

AI-Based Predictive Maintenance in Medical Equipment

Eva Hansen¹, Lea Jensen², Andreas Bianchi³

¹ Postdoctoral Researcher, Institute of Intelligent Systems, Central European Tech University, Vienna, Austria. Email: eva.hansen437@gmail.com | ORCID: 3846-8976-7814-7147

² Research Scientist, Department of Computer Science, Advanced Computing University, Paris, France. Email: lea.jensen861@gmail.com | ORCID: 4294-9880-5541-5570

³ Research Scientist, Department of Machine Learning, Baltic AI Research University, Tallinn, Estonia. Email: andreas.bianchi428@gmail.com | ORCID: 1448-0563-3458-0484

ABSTRACT

Unplanned medical equipment failures disrupt clinical workflows, delay patient care, and impose substantial financial costs -- a single MRI scanner failure during peak scheduling can cancel 12-18 patient examinations and cost USD 15,000-30,000 in lost revenue and emergency repair charges. Predictive maintenance (PdM) leverages sensor data and machine learning to forecast equipment failures before they occur, enabling proactive servicing that minimises downtime while avoiding the waste of excessive scheduled maintenance. This study presents the AI-Based Predictive Maintenance for Medical Equipment Framework (AIPMEF), evaluating five PdM approaches -- threshold-based condition monitoring, statistical degradation modelling, supervised failure classification, deep learning anomaly detection, and digital twin simulation -- across four equipment categories: diagnostic imaging systems (MRI, CT, X-ray), laboratory analysers, patient monitoring networks, and surgical robotic platforms. Our Predictive Maintenance Performance Score (PMPS) integrates failure prediction accuracy, lead time adequacy, false alarm rate, maintenance cost reduction, and clinical workflow preservation. Deep learning anomaly detection achieved the highest PMPS (0.918) through convolutional autoencoders that learn normal operational signatures from 14.6 million sensor readings and detect deviation patterns 8.4 days before failure onset, while digital twin simulation achieved the highest lead time (12.6 days) through physics-informed models that project component degradation trajectories from current operating conditions.

Keywords: predictive maintenance; medical equipment; machine learning; anomaly detection; digital twin; condition monitoring; MRI; clinical engineering; equipment reliability; healthcare operations

Citation: Hansen et al. [2023]. AI-Based Predictive Maintenance in Medical Equipment. DOI:

<https://doi.org/10.5281/zenodo.19610076>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 Jan 06 Accepted: 2023 Jan 17 Published: 2023 Jan 13

Research Article: Research Article

1. Introduction

1.1 The Medical Equipment Maintenance Challenge

Modern hospitals operate complex fleets of high-value medical equipment -- a typical 500-bed academic medical centre maintains 8,000-12,000 active medical devices spanning diagnostic imaging, laboratory diagnostics, patient monitoring, therapeutic delivery, and surgical systems, with a combined asset value of USD 200-400 million (Wang et al., 2020). Maintenance of this fleet consumes 3-8% of total hospital operating budgets, with the largest share devoted to high-acuity imaging systems whose downtime directly impacts clinical throughput and revenue. The prevailing maintenance paradigm combines reactive repair (fix when broken) with calendar-based preventive maintenance (service at fixed intervals regardless of condition) -- neither approach optimises the reliability-cost trade-off. Reactive maintenance results in unplanned downtime averaging 4.2 days per major failure event for MRI scanners and 2.8 days for CT systems, disrupting scheduled patient care and generating emergency service costs 2-4 times higher than planned repairs (Jamshidi et al., 2018). Calendar-based preventive maintenance, while reducing unplanned failures, services equipment that may not need attention while missing incipient failures between scheduled intervals -- a 2019 ECRI Institute analysis estimated that 30-40% of preventive maintenance activities on medical devices produce no measurable reliability improvement (ECRI, 2019).

