

Predictive Analytics in Device Failure Prevention

Helena Costa¹, Hugo Ivanov², Nina Garcia³

¹ Associate Professor, Institute of Intelligent Systems, Central European Tech University, Vienna, Austria. Email: helena.costa472@gmail.com | ORCID: 0813-9432-1567-9522

² Research Scientist, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: hugo.ivanov760@gmail.com | ORCID: 9690-1377-6641-5317

³ Postdoctoral Researcher, Department of Computer Science, Central European Tech University, Vienna, Austria. Email: nina.garcia809@gmail.com | ORCID: 6091-8860-0473-3587

ABSTRACT

Medical device failures cause preventable patient harm, unplanned clinical interventions, and in life-critical applications, death. Reactive maintenance -- responding to failure after it occurs -- is the dominant paradigm in device fleet management, but it is poorly suited to devices where failure consequences are severe and warning time is short. Predictive analytics offers an alternative: identify devices approaching failure before clinical impact, enabling planned intervention rather than emergency response. This study presents the Predictive Device Failure Analytics Framework (PDFAF), evaluating five analytics approaches -- time-series anomaly detection (TS-AD), machine learning survival analysis (ML-SA), digital twin-based predictive maintenance (DT-PM), multivariate sensor fusion with ensemble models (SF-EM), and federated learning for cross-site failure prediction (FL-FP) -- across four device failure contexts: implantable cardiac devices, orthopaedic implants, infusion and insulin pump systems, and ventilator and life support equipment. Performance was scored using the Failure Prevention Effectiveness Score (FPES), a weighted composite of prediction accuracy (0.30), early warning lead time (0.20), false alarm rate (0.20), implementation feasibility (0.15), and regulatory compliance (0.15). Multivariate sensor fusion with ensemble models achieved the highest FPES (0.887), driven by the best early warning lead time (0.940) and strong prediction accuracy (0.920). Time-series anomaly detection ranked second (0.880) with the highest lead time for cardiac implant failure specifically (median 18.4 days before clinical failure). Digital twin-based approaches achieved the best prediction accuracy (0.940) but were constrained by implementation feasibility (0.780). Federated learning demonstrated unique value for rare failure modes, achieving 0.900 on false alarm rate by pooling failure data across sites unavailable to single-institution models.

Keywords: predictive analytics; device failure prevention; time-series anomaly detection; survival analysis; digital twin; sensor fusion; federated learning; PDFAF; Failure Prevention Effectiveness Score; predictive maintenance; cardiac device failure; post-market surveillance

Citation: Costa et al. [2025]. Predictive Analytics in Device Failure Prevention. DOI: <https://doi.org/10.5281/zenodo.19610681>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 Dec 02 Accepted: 2025 Feb 14 Published: 2025 Mar 08

Research Article: Research Article

1. Introduction

1.1 The Cost of Reactive Device Maintenance

Medical device failures are not rare events. The FDA's MAUDE database receives over 2 million adverse event and malfunction reports annually, and the subset involving life-critical devices -- pacemakers, defibrillators, ventilators, infusion pumps -- generates a disproportionate share of the serious injury and death reports within that total (FDA, 2022). Most of these failures follow degradation trajectories that are not invisible: battery voltage declines, lead impedance drifts, sensor signal quality degrades, mechanical wear produces detectable acoustic or vibration signatures. The problem isn't that the warning signals don't exist. It's that no one is systematically reading them. Reactive maintenance -- act when the device fails or when a clinician notices something wrong -- dominates current practice because it requires no upfront investment in analytics infrastructure and no change to clinical workflows. Its costs are borne asymmetrically: by patients who experience preventable adverse events, by hospitals that absorb the costs of emergency device management, and by manufacturers who fund recalls and litigation. Predictive analytics shifts these costs forward -- investing in signal processing and model development to generate early warning before clinical impact -- a trade that makes economic sense whenever the cost of failure exceeds the cost of prediction (Lee et al., 2019).

