International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V7I2P122
Eureka Vision Publication | Volume 7, Issue 2, 160-166, 2026

Original Article

Sensor Fusion Architectures for High-Accuracy Failure Investigation in MedTech

Suchitra Venkatesan
Independent Researcher California, USA.

Received On: 10/03/2026 Revised On: 09/04/2026 Accepted On: 17/04/2026 Published On: 24/04/2026

Abstract - Medical device failures can have life-threatening consequences, yet current failure investigation protocols often rely on single-sensor or manual inspection methods that miss subtle, multi-domain fault signatures. A systematic review of sensor fusion architectures including centralized, decentralized, hierarchical, and hybrid configurations is presented as applied to high-accuracy failure investigation in medical technology (MedTech). Established data-level, feature-level, and decision-level fusion paradigms are examined, their suitability for regulatory environments governed by ISO 13485, IEC 60601, and FDA 21 CFR Part 11 is assessed, and a reference architecture integrating heterogeneous sensing modalities (vibration, acoustic emission, thermal, electrical impedance, and optical) with machine-learning-based anomaly detection is proposed. Key challenges including sensor calibration drift, real-time latency, data provenance, and explainability are discussed. A forward-looking roadmap covering edge-AI deployment, digital-twin integration, and federated learning strategies appropriate for MedTech quality systems is provided.

Keywords - Sensor Fusion, Failure Analysis, Medtech, Fault Detection, Anomaly Detection, ISO 13485, Digital Twin, Machine Learning, Quality Systems, Root Cause Analysis.

1. Introduction

Medical device reliability is a critical patient-safety imperative. The U.S. Food and Drug Administration (FDA) receives tens of thousands of Medical Device Reports (MDRs) annually, and a significant fraction trace back to systematic hardware or software failures that escaped detection during routine quality control [1]. Traditional failure investigation methodologies visual inspection, single-axis vibration analysis, or benchtop electrical testing frequently lack the observability required to isolate multi-domain fault mechanisms characteristic of complex electromechanical assemblies such as infusion pumps, imaging systems, robotic surgical platforms, and implantable cardiac devices.

Sensor fusion, defined as the principled combination of data from multiple sensing modalities to achieve inferences that are superior to those from any individual sensor, offers a path toward more complete fault observability. The aerospace and automotive industries have deployed fusion-based health monitoring for decades [2], but the translation to regulated MedTech environments is complicated by stringent data-integrity requirements, the heterogeneity of failure modes, and the need for auditable, explainable conclusions that can satisfy regulatory scrutiny.

This paper makes the following contributions: (i) a structured taxonomy of sensor fusion architectures relevant to MedTech failure investigation; (ii) a mapping of those architectures to common regulatory requirements; (iii) a reference design illustrating an end-to-end fusion pipeline; and (iv) identification of open research challenges and an applied roadmap.

The remainder of this paper is organized as follows. Section II reviews the landscape of failure modes encountered in MedTech. Section III surveys sensing modalities. Section IV describes fusion architectures and their trade-offs. Section V addresses algorithmic and AI approaches. Section VI covers regulatory considerations. Section VII presents a reference architecture. Section VIII addresses deployment challenges. Section IX outlines future directions, and Section X concludes.

2. Failure Mode Taxonomy in Medtech

Understanding the failure space is a prerequisite to selecting appropriate sensing and fusion strategies. Medical device failures cluster into three primary domains.

2.1. Mechanical Failures

Mechanical failures encompass fatigue crack propagation, wear-induced degradation, seal breaches, and connector fretting. These manifest across many time scales from sudden overload fractures to multi-year fatigue accumulation and may be accompanied by characteristic vibration signatures, acoustic emissions at crack initiation events, or changes in surface topology detectable by optical profilometry [3].

2.2. Electrical and Electronic Failures

Electrical failures include insulation breakdown, capacitor electrolyte dry-out, solder joint cracking under thermal cycling, and electromagnetic interference (EMI)-induced transients. Impedance spectroscopy, partial-discharge detection, and current-signature analysis are established sensing methods for these phenomena. Circuit-board assemblies in life-support equipment may cycle through tens of thousands of thermal excursions, making thermally induced solder fatigue a prevalent failure mode [4].

2.3. Software and Firmware Failures

Although not directly addressable by hardware sensors, software failures often manifest in observable physical symptoms: runaway actuators produce anomalous force or velocity profiles; watchdog-timer resets create characteristic power-cycle signatures; and memory corruption may alter duty-cycle patterns in power converters. Correlating software-event logs with physical sensor streams is therefore a key use case for multi-modal fusion.