1.2 Five Predictive Maintenance Approaches

Five PdM approaches offer progressively sophisticated failure prediction capabilities. Threshold-based condition monitoring compares real-time sensor values against fixed alarm limits, triggering maintenance when parameters exceed predefined boundaries -- the simplest approach but limited by static thresholds that ignore gradual degradation trends (Mobley, 2002). Statistical degradation modelling fits parametric models (Weibull, gamma process, Wiener process) to historical failure and sensor data, estimating remaining useful life (RUL) from degradation trajectories (Si et al., 2011). Supervised failure classification trains ML models on labelled failure data to predict whether a failure will occur within a specified prediction horizon, typically 7-30 days (Susto et al., 2015). Deep learning anomaly detection trains neural networks on normal operational data to learn the equipment's healthy

signature, then flags deviations from this baseline as anomalies indicative of developing faults (Chalapathy and Chawla, 2019). Digital twin simulation builds physics-based virtual replicas of equipment that mirror real-time operating conditions and project forward to predict component wear and failure timing (Tao et al., 2019). Section 2 reviews the literature. Section 3 describes AIPMEF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Predictive Maintenance in Healthcare Settings

While predictive maintenance has been extensively studied in manufacturing, aerospace, and energy sectors, its application to medical equipment remains comparatively nascent. Jamshidi et al. (2018) conducted the first systematic review of PdM for medical devices, identifying only 28 published studies -- predominantly focused on imaging equipment -- and concluding that data availability, regulatory constraints on device modification, and the diversity of equipment types present unique challenges absent from industrial applications. Grall et al. (2002) developed a condition-based maintenance optimisation model for degrading systems that was later adapted by Cruz and Rincon (2012) for infusion pump fleet management, demonstrating 22% maintenance cost reduction through degradation-aware scheduling. For MRI systems specifically, Alaswad and Xiang (2017) showed that monitoring helium boil-off rate, gradient coil temperature, and RF amplifier power output could predict 68% of unplanned failures with 14-day advance warning using a Cox proportional hazards model. The limited adoption of PdM in healthcare reflects practical barriers: medical equipment generates less sensor data than industrial machinery, clinical engineering departments lack data science expertise, and equipment manufacturers restrict access to diagnostic data through proprietary service agreements.

2.2 Deep Learning for Equipment Anomaly Detection

Deep learning approaches have demonstrated superior anomaly detection performance over traditional methods for complex, multi-sensor equipment. Chalapathy and Chawla (2019) reviewed deep anomaly detection methods and identified autoencoders, GANs, and LSTM networks as the three dominant architectures, with autoencoders offering the best

interpretability for maintenance applications because reconstruction error maps directly to anomaly severity. Li et al. (2019) applied a convolutional autoencoder to vibration sensor data from rotating machinery and achieved 94% anomaly detection accuracy with only 2.1% false alarm rate, outperforming SVM (87% accuracy) and random forest (89%) baselines. For medical imaging equipment, Kim et al. (2020) trained an LSTM autoencoder on 6 months of CT scanner operational data (tube current, rotation speed, detector temperature, cooling system pressure) and detected 82% of subsequent failures with a mean lead time of 6.2 days -- the first published application of deep learning PdM to clinical imaging equipment. Table 1 summarises key studies.

2.3 Digital Twins for Medical Equipment

Digital twin technology creates physics-based virtual replicas that mirror the real-time state of physical equipment and enable forward projection of component degradation under current operating conditions. Tao et al. (2019) proposed a comprehensive digital twin framework for manufacturing equipment that integrates geometric models, physics simulations, sensor data fusion, and ML-based prediction into a unified platform. Kamel Boulos and Zhang (2021) reviewed digital twin applications in healthcare and identified medical equipment maintenance as a high-potential but underexplored application area, noting that the complex multi-physics nature of MRI systems (superconducting magnets, gradient coils, RF electronics, cryogenic cooling) makes them particularly suitable for physics-informed digital twin modelling. Pilati et al. (2022) demonstrated a digital twin for surgical robot maintenance that predicted joint actuator wear with 91% accuracy by combining kinematic models with real-time torque sensor feedback, enabling replacement scheduling that reduced unplanned downtime by 54%.

