesize phase, there are a few key elements and personas required for completion. These include representation from the system or component engineering team, including hardware and software engineers, to help provide schematics, nominal and off nominal operating ranges of the related system or hardware under investigation. The hypotheses should be developed with SMEs input to include an understanding of the system, behavior patterns and failure conditions in-service, and data science methodologies. It is common for multiple hypotheses to be created and refined for investigation. An understanding of the physics of the system through Interface Control Documents, schematics and drawings, system description documents, etc., should be gained from Design Engineers who are experts of the target system or component.

Information surrounding the in-service failure mode(s) and conditions can be obtained from reliability engineers (from airlines, suppliers, and OEMs), including hardware teardown reports, failure descriptions, and maintenance logbooks. At this state, it is also important to establish targeted outcomes for the analysis. It has been our experience over the past 5 years that it is critical to have design, hardware, electrical and software engineers as part of the team when hypothesizing failure modes and operational scenarios to review from the beginning in predictive maintenance issue research.

Examples of success include an air cycle machine prognostic model, which has led to over 75 proactive removals. Inclusion of the design engineers also allowed them to make changes to system software, which improved air cycle machine performance. Finally, before moving to the development phase it is critical to be fully kitted, that means having robust hypothesis to test, enough sensor data to explore, relevant reliability data to extract failure and failure modes, and initial prediction targets.

The Develop phase is where the majority of data science analysis and iteration takes place as data scientists apply data cleansing, labeling, and machine learning techniques to research potential prognostic signatures. In general, there are several approaches to be considered:

Engineering driven approach

- Engineering knowledge helps to down-select sensors and parameters
- Engineering hypotheses, if validated, are interpretable and robust

Engineering to data approach

- Research scientists develop physics-based models with component efficiency metrics
- Calculate component efficiency from sensor data input to predict failures

Data to engineering approach

- Use flight test data to build a model for predicting component internal state parameters

- Use predicted parameters to test engineering hypotheses

Single-stage advanced Machine Learning models

- Engineering input: event identification and parameter selection
- Long Short-Term Memory, Auto-Encoder, Generative Adversarial Network

The above approach is referenced from Boeing Operationalized Aircraft Predictive Maintenance "Lessons Learned from Aircraft Component Failure Prediction using Full Flight Sensor Data" [Yuan (2022)]. As a leader of data scientists and systems engineers, it is important in this phase to balance impact of the prognostic to airline customers, OEMs, and suppliers with how much time and effort is spent on analysis. As an example, the time and resources spent on a significant fleet issue that impacts thousands of aircraft resulting in air turn backs or significant asset downtime will be different from those spent on an issue that impacts a small fleet and has aircraft dispatch relief capability.

A lesson learned as well is to conduct consistent (from a persona standpoint) peer reviews such that the data science team can quickly get feedback on exploratory data results and model validation. This can lead to new hypotheses to test out. It can also lead to the requirement of more data and resources.

From a leadership perspective, it has been observed that there is a point of diminishing returns when additional data or hypothesis generation will not lead to a predictive signature. This was the case when partnering with operators to try and develop a windshield shatter event detection signature. Ultimately, the team was unable to produce an acceptable model for prediction and the analysis was put on hold. It is important even in these cases to ensure model, code, and documentation are saved, as many times the analysis could be re-opened or the data leveraged for a future prognostic.

Another learning is: existing sensors might not be sufficient, and additional sensors can be proposed for future design improvement and prognostic success for the next component update or airplane model/sub-model. In some cases (hopefully many) a prognostic signature can be achieved and then additional operator data sets are tested which takes us to the generalize phase.

In the Generalize phase, the predictive maintenance team has enough confidence in a prognostic model that they test it with other airline data partners to determine if initial model performance is sustained. At this juncture, initial prognostic research targets may have been adjusted as model features were developed and tested. It is also likely at this point the model is running for a limited set of operators and proactive on wing testing might be occurring to further validate the prognostic performance.