2.4. Biological Interface Failures

Mechanical impedance of implantable sensors; biofouling degrades optical transducer throughput; and protein adsorption shifts electrochemical electrode potentials. These modes demand specialized electrochemical and optical sensing channels not found in conventional industrial health-monitoring systems.

TABLE I summarizes the primary failure mode categories, associated sensing modalities, and typical detection latencies.

Table 1. Failure Mode – Sensing Modality Mapping

Failure Category	Dominant Sensing Modality	Typical Detection Latency
Mechanical fatigue	Vibration, AE, strain gauge	Hours – weeks
Wear / fretting	AE, surface optical, tribometry	Days – months
Solder / connector fatigue	Thermal IR, impedance, AE	Weeks – years
Insulation degradation	Partial discharge, impedance	Minutes – months
Software/firmware anomaly	Power signature, actuator motion	Milliseconds – seconds
Biofouling / capsule formation	Optical transmittance, EIS	Days – weeks

3. Sensing Modalities for Failure Investigation

A well-designed fusion system begins with the thoughtful selection of complementary sensing modalities whose information content is partially overlapping providing redundancy and partially orthogonal providing breadth of fault observability.

3.1. Vibration and Inertial Sensing

Micro-electromechanical systems (MEMS) accelerometers and piezoelectric accelerometers capture structural vibration signatures in the 0.1 Hz to 20 kHz range. Power spectral density (PSD) analysis, envelope detection, and time-frequency representations such as the short-time Fourier transform (STFT) or wavelet decomposition reveal bearing defect frequencies, gear mesh harmonics, and imbalance signatures. In MedTech applications, vibration sensing has been applied to centrifuge motors in haematology analysers, compressors in anaesthesia workstations, and drive mechanisms in robotic surgical systems [5].

3.2. Acoustic Emission

Acoustic emission (AE) sensors detect stress waves in the 100 kHz–1 MHz range generated by crack propagation, plastic deformation, and phase transitions. AE is particularly sensitive to early-stage fatigue cracking and is used in structural-health monitoring of pressure vessels and catheter hubs. The high bandwidth of AE data demands careful signal conditioning and feature extraction prior to fusion [6].

3.3. Thermal and Infrared Sensing

Infrared cameras and point thermistors reveal hot spots indicative of increased electrical resistance, coolant flow loss, or excessive friction. Thermal drift in reference-circuit components can be tracked over operational lifetime to infer remaining useful life (RUL). In power converters aboard ventilators and external defibrillators, junction temperatures of switching transistors are strongly correlated with long-term reliability [7].

3.4. Electrical Impedance Spectroscopy

Electrical impedance spectroscopy (EIS) applies a swept sinusoidal excitation and measures the complex impedance as a function of frequency. It can characterise electrode-electrolyte interfaces, detect moisture ingress in hermetic enclosures, and quantify dielectric aging in polymer insulators. EIS is well-established for battery state-of-health monitoring and is increasingly applied to implantable device characterisation [8].

3.5. Optical and Vision-Based Sensing

Machine-vision systems with structured-light or confocal optical profilometers capture surface morphology at sub-micron resolution. Optical coherence tomography (OCT) provides depth-resolved imaging of layered structures such as catheter walls or multilayer circuit-board laminates. Hyperspectral cameras can discriminate material composition changes indicative of corrosion or contamination.

3.6. Pressure, Force, and Flow

Pressure transducers and differential flow meters are essential for pneumatic and hydraulic subsystems in devices such as ventilators, infusion pumps, and dialysis machines. Deviations from nominal flow-pressure relationships are diagnostic of pump degradation, tube occlusion, or valve seat erosion.

4. Sensor Fusion Architectures

Sensor fusion architectures are conventionally classified by the processing level at which information is combined: data-level (early), feature-level (intermediate), or decision-level (late). Superimposed on this taxonomy is the topological arrangement of fusion nodes: centralised, decentralised, hierarchical, or hybrid.

4.1. Data-Level (Early) Fusion

In data-level fusion, raw or minimally pre-processed signals from multiple sensors are concatenated into a joint representation before feature extraction and inference. This approach preserves the maximum mutual information across modalities and is theoretically optimal when the joint likelihood model is tractable. However, it requires synchronised, co-registered sensor streams and imposes high bandwidth and storage demands. Calibration drift in any constituent sensor can corrupt the entire fused representation. In MedTech failure investigation, data-level fusion is most appropriate when sensors share a common physical substrate for example, a multi-electrode impedance array on a circuit board or when time-series alignment is guaranteed by hardware triggering.