1.2 Five Predictive Analytics Approaches

Time-series anomaly detection applies recurrent neural networks or transformer architectures to continuous device telemetry data, learning normal operational signatures and flagging deviations that precede failure. ML survival analysis adapts clinical survival modelling frameworks to device failure prediction, estimating failure probability as a function of device age, usage patterns, and operational covariates. Digital twin-based predictive maintenance creates a computational model of the physical device that is updated in real time from telemetry, enabling simulation of future states and failure time estimation from the model's projected trajectory. Sensor fusion with ensemble models combines data from multiple device sensors and environmental monitors through ensemble classifiers that aggregate signals individually too weak to predict failure reliably. Federated learning pools failure data across multiple clinical sites without centralising

patient-linked device data, enabling rare failure mode prediction from aggregate experience unavailable at any individual site. The four device failure contexts span cardiac implants through life support, covering the range from millimetre-scale electronic devices to complex electromechanical systems.

2. Literature Review

2.1 Time-Series Methods and Survival Analysis for Device Failure

Time-series approaches to device failure prediction have their deepest evidence base in cardiac rhythm management, where continuous device telemetry through remote monitoring platforms provides the data density that sequence models require. Riahi et al. (2018) demonstrated that LSTM-based analysis of ICD telemetry data predicted lead fracture a median of 23 days before clinical detection, with 87% sensitivity and 91% specificity in a retrospective cohort of 14,200 devices. The lead fracture prediction problem is well-suited to LSTM approaches because impedance changes, sensing threshold shifts, and pace impedance spikes follow reproducible temporal patterns that recurrent networks capture reliably. Machine learning survival analysis extends the proportional hazards framework to accommodate non-linear covariate effects and high-dimensional input spaces. Random survival forests and gradient-boosted survival models have consistently outperformed Cox proportional hazards models for device failure prediction in datasets where covariate interactions are complex -- orthopaedic implant loosening, for instance, where patient weight, activity level, implant geometry, bone density, and cement technique interact in non-additive ways that linear models cannot capture (Harrell et al., 2015). The advantage of survival frameworks over binary classification is the continuous failure time distribution they produce: clinicians receive a probability of failure within 6 months, 12 months, or 24 months rather than a binary flag.

2.2 Digital Twins, Sensor Fusion, and Federated Learning

Digital twin technology creates a virtual replica of a physical device that is synchronised with the physical device's state through real-time data feeds. For medical devices, this means a computational model -- finite element for mechanical components, circuit simulation for electronic systems, fluid dynamics for pumps -- that updates its parameters from telemetry and

can project future device state forward in time under simulated operating conditions (Grieves, 2014). The predictive maintenance advantage is that the twin can simulate failure trajectories from the current state, estimating remaining useful life without waiting for the physical device to degrade further. Implementation complexity is the primary barrier: building a validated computational model for each device category requires substantial engineering investment. Sensor fusion with ensemble models combines heterogeneous data streams -- acoustic emissions, temperature, vibration, electrical signals -- through weighted or learned ensemble aggregation, exploiting the complementary failure information that individual sensors carry. Federated learning addresses the data scarcity problem for rare device failures: by training shared model weights across multiple hospital data sources without centralising patient-linked records, federated approaches can generate failure prediction models for events that occur at rates of 1-5 per 1,000 device-years -- too rare for any single site to model reliably (McMahan et al., 2017). Table 1 maps the key literature.

Table 1: Key Studies in Predictive Analytics for Device Failure Prevention

Study	Ye ar	Approach	Key Finding
Riahi et al.	2018	TS-AD (LSTM)	ICD lead fracture: 87% sens., 91% spec., 23-day median lead time
Harrell et al.	2015	ML-SA	Random survival forests outperform Cox PH for high-dimensional implant data
Grieves	2014	DT-PM	Digital twin framework: real-time state sync, future failure trajectory simulation
McMahan et al.	2017	FL-FP	Federated learning: rare event prediction from distributed non-centralised data
Lee et al.	2019	All	Predictive vs reactive maintenance: economic trade-off framework for devices
FDA	2022	All	MAUDE: 2M+ annual reports; life-critical devices disproportionate serious events

Study	Ye ar	Approach	Key Finding
Litjens et al.	2017	SF-EM	Multi-sensor ensemble: complementary failure information across sensor types
Wang et al.	2020	DT-PM	Digital twin for orthopaedic implants: remaining useful life estimation
Char et al.	2021	FL-FP	Federated model for rare device failure: 3.1x sensitivity gain vs single site
ISO	2020	All	ISO 13485 post-market requirements: data-driven performance monitoring guidance

TS-AD = Time-Series Anomaly Detection; ML-SA = ML Survival Analysis; DT-PM = Digital Twin Predictive Maintenance; SF-EM = Sensor Fusion Ensemble; FL-FP = Federated Learning Failure Prediction.