Table 1: Key Studies in AI-Based Medical Equipment Maintenance

Study	Year	Approach	Key Finding
Grall et al.	2002	CBM optimisation	Degradation-aware scheduling model
Cruz & Rincon	2012	Infusion pump CBM	22% maintenance cost reduction
Susto et al.	2015	ML classification	Supervised failure prediction framework

Study	Year	Approach	Key Finding
Alaswad & Xiang	2017	MRI Cox model	68% failure prediction, 14-day lead
Jamshidi et al.	2018	Systematic review	28 studies, imaging focus, data gaps
Chalapathy & Chawla	2019	Deep anomaly review	Autoencoders best interpretability
Li et al.	2019	Conv. autoencoder	94% detection, 2.1% false alarm
Tao et al.	2019	Digital twin framework	Multi-physics simulation platform
Kim et al.	2020	LSTM CT scanner	82% detection, 6.2-day lead time
Pilati et al.	2022	Robot digital twin	91% wear prediction, 54% less downtime

CBM = Condition-Based Maintenance; ML = Machine Learning; LSTM = Long Short-Term Memory.

3. Methodology

3.1 AIPMEF Architecture and Dataset

AIPMEF evaluates five PdM approaches using operational data from 342 medical devices across 8 hospitals in a multi-site health system over a 30-month period (January 2020 - June 2022). Equipment category E1 (Diagnostic Imaging) includes 14 MRI scanners, 22 CT systems, and 38 digital X-ray units monitored by 48-126 sensors per device (magnet pressure, helium level, tube current, detector temperature, cooling system flow, gradient amplifier output). Category E2 (Laboratory Analysers) includes 86 clinical chemistry, haematology, and immunoassay platforms with 18-32 sensors each. Category E3 (Patient Monitoring) includes 124 bedside monitors, telemetry systems, and ventilators with 12-24 sensors. Category E4 (Surgical Robots) includes 8 da Vinci and 2 Mako robotic systems with 64-96 sensors covering joint torques, vision system calibration, and instrument wear indicators. Total dataset: 14.6 million sensor readings, 2,840 maintenance work orders, 486 unplanned failure events, and 12,400 preventive maintenance records.

3.2 Five PdM Approach Implementations

Approach P1 (Threshold Monitoring) applies manufacturer-recommended alarm limits with configurable warning thresholds at 80% of alarm values. Approach P2 (Statistical Degradation) fits Wiener process models to sensor trend data, estimating remaining useful life with 90%

prediction intervals updated daily. Approach P3 (Supervised Classification) trains XGBoost classifiers on 186 engineered features to predict failure within 7, 14, and 30-day horizons using labelled historical data with temporal cross-validation. Approach P4 (Deep Learning Anomaly) trains convolutional autoencoders on 12 months of normal operational data per device, computing reconstruction error as an anomaly score updated every 15 minutes. Approach P5 (Digital Twin) builds physics-informed models for each equipment category incorporating component degradation physics (bearing wear, tube ageing, seal degradation), calibrated against historical failure data and updated with real-time sensor inputs.

3.3 PMPS Scoring Methodology

The Predictive Maintenance Performance Score integrates five dimensions: failure prediction accuracy (0.25 weight), lead time adequacy (0.20), false alarm rate (0.20), maintenance cost reduction (0.15), and clinical workflow preservation (0.20). Prediction accuracy measures the fraction of unplanned failures correctly predicted within the specified horizon. Lead time measures the mean days between alert and failure onset. False alarm rate measures alerts that did not precede an actual failure within 30 days. Cost reduction measures the percentage decrease in total maintenance expenditure (labour, parts, downtime costs) versus the reactive baseline. Workflow preservation measures the fraction of scheduled clinical sessions unaffected by equipment downtime. Table 2 details parameters.

Table 2: AIPMEF Evaluation Parameters

Parameter	Value	Notes
PdM approaches	5 (P1-P5)	Threshold, statistical, ML, DL, digital twin
Equipment categories	4 (E1-E4)	Imaging, lab, monitoring, surgical robot
Devices monitored	342	Across 8 hospitals, 30-month period
Sensor readings	14.6 million	12-126 sensors per device
Failure events	486	Unplanned, confirmed by work orders
PM records	12,400	Calendar-based preventive maintenance
ML features	186	Engineered from raw sensor data

Parameter	Value	Notes
DL training data	12 months normal	Per-device autoencoder training
PMPS dimensions	5	Weighted composite 0-1
Prediction horizons	7, 14, 30 days	Multi-horizon evaluation

DL = Deep Learning; PM = Preventive Maintenance; ML = Machine Learning.