Figure 5 illustrates a type of matrix used to determine if precision and recall (as an example) are in range to make a deployment decision. In the commercial aerospace industry, it is important to consider airline operation impacts, time to replace, dispatch relief, component replacement and repair cost as factors when deploying a prognostic.

It has been the experience of the authors to consider recall scores of at least 50% and precision scores of at least 75% as targeted prognostic performance. This does vary depending on several key factors, including the ability to inspect and test the component prior to removal to improve confidence of part degradation.

Another factor is whether the prognostic is supplementing a scheduled based maintenance task, as this could allow for relaxed (lower) performance scores such that risk of a no fault found condition would be during planned maintenance. It should be noted the use of condition-based monitoring is a quickly emerging domain and presents opportunity for further optimized maintenance.

Component replacement costs and shop repair costs are also significant data points to consider when considering prognostic deployment decisions. Aerospace component costs range from hundreds to millions of dollars. An expensive part with no on wing test should require a high performing prognostic such that no fault found costs are minimized.

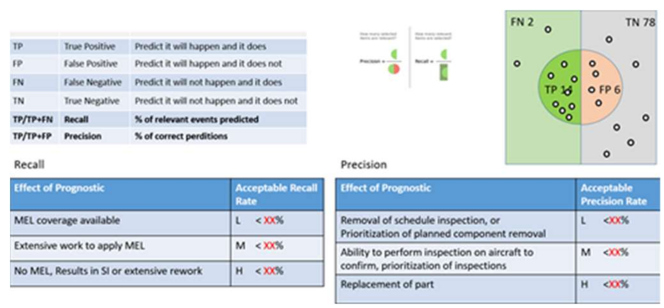

Figure 5: Sample Deployment Decision Matrix

3. Measuring Success

Once prognostic models are deployed, there can be several ways to determine how effective and successful a prognostic is performing in the field. It is also very time consuming from a data collection stand point to correlate prognostic driven actions to actual airplane and component root cause corrective action.

Figure 6 is a real world summary of prognostic performance for Boeing 787 air cycle machine. This prognostic has been in operation since August of 2018. Over a 4-year period, it was deployed to approximately 200 787 aircraft. In that time, it correctly predicted 75 failures leading to proactive inspections and removals as required, but missed 103 failure events.

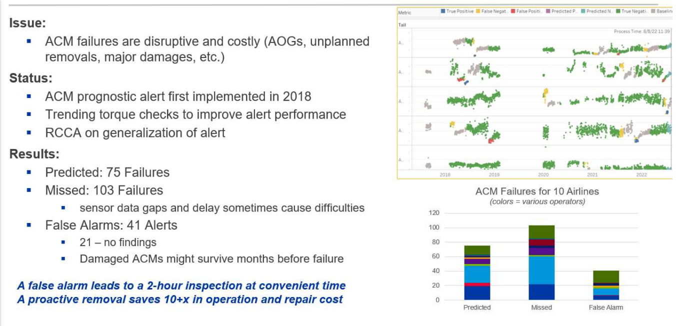

Figure 6: Boeing 787 Air Cycle Machine Prognostic Performance

In the case of the missed events, some are due to sensor data dropouts or delayed data transmission. In addition, the prognostic model requires some component on-wing time to accumulate enough historical data built up to establish per-component baseline behavior to make predictions.

The approach for prognostic feedback required manual data sharing along with person-to-person email contact to build up the action taken and root cause data library for prognostic performance assessments. The air cycle machine prognostic use case has helped to establish feedback inputs into health management solutions such that it is easier for maintenance technical operation personal to quickly input feedback on prognostic actions.

Another form of early validation is to request airline partners to document early release prognostics with pictures of findings. It is also important (where possible) to compare before and after data insight behavior to confirm the prognostic model is working as expected. Figure 7 demonstrates a heat exchanger that was exhibiting clogged behavior, which was confirmed during the planned maintenance (cleaning) event.