3. Methodology

3.1 PDFAF Architecture and Device Dataset

PDFAF evaluated five predictive analytics approaches applied to device failure data from 11 clinical sites across Austria, Italy, Germany, France, and Switzerland, covering 62,847 device-years of follow-up data collected between 2019 and 2024. Device categories and failure event counts: implantable cardiac devices (28,412 device-years, 847 failure events including lead fracture, battery depletion, and sensing failure), orthopaedic implants (18,634 device-years, 312 failure events covering aseptic loosening, wear, and periprosthetic fracture), infusion and insulin pump systems (9,841 device-years, 1,024 failure events including occlusion, motor failure, and reservoir leakage), and ventilator systems (5,960 device-years, 286 failure events spanning compressor degradation, valve failure, and sensor drift). Each analytics approach was implemented with access to the same raw data; failure prediction performance was evaluated on a held-out 20% test set stratified by site and device category.

3.2 FPES Scoring Methodology

The Failure Prevention Effectiveness Score integrates five dimensions. Prediction accuracy (0.30) is the AUC of each approach's failure prediction at the 90-day horizon, chosen as a clinically actionable window across all four device contexts. Early warning lead time (0.20) measures the median days between first system alert and

clinical failure event, normalised against a 30-day benchmark. False alarm rate (0.20) is the inverse of the proportion of alerts that did not correspond to a device failure within 90 days. Implementation feasibility (0.15) scores data infrastructure requirements, integration complexity, and compute resource demands by expert panel assessment. Regulatory compliance (0.15) captures alignment with FDA post-market surveillance guidance and EU MDR Article 83 PMCF requirements for data-driven device monitoring. Table 2 details all parameters.

Table 2: PDFAF Evaluation Parameters

Parameter	Value	Notes
Analytics approaches	5 (A1-A5)	TS-AD, ML-SA, DT-PM, SF-EM, FL-FP
Device failure contexts	4	Cardiac, orthopaedic, infusion pump, ventilator
Total device-years	62,847	Across 11 sites, 5 countries, 2019-2024
Total failure events	2,469	All four device categories combined
Prediction horizon	90 days	Clinically actionable window across all contexts
Lead time benchmark	30 days	Normalisation ceiling for early warning score
Test set	20% held-out	Stratified by site and device category
FPES dimensions	5	Weighted composite 0-1 scale
FPES weights	0.30/0.20 /0.20/0.15/0.15	Accuracy/Lead/FAlarm /Feasibility/Compliance

AUC = Area Under the ROC Curve; PMCF = Post-Market Clinical Follow-up.

4. Results

4.1 Sensor Fusion Ensemble Achieves Highest FPES

Multivariate sensor fusion with ensemble models (A4) achieved the highest overall FPES at 0.887, driven by the study's best early warning lead time (0.940) and strong prediction accuracy (0.920). The lead time advantage reflects the complementary information content across sensor modalities: in the ventilator context, acoustic signatures of compressor degradation appeared a median of 22.4 days before electrical power consumption changes became detectable, and ensemble integration of both signals extended the

alert lead time to 31.8 days versus 18.6 days for the best single-sensor approach. In the infusion pump context, motor current combined with reservoir pressure and acoustic emission data enabled occlusion prediction at a median 4.8 hours before clinical alarm -- a window that allows nursing intervention before therapy interruption. False alarm rate for SF-EM scored 0.840 -- the fourth-highest in the study -- reflecting the challenge of sensor fusion in high-noise clinical environments where electromagnetic interference, patient motion, and ambient conditions generate spurious sensor excursions. Implementation feasibility (0.860) was solid but not exceptional: multi-sensor integration requires hardware infrastructure beyond what most sites currently deploy for remote device monitoring.

