Can Machine Learning and Explainable Artificial Intelligence Help to Improve an Expert Model for Predicting Thermomechanical Fatigue?

Stefan Glaser, Thomas Seifert, and Daniela Oelke

Offenburg University, Germany
firstname.lastname@hs-offenburg.de

Abstract. Machine learning (ML) models are increasingly used for predictive tasks, yet traditional data-based models relying on expert knowledge remain prevalent. This paper examines the enhancement of an expert model for thermomechanical fatigue (TMF) life prediction of turbine components using ML. Using explainable artificial intelligence (XAI) methods such as Permutation Feature Importance (PFI) and SHAP values, we analyzed the patterns and relationships learned by the ML models. Our findings reveal that ML models can outperform expert models, but integrating domain knowledge remains crucial. The study concludes with a proposal to further refine the expert model using insights gained from ML models, aiming for a synergistic improvement.

Keywords: explainable artificial intelligence (XAI), TMF life prediction

1 Introduction

Predictive models using machine learning (ML) are increasingly applied across various fields, although data-based models are not new in many areas, relying on expert knowledge and statistical methods. ML models have the advantage of uncovering unknown relationships, but they may also inherit data biases, especially with limited training data. This is particularly pertinent in engineering domains like thermomechanical fatigue (TMF) of materials in high-temperature applications, where experimental data is scarce. This paper explores enhancing an expert model with an ML model by training it on data of the turbine-blade nickelbase superalloy MAR-M247. We evaluate the ML model using explainable artificial intelligence (XAI) and exploratory data analysis to understand the patterns and relationships it learns, aiming to improve the expert model effectively, thus combining the strengths and weaknesses of both approaches.

2 Domain background

High-temperature turbine components endure start-up and shut-down cycles that cause progressive material damage, leading to failure after a certain number of cycles. To predict the number of cycles to failure (N_{f}), engineers require a TMF life assessment model that has been fitted to specific materials through various loading condition tests.

Different features based on material tests are used to predict the number of cycles to failure (N_{f}). We distinguish between simple features which are direct measurements (e.g. mechanical/plastic strain ranges $\varepsilon_{\mathrm{me/pl}}$, minimum/maximum stress $\sigma_{\mathrm{min/max}}$) or basic combinations (e.g. stress ratio R, stress range $\Delta\sigma$) that require no domain knowledge and expert features that incorporate domain knowledge or additional information. In this work, a fracture-mechanics based expert model for TMF life prediction is considered which in the following is called the D_{TMF} regressor. The model allows to assess damage contributions from fatigue (D_{fat}) and creep (D_{creep}).

3 Can a ML model replace the expert model?

In total, we trained 10 ML models using 8 different ML algorithms and evaluated them using the logarithmic Root Mean Squared Error (RMSL10E). The D_{TMF} regressor achieves an RMSL10E score of 0.496. All, or when using simple features all but two ML-based approaches, were able to achieve even better results than the D_{TMF} regressor. For all further analyses, we worked with the Random Forest, which was one of the best approaches for both the expert features and the simple features with an RMSL10E value of 0.340 and 0.337 respectively. Additional tests were performed with a Support Vector Regressor, but are not detailed in this abstract due to space limitations.

4 Exploration of the ML models with methods of XAI

We used two methods of XAI to explore the functionality of the ML models: Permutation Feature Imporance (PFI) and SHAP (SHapley Additive exPlanations) values. PMI and SHAP agree that when adding the expert features, D_{TMF} is seen as particularly important for the Random Forest showing that expert knowledge has actually been used to create a feature that has a high information content with regard to the TMF life prediction. This is followed by R, D_{creep}, and σ_{min} (PFI), respectively σ_{min}, D_{creep}, and D_{fat} (SHAP).

When using simple features only, both agree that $\varepsilon_{\mathrm{pl}}$ is most important, followed at a distance by $\varepsilon_{\mathrm{me}}$ and σ_{min}. From the perspective of a domain expert, the high importance of simple features like $\varepsilon_{\mathrm{pl}}$ or $\varepsilon_{\mathrm{me}}$ is understandable as there are even simple life prediction models which have been built by domain experts that rely on those features.

What is more interesting is the fact that both models attribute importance to the feature σ_{min}. This feature is indirectly represented in the expert model by the features $\Delta\sigma$ and R, nevertheless it is even considered important when the D_{TMF} is used as a feature. This raises the question whether it contains some information that is not covered by the other features. However, our ML-based experiments showed that the results do not deteriorate if σ_{min} is not given to the algorithms. Apparently, its information is also contained in other features and the preference for σ_{min} might also be due to a bias in the data.

One of the advantages of expert models can be seen in the fact that the hold time was not considered important by the models, presumably because of the small number of instances with high values for the feature. However, as the relationship between hold time and the number of cycles to failure is known in the domain, an expert model can also take this into account.

To conclude: The principal trends identified by XAI are also predicted by the expert model in this way. Since additional simple features are important even when the expert features are used, there seems to be additional information included in the simple features that are not yet represented in the expert features. In this first analysis, we did not succeed in identifying the relationships which the ML models discovered in a way that they could be directly used to improve the expert model. Therefore, as a next step we intend to the train a boosting model that directly builds on the D_{TMF} regressor hoping that this way we can gain better hints on what to add.