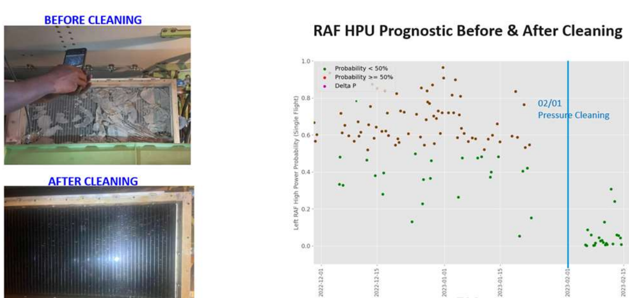

Figure 7: single prognostic alert & maintenance action before and after driven by prognostic

4. Predictive Maintenance Skill & Technology

When building a predictive maintenance team, it is important to identify individuals with a passion for data analytics and a passion for how their data models and insights can improve vehicle, system, and component operation. From a leadership perspective, a cross-functional team composition has been shown to be effective when building an integrated predictive maintenance team. There should be a strong data science function that serves as a center of excellence for data science best practices, training, methodology research, technology evaluation and implementation. This data science function also helps align different data science skills to new predictive maintenance requests, which could be time series data focused or wear out, remaining useful life focused.

It is also ideal that a strong software implementation and deployment team be closely integrated with the data science team. An early learning was that it is very difficult to scale a predictive maintenance team when data scientists are required to sustain machine learning models while trying to research new prognostics. The close integration of data science, software development and development operations (DevOps) is crucial for both resources and technology. If data scientists have code and research in a different platform or cloud environment from the one for deployment, there are non-value-added code translation efforts that may be required for model deployment.

The use of Application Programming Interfaces (APIs) can also be streamlined when research and deploy environments share common code bases. The technology platform for data science research must be closely integrated with the machine learning development operations environments. Well-documented machine learning models utilizing standard code repositories can allow for smooth transitions between data scientists and DevOps personas. This also ensures future updates or platform technology changes can be specified to support existing and new prognostic models.

Another challenge and lesson learned is keeping data scientists and engineers engaged on prognostic research projects that are interesting, while balancing data available for research and analysis tasks. One approach is to allow data scientists to focus on system / component issue prognostic research, development and deployment and try out new tools, technology as part of the research process. There are times when system migration to new platforms is required and as such focus for the team shifts to new processes, training for new application implementation.

Another critical learning is to connect data scientists to design engineers as well as airline technical operation personas such that everyone understands the target performance and steps to make go or no-go decisions on analysis and deployment continuation.

A final skill and technology topic for consideration is focused on data pipeline architecture and data engineering skillsets required for sensor data and reliability data feeds. Predictive maintenance teams that invest in data pipeline monitoring for efficient data dropout troubleshooting can reduce rework as well. Early in our predictive maintenance journey, we would have to do manual large data ingestions when we did not quickly react to data dropouts. This led to the creation of data ingestion and monitoring development operations teams focused on data pipeline set up and sustainment. This also allows the data science team to stay focused on machine learning research for prognostic issues instead of reacting to data coverage or data delay issues.

It is also important to ensure airline information technologists (IT) partners are aware of data outages as quickly as possible as part of the root cause correct action process supporting the data pipeline. It is also important that airplane avionics engineers are part of the data ingestion team in case the issue is airplane related. Our predictive maintenance team has also developed unique data monitoring functions that correlate Automatic Dependent Surveillance Broadcast (ADS-B) aircraft data with flight sensor data files to ensure that a sensor data file is received when an airplane flies.