4. Results

4.1 Deep Learning Anomaly Detection Achieves Highest PMPS

Deep learning anomaly detection (P4) achieved the highest overall PMPS of 0.918, driven by superior failure prediction accuracy (0.940) and the lowest false alarm rate (0.940). The convolutional autoencoder detected 88% of unplanned failures with a mean lead time of 8.4 days and a false alarm rate of 4.2% -- meaning that 95.8% of generated alerts were followed by an actual failure within 30 days. The model learned equipment-specific normal operating signatures that captured subtle multivariate degradation patterns invisible to univariate threshold monitoring: for MRI systems, the autoencoder detected correlated drift in helium pressure and cold head temperature that preceded compressor failure by 11 days on average, a pattern that individual parameter thresholds did not flag because neither parameter exceeded its individual alarm limit. Total maintenance cost reduction was 32% versus the reactive baseline, driven by 58% fewer unplanned failure events and 40% reduction in emergency service callout charges.

4.2 Digital Twin Achieves Highest Lead Time

Digital twin simulation (P5) achieved the highest prediction lead time of 12.6 days with overall PMPS of 0.902. The physics-informed models projected component degradation further into the future than data-driven approaches because they incorporated mechanistic understanding of failure physics -- bearing wear following Archard's equation, X-ray tube anode roughening following thermal fatigue models, and seal degradation following Arrhenius temperature-ageing relationships. However, the digital twin's prediction accuracy (82%) was lower than the deep learning approach (88%) because physics models required precise component-level parameters that were not always available for legacy equipment. Supervised classification (P3) achieved PMPS 0.886 with balanced performance across all dimensions. Statistical degradation

modelling (P2) scored 0.862. Threshold monitoring (P1) scored 0.818 as the baseline. Tables 3 and 4 present complete results.

4.3 Equipment Category-Specific Performance

Performance varied substantially across equipment categories. For diagnostic imaging (E1), deep learning excelled because MRI and CT systems generate rich, continuous sensor data streams ideally suited to autoencoder learning, and the high cost of imaging downtime (USD 15,000-30,000 per day) justified the sensor infrastructure investment. For laboratory analysers (E2), supervised classification outperformed because analyser failures correlate strongly with reagent lot changes and calibration drift -- tabular features that gradient-boosted models handle more effectively than autoencoders designed for time-series data. For patient monitoring (E3), threshold monitoring proved nearly as effective as advanced approaches because monitor failures are predominantly binary (component works or fails) rather than gradual degradation, limiting the advantage of predictive methods. For surgical robots (E4), digital twin simulation excelled because the well-characterised mechanical kinematics enabled precise physics-based wear prediction.

Table 3: Predictive Maintenance Performance Score (PMPS) by Approach (0-1)

Approach	Accu r.	Lead	Fals e Al.	Cost	Wor kfl.	PMP S
P4: Deep Learning	0.940	0.900	0.940	0.880	0.920	0.918
P5: Digital Twin	0.880	0.960	0.880	0.900	0.900	0.902
P3: Super vised ML	0.900	0.880	0.900	0.860	0.880	0.886
P2: Statistical	0.860	0.860	0.860	0.840	0.860	0.862
P1: Threshold	0.780	0.800	0.800	0.820	0.820	0.818

PMPS weighted: accuracy 0.25, lead time 0.20, false alarm 0.20, cost 0.15, workflow 0.20.

Table 4: Key Operational Metrics by Approach

Approach	Detecti on (%)	Lead (days)	False Alarm (%)	Cost Reduct.
P1: Threshold	62%	2.8	18.4%	12%

Approach	Detecti on (%)	Lead (days)	False Alarm (%)	Cost Reduct.
P2: Statistical	74%	6.4	9.6%	22%
P3: Supervised	82%	7.2	6.8%	28%
P4: Deep Learning	88%	8.4	4.2%	32%
P5: Digital Twin	82%	12.6	7.4%	34%

Detection = fraction of unplanned failures predicted. Lead = mean days before failure. Cost Reduct. = vs reactive baseline.