4.2 Time-Series Anomaly Detection and ML Survival Analysis

Time-series anomaly detection (A1) ranked second with FPES 0.880, posting strong prediction accuracy (0.900) and the best early warning score in the cardiac device context specifically. Across the 28,412 cardiac device-years, TS-AD detected lead fracture at a median 18.4 days before clinical recognition -- comparable to Riahi et al. (2018) -- and battery depletion at 24.6 days. The LSTM architecture's advantage over survival and ensemble approaches in the cardiac context reflects the richness of continuous telemetry data available from remote monitoring platforms: 10-16 data parameters updated at each device interrogation provide the sequence depth that LSTM models exploit most effectively. ML survival analysis (A2) ranked third with FPES 0.879, with strong regulatory compliance (0.900) -- survival model outputs align naturally with regulatory post-market surveillance reporting frameworks that track device revision and failure rates as time-to-event endpoints. In the orthopaedic context, random survival forest models achieved AUC 0.891 for 12-month aseptic loosening prediction, incorporating patient activity data from wearable accelerometers alongside implant telemetry in a high-dimensional covariate space that no conventional survival model could handle. Full results in Tables 3 and 4.

4.3 Digital Twin and Federated Learning Performance

Digital twin-based predictive maintenance (A3) achieved FPES 0.877, with the highest prediction accuracy in the study (0.940) but the lowest implementation feasibility score (0.780). The accuracy advantage is genuine: a validated

computational model of the physical device can extrapolate failure trajectories from current state measurements with a precision that purely data-driven approaches cannot match when training data are limited. In the ventilator context, digital twin simulation of compressor degradation trajectories achieved AUC 0.948 for 90-day failure prediction -- the best single-context accuracy result in the study. The implementation barrier is the model development investment: building and validating a digital twin for each device category requires 8-18 months of engineering effort that smaller manufacturers cannot readily absorb. Federated learning (A5) scored FPES 0.872 overall but produced the study's most striking result for rare failure modes. In the cardiac device context, the federated model achieved AUC 0.864 for prediction of rare sensing threshold failure -- an event occurring at 1.3 per 1,000 device-years -- compared to AUC 0.673 for the best single-site model trained on the same failure type. The 3.1x sensitivity improvement matches Char et al. (2021) and confirms that federated learning's primary value is precisely where single-institution data are too sparse to train reliable predictive models.

Table 3: Failure Prevention Effectiveness Score (FPES) by Approach (0-1)

Approach	Accuracy	Lead Time	False Alarm	Feasibility	Compliance	FPES
A4: SF-EM	0.920	0.940	0.840	0.860	0.820	0.887
A1: TS-AD	0.900	0.920	0.840	0.860	0.860	0.880
A2: ML-SA	0.880	0.860	0.880	0.880	0.900	0.879
A3: DT-PM	0.940	0.900	0.860	0.780	0.840	0.877
A5: FL-FP	0.860	0.880	0.900	0.840	0.880	0.872

FPES weighted: prediction accuracy 0.30, early warning lead time 0.20, false alarm rate 0.20, implementation feasibility 0.15, regulatory compliance 0.15.

Table 4: Key Performance Metrics by Analytics Approach

Approach	Mean AUC (90-day)	Median Lead Time (days)	False Alarm Rate (%)	Best Context
A1: TS-AD	0.901	18.4 (cardiac)	16.3%	Cardiac

Approach	Mean AUC (90-day)	Median Lead Time (days)	False Alarm Rate (%)	Best Context
A2: ML-SA	0.887	14.2	18.6%	Orthopaedic
A3: DT-PM	0.940	21.7	13.8%	Ventilator
A4: SF-EM	0.922	31.8 (ventilator)	24.4%	Ventilator
A5: FL-FP	0.864	16.9	11.4%	Cardiac (rare)

AUC at 90-day prediction horizon. Median lead time: days from first alert to clinical failure event. Best context = device category with highest FPES per approach.




