Data-Driven Maintenance Forecasting for Medical Devices: A Decision Support System for Ziemer

Degree programme: BSc in Industrial Engineering and Management Science
Specialisation: Business Engineering
Thesis advisor: Prof. Dr. Bastian Widenmayer
Expert: Prof. Dr. Niki Roger Zumbrunnen
Industrial partner: Ziemer Ophthalmic Systems AG, Port

To support strategic lifecycle decisions within its placement model, Ziemer Ophthalmic Systems AG requires a data-driven foundation. This thesis develops a decision support system that forecasts device-specific operating costs based on historical maintenance data. The system enables evidence-based recommendations on whether to continue, refurbish, or retire medical devices, thereby enhancing service planning and profitability.

Michaël Manuel Sebastiao

Introduction and Objectives

Ziemer Ophthalmic Systems AG operates a placement business model in which ophthalmic laser systems are installed at hospitals and clinics and billed per procedure. Strategic decisions on whether to continue using a device, refurbish it, or retire it are still made manually, based on fragmented cost insights and limited historical context. This creates uncertainty in service profitability and risk exposure. The objective of this thesis is to develop a data-driven decision support system that uses historical maintenance, repair, and operating cost data to support device-level lifecycle decisions. By forecasting future costs, the system provides structured recommendations that help optimize planning, minimize unprofitable placements, and improve long-term margin control.

Research Design

Due to confidentiality constraints, synthetic data was simulated to represent maintenance events, service costs, and usage behavior. The system integrates five forecasting models (median, linear regression, ridge regression, LightGBM, global median), each predicting monthly repair and personnel costs over 36 months. Model training and testing were performed using k-fold cross-validation, and the best model per device was selected based on Mean Absolute Error (MAE).

Results

Model evaluation shows that LightGBM provided the highest accuracy in 76 of 80 cases (best MAE: CHF 125). Simpler models served as robust fallbacks when historical patterns were weak. The system identifies cost-intensive devices, detects degradation trends, and flags refurbishment candidates based on forecasted OPEX. Each device receives a fully automated PDF report, generated in Python, containing technical master data, financial KPIs, top components, forecast metrics, and a final recommendation. Life-cycle assessments are based on discounted cost projections matched with revenue estimates. Differences between models are reflected in prediction variance, sensitivity to usage, and responsiveness to cost spikes.

Implications and Recommendations

The system provides a scalable foundation for data-driven lifecycle management. Once connected to real data, it enables standardized evaluations of devices, supports early cost-risk identification, and improves planning accuracy. It is recommended to integrate the system gradually into Ziemer's internal review process, starting with pilot devices. Regular model updates and extension to further device types will ensure long-term usability.

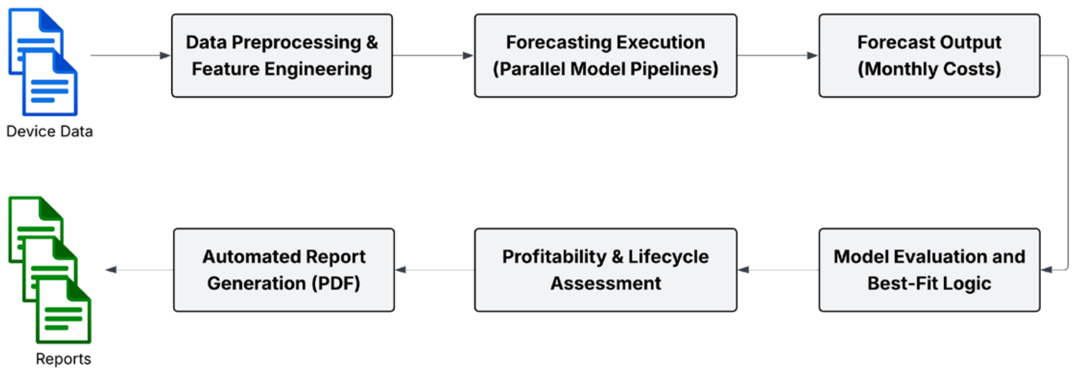

Figure 1: DSS Process