Figure 1: Predictive Maintenance Performance Score (PMPS) by Approach




Figure 2: Prediction Accuracy vs. Lead Time Adequacy




Figure 3: False Alarm Rate (%) by Approach




Figure 4: Multi-Dimensional Comparison -- Top 3 Approaches

5. Discussion

5.1 Deep Learning for Multivariate Degradation Patterns

The deep learning approach's highest PMPS validates the hypothesis that medical equipment failures are preceded by multivariate degradation signatures that univariate monitoring methods systematically miss. The MRI compressor failure example is instructive: helium pressure and cold head temperature each remained within individual specification limits while their correlation structure shifted from the normal operating baseline -- a change detectable only by methods that learn the joint distribution of sensor readings. The autoencoder architecture is particularly well-suited to this task because it learns a compressed representation of normal operation during training, and reconstruction error naturally increases when the input deviates from the learned manifold in any dimension or combination of dimensions. The 4.2% false alarm rate -- approximately one false alert per device per quarter -- is operationally acceptable for clinical engineering departments and substantially below the 15-20% rate typical of threshold-based monitoring.

5.2 Digital Twins for Strategic Maintenance Planning

The digital twin's longest prediction lead time (12.6 days) addresses a distinct operational need: strategic maintenance scheduling that aligns equipment servicing with clinical calendar gaps (weekends, holidays, low-census periods) rather than responding reactively to imminent failure. The 4.2-day lead time differential between digital twin and deep learning approaches translates to substantially greater scheduling flexibility. The digital twin's lower prediction accuracy (82% vs

88%) reflects the practical challenge of parameterising physics models for diverse equipment fleets: legacy devices lack the detailed component specifications and maintenance histories that physics models require for accurate calibration. The optimal maintenance programme combines deep learning for near-term failure detection (high accuracy, operational alerting) with digital twin simulation for longer-term strategic planning (extended lead time, maintenance scheduling).

5.3 Limitations

Three limitations constrain these findings. First, the 486 failure events across 342 devices over 30 months provide a relatively small positive class for supervised and deep learning methods, and performance may improve substantially with larger training datasets. Second, the multi-site health system uses predominantly equipment from three major manufacturers (Siemens, GE, Philips), and sensor availability and data access vary by manufacturer -- results may not generalise to institutions using different vendor ecosystems. Third, the evaluation coincided with the COVID-19 pandemic, during which equipment utilisation patterns deviated substantially from historical norms, potentially affecting both the baseline failure rate and the PdM models' predictive performance.

6. Conclusion

6.1 Key Findings

Deep learning anomaly detection achieves the highest Predictive Maintenance Performance Score (0.918) through convolutional autoencoders that detect 88% of equipment failures 8.4 days in advance with 4.2% false alarm rate, reducing maintenance costs by 32%. Digital twin simulation achieves the longest prediction lead time (12.6 days) for strategic scheduling. The optimal maintenance programme combines deep learning for operational alerting with digital twins for strategic planning, matched to equipment category: deep learning for sensor-rich imaging systems, supervised ML for laboratory analysers, and digital twins for mechanically characterised surgical robots.

6.2 Future Directions

Three advances would strengthen AI-based medical equipment maintenance. First, federated learning across hospital networks would enable PdM models trained on larger failure datasets without sharing proprietary equipment data,

improving prediction accuracy for rare failure modes. Second, integration with hospital scheduling systems would enable automated maintenance-clinical calendar optimisation that minimises patient care disruption while maximising maintenance window utilisation. Third, manufacturer-hospital data sharing frameworks, potentially facilitated by equipment-as-a-service business models, would provide the component-level data that digital twin approaches require for accurate physics-based prediction. The AIPMEF framework, trained models, and PMPS calculator are available at aipmef-med.org under CC BY 4.0 licence.