5. Next Phase of Predictive Maintenance

Today's approach to health management and predictive maintenance is predominantly engineering based and generally reactive. Items to be considered, investigated and monitored are initiated either as operational problems or as items of risk identified at the beginning of a program or during in service operations. Each time one of these areas of focus is identified, a certain level of research, hypothesis development, discovery and validation process is initiated. This usually starts with the question of what data is currently available. While big data and machine learning techniques are being utilized in the area of prognostics and health management, these are individual efforts and are utilized when an investigation is initiated to either to reverse engineer a behavioral model or to determine parameters of interest that will indicate component health or behavior.

Model based Engineering (MBE), digital system models, digital twins and digital threads have the ability to enhance the arena of health management and predictive maintenance beyond the existing engineering knowledge and logic-based approach.

As we move into a world of MBE, we can create a system-wide and automated approach to health management, This is something that digital twins and threads will allow us to do. When we understand how a component or system is designed to work and compare that to how it is actually working, a value can be applied to the difference, trends determined and health measures eventually assigned. With computing capabilities and advanced data science / artificial intelligence / machine learning techniques applied to multiple data sets, these health measures can consider seemingly unimportant and disparate data types to provide accurate health management status, trusted predictions and actionable intelligence, this is the concept of utilizing the Operational Digital Twin for predictive maintenance and health management.

6. Conclusion

In this paper, the authors have discussed the approach to build a predictive maintenance team, the process for building prognostics, and criteria to make go or no-go decisions. Real world prognostic performance and feedback inputs were discussed and experiences shared. Leading a predictive maintenance team is a very rewarding experience. The opportunity to lead and work with data scientist, engineers to ingest, translate, analyze, deploy prognostics that drive actionable insights for customers is very fulfilling. Facing the ever-growing demand to optimize asset availability, leaders that embrace predictive maintenance philosophies will continue to outpace those that do not.

Acknowledgement

We would like to thank our data team, engineering team, and airline collaborators.

References

Yuan, J. (2022). Boeing Operationalized Aircraft Predictive Maintenance. *AAAI Fall Symposia Series on Artificial Intelligence for Predictive Maintenance*. Nov 19, 2022.

Changzhou Wang received his Ph.D. in Information Technology from George Mason University, Fairfax, Virginia, USA in 2000. He joined Boeing in Sept 2000 and is now an Associate Technical Fellow and data scientist in Boeing Global Services. His research interests include temporal data analytics tool development, adapting advanced machine learning methods and combining large-scale flight sensor data and deep engineering knowledge for prognostic modeling. He is a member of ACM and AGIFORS.

Darren Macer is a Senior Technical Fellow specializing in Predictive Maintenance and Health Management for both commercial and military platforms. In this role he leads the research, development and maintenance of capabilities utilizing engineering knowledge, big data techniques and Model Based Engineering techniques and applying them to operational and maintenance data to understand components, systems or aircraft health. Darren also leads the enterprise effort defining the Product Support digital thread/digital twin, in this role he collaborates across the enterprise in defining the product support digital thread/twin strategy, providing technical oversight and guidance and enabling effective solutions that meet business and customer needs. In this position Darren draws upon a career in the aviation industry that spans aircraft support, operation, design, repair, modification and maintenance.

Mark Mazarek earned B.S. in Aerospace Engineering from Embry Riddle Aeronautical University (Daytona Beach, Florida, USA) and M.S. in Systems Engineering from Missouri Science and Technology (Rolla, Missouri, USA). A senior manager in Boeing Global Services (BGS)

Engineering focused on Commercial Airplane Fleet Performance & Lifecycle data, he leads a team of managers, individual contributors and integrated project teams in the development of digital, data and analytics engineering capabilities that improve customer operations, optimize lifecycle cost of airplanes and platforms, drive operational excellence, and further strengthen product and services quality and safety. He also leads the development team focused on operational digital twin and operational digital thread initiatives, as BGS focuses on scaling digital engineering capabilities. Most recently, he led the development and execution of the commercial predictive maintenance capability within BGS Engineering collaborating with Airplane Health Management, Boeing Commercial Airplanes Customer Support, BCA Design Engineering and airline customers.