Improving Computational Efficiency of Prognostics Algorithms in Resource-Constrained Settings

Katelyn J. Jarvis, Chetan S. Kulkarni, Wendy A. Okolo, Chris Teubert

In engineering and aerospace applications, it is vital to operational success to have insight into the expected performance and health of physical systems. The field of prognostics and health management provides quantitative methods for monitoring, predicting, and managing system health. Prognostics algorithms can be employed to assess the current state of a system, propagate the system throughout time, and predict potential anomalies or failures that may occur. While they can provide accurate prediction results, effective prognostics algorithms can be challenging to use in resource-constrained settings due to computational limitations and high computational latency, leading to obsolete predictions. Thus, computationally efficient and accurate algorithms are necessary for future remaining useful life predictions. In this work, we implement new algorithmic approaches for prediction, quantitatively compare them via a battery degradation use-case, and provide recommendations of potential improvements to a prognostics framework.

One approach to prediction is through sampling, whereby the current state of a physical system is sampled many times and each sample is propagated forward until failure is reached, resulting in a distribution of failure values. To improve the efficiency of this process, we implemented five new algorithmic approaches to prediction, including three distinct sampling methods (standard Monte Carlo, Quasi-Monte Carlo, and Latin Hypercube Sampling), a variable time step algorithm, and a variable sample size algorithm. To compare the algorithms, we employ a variety of metrics designed specifically to analyze both computational efficiency and model accuracy. Our metrics include accuracy to compare the average predicted value to ground truth, mean absolute deviation to illustrate dispersion, specific percentile error to describe accuracy within a user-defined risk tolerance, and code run-time. To quantitatively analyze our results, we employ a use-case of degradation of a Lithium-ion battery. We use an electrochemistry-based model to describe the current health state of the battery, and implement our prediction algorithms to propagate forward in time until end-of-discharge (EOD) is reached.

Notably, through this work it was found that none of our sampling approaches had a significant impact on computational efficiency or model accuracy in predicting EOD of the battery. We find that while the sampling methods are unique, the distributions they generate are similar, ultimately producing final predictions that are nearly identical. In exploring the effect of the time step within the prediction algorithm, we found that prediction accuracy was highly dependent on the time step used, and that implementing a variable time step within a particular prediction may provide an increase in computational efficiency while also maintaining prediction accuracy. Finally, implementing a variable sample size also affected prediction, and our results show that tuning both the magnitude and timing of the sample size adjustment can result in improved computation speed and maintained prediction accuracy. Taken together, our findings highlight the challenge of performing prognostics in resource-constrained settings, and illustrate the potential of developing new prediction algorithms to improve computational efficiency.