



Contribution ID: 49

Type: **not specified**

Remaining Useful Life Prediction: a New Temporal Difference Learning Perspective

Predicting the Remaining Useful Life (RUL) of equipment is critical for enabling proactive maintenance and mitigating unexpected failures. Traditional RUL prediction methods often rely on direct regression from sensor data to failure time, resembling Monte Carlo (MC) approaches in reinforcement learning, which require full run-to-failure trajectories and can exhibit high variance. This paper introduces a novel perspective by formulating RUL prediction as a policy evaluation problem within a Markov Reward Process (MRP). In this framework, a reward of +1 is assigned for each operational time step, and the RUL is the expected sum of future rewards. We propose leveraging Temporal Difference (TD) learning, specifically $TD(\lambda)$ with function approximation, to estimate the RUL. This approach allows for learning from incomplete trajectories (bootstrapping), leading to improved sample efficiency, online updates, and reduced variance compared to MC-based methods. We provide theoretical analysis for the convergence and asymptotic properties of TD methods in the RUL context. Empirical evaluations on the NASA turbofan engine dataset demonstrate that our $TD(\lambda)$ approach outperforms traditional MC methods, particularly when full run-to-failure data is scarce and only partial trajectories are available. The results highlight the benefits of TD learning in terms of accuracy and data efficiency, offering a robust and scalable solution for RUL prediction.

Special/ Invited session

Classification

Both methodology and application

Keywords

Reinforcement Learning, Prognostics, Remaining Life Prediction, TD Learning

Primary author: YAN, Hao (Arizona State University)

Presenter: YAN, Hao (Arizona State University)

Track Classification: QSR-Informs