"Multi-Model Methods for Explainable Predictive Maintenance of Aeronautical Systems"

This Thesis addresses the need for Predictive Maintenance (PM) solutions that are not only accurate but also trustworthy. While Machine Learning (ML) has fostered increasingly advanced predictive capabilities, many state-of-the-art models remain opaque, limiting their adoption in safety-critical domains such as Aerospace and Aeronautics. In these contexts, where regulatory oversight and human safety are paramount, transparency must be balanced against accuracy.

To meet this challenge, we put forth a hierarchy of needs for PM applications in high-risk settings, emphasizing interpretability, explainability, and actionability as primary requirements. Within this framework, the Mixture of Experts (MoE) architecture is identified as a starting point for reconciling fidelity with intelligibility.

A cornerstone of the work is the introduction of CoCoAFusE, a novel Bayesian extension of MoEs.

CoCoAFusE explicitly addresses the interplay between expert sub-models, offering analysts richer insights and broader representational capabilities towards complex data-generating mechanisms.

Additional explainability tools are developed for both MoEs and CoCoAFusE, enhancing their applicability in dynamical system modeling and fault detection tasks.

Empirical validation on synthetic data, public benchmarks, and industrial case studies demonstrates that CoCoAFusE can deliver accurate, transparent, and actionable predictions.

Overall, this work aims at shrinking the gap between black-box effectiveness and human insights, contributing to the design of responsible AI systems that are both effective and trustworthy in critical engineering domains.