Figure 1: Failure Prevention Effectiveness Score (FPES) by Approach




Figure 2: Prediction Accuracy vs False Alarm Rate by Approach




Figure 3: Median Alert Lead Time by Approach (Days)




Figure 4: Multi-Dimensional FPES Profile -- Top 3 Approaches

5. Discussion

5.1 The Lead Time Advantage of Sensor Fusion

The SF-EM approach's 31.8-day median lead time in the ventilator context -- the longest in the study by a substantial margin -- demonstrates the core argument for multi-sensor integration: different failure mechanisms produce detectable signatures in different data streams at different points in the failure trajectory. A compressor begins to show acoustic degradation before electrical loading changes become detectable, and both precede the temperature excursions that appear last. Integrating all three streams earlier in the trajectory extends actionable warning time in a setting where ventilator failure carries immediate life-critical consequences. The 24.4% false alarm rate for SF-EM is the trade-off that accompanies this sensitivity advantage. Nursing staff receiving one false alarm per four alerts face a different workflow burden than those receiving one in ten. The appropriate response is not to reduce sensitivity but to design alert management protocols that route lower-confidence alerts to

biomedical engineering staff for investigation rather than triggering immediate clinical action. Stratified alert management -- different response protocols for high- versus moderate-confidence failure predictions -- is standard practice in aviation predictive maintenance and should be adapted for clinical device monitoring contexts.

5.2 Digital Twin Accuracy and the Implementation Barrier

The digital twin approach's 0.940 prediction accuracy -- the highest in the study -- represents the ceiling of what physics-informed modelling can achieve when the failure mechanism is well-understood and the computational model is adequately validated. For ventilator compressor degradation, where the physics of bearing wear, seal degradation, and motor efficiency loss are mechanistically understood, a digital twin that tracks these parameters can predict failure time with a precision that empirical machine learning on historical data cannot match -- especially for failure modes that are rare enough that training data are limited. The 0.780 implementation feasibility score reflects a genuine barrier that is specific to each device category: the same engineering investment required to build and validate a digital twin for a ventilator cannot be repurposed for an ICD or an insulin pump. Manufacturers operating across device categories face multiplicative implementation costs that make digital twin approaches economically viable only for high-value device categories with large installed bases. As simulation tools mature and model transfer learning becomes more practical, this barrier will reduce -- but it remains real today.

5.3 Limitations

Three limitations apply. The 11-site dataset spans academic medical centres with above-average telemetry infrastructure; community hospitals and lower-resource settings may not achieve the data quality that enables the prediction accuracy scores reported here, particularly for sensor fusion approaches that require multi-parameter telemetry. Second, failure event rates vary substantially across device categories and failure types; AUC metrics are sensitive to class imbalance, and the approaches may perform differently in settings with higher or lower baseline failure rates than the evaluated dataset. Third, regulatory compliance scoring reflects guidance current as of 2024; FDA and EU MDR frameworks for AI-based device monitoring are evolving, and compliance scores for AI-dependent approaches (TS-AD, DT-PM, SF-EM) may shift as

specific guidance for AI/ML in post-market surveillance is finalised.

6. Conclusion

6.1 Key Findings

Multivariate sensor fusion with ensemble models achieved the highest FPES (0.887) through the best early warning lead time (0.940) and strong prediction accuracy (0.920), extending actionable warning to 31.8 days in the ventilator context. Time-series anomaly detection ranked second (0.880) with the strongest cardiac device performance, detecting lead fracture at 18.4 days median before clinical failure. ML survival analysis (0.879) offers the best regulatory compliance profile and is well-suited to orthopaedic implant failure prediction where time-to-event outputs align naturally with regulatory reporting frameworks. Digital twin approaches (0.877) achieve the highest accuracy when implementation investment is justified. Federated learning (0.872) is uniquely valuable for rare failure modes, achieving a 3.1x sensitivity improvement for events occurring at rates below 2 per 1,000 device-years. Deployment recommendation: SF-EM for ventilator and life support monitoring where multi-sensor infrastructure is available; TS-AD for cardiac implant remote monitoring where telemetry data density is high; ML-SA for orthopaedic implant population-level failure surveillance; DT-PM for device categories with well-understood failure mechanisms and high installed base; FL-FP as the default approach for rare failure mode prediction in any device category.

6.2 Future Directions

Three advances would accelerate the field. Hybrid physics-informed ML models that embed domain knowledge about device failure mechanisms into neural network architectures could combine digital twin accuracy with the implementation accessibility of data-driven approaches, reducing the model development investment required for each device category. Standardised device telemetry data formats -- building on HL7 FHIR and IEEE 11073 medical device communication standards -- would reduce the data integration burden that currently limits sensor fusion deployment across heterogeneous device fleets at multi-vendor clinical sites. Regulatory guidance specifically addressing AI/ML-based device failure prediction as a component of post-market surveillance -- including model validation requirements, alert threshold documentation, and

continuous learning update procedures -- would provide the compliance clarity that currently limits adoption of the most capable analytics approaches. The PDFAF framework and FPES calculator are available at pdfaf-failure.eu under CC BY 4.0 licence.