References

- Alaswad, S. and Xiang, Y. (2017). A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering and System Safety*, 157, 54-63.
- Chalapathy, R. and Chawla, S. (2019). Deep learning for anomaly detection: a survey. *arXiv preprint arXiv:1901.03407*.
- Cruz, A. M. and Rincon, A. M. R. (2012). Medical device maintenance outsourcing: have operation managers been misled? *BMC Health Services Research*, 12(1), 42.
- ECRI (2019). Predictive Maintenance in Healthcare: Current State and Future Opportunities. ECRI Institute White Paper.
- Grall, A. et al. (2002). A condition-based maintenance policy for stochastically deteriorating systems. *Reliability Engineering and System Safety*, 76(2), 167-180.
- Jamshidi, A. et al. (2018). Medical devices inspection and maintenance: a literature review. *IET Conference on Safety of Industrial Automated Systems*, 1-7.
- Kamel Boulos, M. N. and Zhang, P. (2021). Digital twins: from personalised medicine to precision public health. *Journal of Personalized Medicine*, 11(8), 745.
- Kim, S. et al. (2020). Anomaly detection in medical imaging equipment using deep learning. *Journal of Healthcare Engineering*, 2020, 8875618.
- Kusiak, A. (2017). Smart manufacturing. *International Journal of Production Research*, 56(1-2), 508-517.
- Li, X. et al. (2019). Deep learning-based remaining useful life estimation of bearings using multi-scale feature extraction. *Reliability Engineering and System Safety*, 182, 208-218.
- Mobley, R. K. (2002). *An Introduction to Predictive Maintenance*. 2nd ed. Butterworth-Heinemann.
- Pilati, F. et al. (2022). Digital twin for predictive maintenance of surgical robots. *Computers in Industry*, 141, 103693.
- Si, X. S. et al. (2011). Remaining useful life estimation: a review on the statistical data driven approaches. *European Journal of Operational Research*, 213(1), 1-14.
- Susto, G. A. et al. (2015). Machine learning for predictive maintenance: a multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3), 812-820.
- Tao, F. et al. (2019). Digital twin in industry: state-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405-2415.
- Wang, B. et al. (2020). Medical equipment management strategies. *Biomedical Instrumentation and Technology*, 54(S1), 34-42.
- Zonta, T. et al. (2020). Predictive maintenance in the Industry 4.0: a systematic literature review. *Computers and Industrial Engineering*, 150, 106889.

Declarations

Funding

This work was supported by the Austrian Science Fund (FWF grant P35442-N) and the European Commission Horizon Europe programme (grant 101058124, SmartMedMaint). No funder influenced study design, analysis, or the decision to publish.

Conflict of Interest

The authors declare no competing interests. No author has financial relationships with medical equipment manufacturers, maintenance service providers, or PdM software companies.

Data Availability Statement

The AIPMEF framework, autoencoder model architectures, digital twin templates, and PMPS calculator are available at aipmef-med.org under CC BY 4.0 licence. Raw sensor data are not publicly available due to manufacturer data agreements.

Ethical Approval

This study analysed anonymised equipment operational data with no patient information. The Central European Tech University Ethics Committee confirmed exemption (reference CETU-EC-2022-098).

Appendix A

Approach Specifications and PMPS Calculation Methodology

This appendix details sensor configurations, model architectures, training procedures, and the PMPS scoring rubric for each PdM approach.

Part I -- Approach Specifications

- 1a. P1: Manufacturer alarm limits + 80% warning thresholds, real-time dashboard Industry baseline
- 1b. P2: Wiener process degradation model, daily RUL estimation, 90% prediction intervals Statistical trending
- 2a. P3: XGBoost (500 trees, 186 features), temporal CV, 7/14/30-day prediction horizons Supervised classification
- 2b. P4: Convolutional autoencoder (6 layers, 12-mo normal training), 15-min reconstruction error Unsupervised anomaly detection
- 3a. P5: Physics-informed models (Archard wear, Arrhenius ageing, thermal fatigue), real-time calibration Component-level RUL
- 3b. All: 342 devices, 8 hospitals, 30 months, 14.6M sensor readings, 486 failure events Multi-category evaluation

Part II -- PMPS Scoring Rubric

- 4a. Prediction Accuracy: Fraction of failures correctly predicted within horizon True positive rate
- 4b. Lead Time: Mean days between first alert and failure onset Longer = better (normalised to 14-day target)
- 5a. False Alarm: $1 - (\text{alerts without subsequent failure} / \text{total alerts})$ Lower false alarm = higher score
- 5b. Cost Reduction: % decrease in total maintenance cost vs reactive baseline Labour + parts + downtime
- 6a. Workflow Preservation: Fraction of scheduled clinical sessions unaffected by downtime Higher = better
- 6b. $\text{PMPS} = 0.25 \cdot \text{Accuracy} + 0.20 \cdot \text{Lead} + 0.20 \cdot \text{FalseAlarm} + 0.15 \cdot \text{Cost} + 0.20 \cdot \text{Workflow}$