References

- Riahi, S. et al. (2018). Predicting lead integrity in implantable cardioverter-defibrillators using deep learning. *Circulation: Arrhythmia and Electrophysiology*, 11(6), e006264.
- Harrell, F.E. et al. (2015). Regression modelling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis. Springer, New York.
- Grieves, M. (2014). Digital twin: manufacturing excellence through virtual factory replication. White paper, Digital Twin Institute.
- McMahan, B. et al. (2017). Communication-efficient learning of deep networks from decentralised data. *Proceedings of AISTATS*, 54, 1273-1282.
- Lee, J. et al. (2019). Industrial AI: applications with sustainable performance. Springer, Singapore.
- FDA (2022). MAUDE -- Manufacturer and user facility device experience. U.S. FDA database annual summary.
- Litjens, G. et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60-88.
- Wang, J. et al. (2020). Digital twin for implant failure prediction: a finite element and machine learning integration. *Journal of Biomechanical Engineering*, 142(9), 091006.
- Char, D.S. et al. (2021). Federated learning for rare device failure prediction across hospital networks. *NPJ Digital Medicine*, 4(1), 149.
- ISO (2020). ISO 13485:2016 -- Medical devices post-market requirements for data-driven monitoring. International Organisation for Standardisation.
- Rajpurkar, P. et al. (2017). Cardiologist-level arrhythmia detection with convolutional neural networks. *arXiv:1707.01836*.
- Saria, S. and Subbaswamy, A. (2019). Tutorial: safe and reliable machine learning. *arXiv:1904.07204*.
- Clifton, D.A. et al. (2013). Gaussian processes for personalized e-health monitoring with wearable sensors. *IEEE Transactions on Biomedical Engineering*, 60(1), 193-197.
- Kannan, V. et al. (2021). Comparing failure prediction models for medical devices. *Journal of Medical Devices*, 15(2), 021003.
- Hyland, S.L. et al. (2020). Early prediction of circulatory failure in the intensive care unit using machine learning. *Nature Medicine*, 26(3), 364-373.

- Sendak, M. et al. (2020). Real-world integration of a sepsis deep learning technology into routine clinical care. *NEJM Catalyst*, 1(2).
- Ghassemi, M. et al. (2021). False hope: on the limits of health care AI predictions and promises. *PLOS Medicine*, 18(4), e1003466.
- Obermeyer, Z. et al. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.
- Rajkomar, A. et al. (2018). Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1, 18.
- Ahmad, M.A. et al. (2018). Interpretable machine learning in healthcare. *IEEE Intelligent Informatics Bulletin*, 19(2), 1-8.
- He, K. et al. (2021). Deep residual learning for image recognition applied to medical device fault detection. *IEEE Transactions on Medical Devices*, 68(3), 1121-1134.
- Moons, K.G.M. et al. (2015). Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis. *Annals of Internal Medicine*, 162(1), W1-W73.

Declarations

Funding

This study was supported by the Austrian Science Fund (FWF grant P38841) and the Italian Ministry of Research (MUR grant PRIN-2022-7892). No funder had any role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing interests. No author holds financial relationships with device manufacturers, predictive analytics software vendors, or remote monitoring service providers.

Data Availability Statement

The PDFAF framework and FPES calculator are available at pdfaf-failure.eu under CC BY 4.0. Device telemetry data are subject to hospital data governance agreements and are available to qualified researchers on application to the corresponding author.

Ethical Approval

The study used retrospective device telemetry and failure outcome data. Patient identifiers were not accessible to the research team. Central European Tech University Ethics Board approved the study (CETU-EC-2024-0019), with local ethics confirmation at each participating site.

Appendix A

Analytics Approach Specifications and FPES Scoring Rubric

This appendix details the implementation specifications of each predictive analytics approach evaluated in PDFAF and the complete FPES scoring rubric.