4.2. Feature-Level (Intermediate) Fusion

Feature-level fusion extracts descriptive statistics or latent representations from each sensor channel independently and then concatenates or projects them into a shared feature space for joint classification or regression. This is the most commonly deployed paradigm in industrial condition monitoring because it: (i) reduces bandwidth requirements by transmitting feature vectors rather than raw waveforms; (ii) accommodates heterogeneous sampling rates and measurement units; and (iii) allows per-modality pre-processing pipelines to be validated independently before integration. Common feature types include spectral band energies, wavelet packet node energies, statistical moments (mean, variance, kurtosis, skewness), and autoregressive model coefficients [9].

4.3. Decision-Level (Late) Fusion

Decision-level fusion combines the output classification labels or probability distributions of modality-specific classifiers. Methods include majority voting, Bayesian combination of posterior probabilities, Dempster-Shafer evidence theory, and fuzzy logic aggregation. Late fusion is the most architecturally modular approach: individual sensing channels can be added, replaced, or taken offline without retraining the other channel classifiers. This property is valuable in MedTech contexts where a single device variant may ship with or without optional sensing hardware. The principal disadvantage is the loss of cross-modal correlations that are discarded when reducing each channel to a scalar label or probability vector.

4.4. Hierarchical Fusion

Hierarchical architectures combine the above paradigms across multiple tiers. A representative structure might perform: (i) data-level fusion of co-located vibration and acoustic emission channels at a local processing node; (ii) feature-level fusion of vibration, thermal, and impedance feature vectors at a device-level aggregator; and (iii) decision-level fusion across multiple device instances at a fleet-level monitoring server. This mirrors the physical hierarchy of complex MedTech systems module, assembly, system, and deployed fleet and distributes computational load proportionally [10].

4.5. Hybrid and Adaptive Architectures

Adaptive fusion architectures dynamically select or weight fusion levels based on estimated sensor reliability, communication bandwidth availability, or computational budget. Quality-of-information (QoI) metrics derived from self-test routines, cross-sensor consistency checks, or uncertainty estimates from Bayesian models gate the contribution of each sensor channel. Adaptive weighting is especially important in MedTech failure investigation because a device undergoing the investigation may itself have sensor channels that are partially degraded, introducing the challenge of fault-in-the-sensor during fault-in-the-device scenarios.

Table 2. Fusion Architecture Comparative Summary

Architecture	Information Preserved	Bandwidth Demand	Modularity	Typical MedTech Use Case
Data-level	Maximum	Very High	Low	Co-located multi-electrode arrays
Feature-level	High	Moderate	Moderate	Multi-modal device health dashboards
Decision-level	Moderate	Low	High	Fleet-level anomaly dashboards
Hierarchical	Adaptive	Distributed	High	Complex capital equipment (imaging, robotics)
Hybrid/Adaptive	Context-dependent	Variable	Very High	Intelligent portable diagnostic devices

5. Algorithmic Approaches and AI Integration

The choice of fusion architecture determines the algorithmic options available at each fusion node. This section surveys the principal algorithmic families applicable to MedTech failure investigation.

5.1. Classical Statistical Methods

Principal component analysis (PCA) and its kernel extensions reduce high-dimensional feature vectors to lower-dimensional subspaces that capture maximum variance, facilitating anomaly detection as Hotelling's T^2 or Q-statistic exceedances. Linear discriminant analysis (LDA) learns class-discriminant projections when labelled failure data are available. These methods offer interpretability and low computational overhead, making them suitable for resource-constrained embedded systems and for initial hypothesis generation during a structured failure investigation [11].

5.2. Machine Learning Classifiers

Support vector machines (SVMs) with radial basis function kernels perform well in moderate-dimensional, limited-sample scenarios typical of MedTech failure datasets where labelled failure events are inherently rare. Random forests provide ensemble diversity, implicit feature importance ranking, and robustness to missing modalities a valuable property when sensor channels are intermittently unavailable. Gradient-boosted trees have demonstrated state-of-the-art accuracy on tabular feature vectors in predictive maintenance benchmarks [12].

5.3. Deep Learning Approaches

Convolutional neural networks (CNNs) applied to time-frequency spectrograms learn translation-invariant fault-pattern detectors without manual feature engineering. Long short-term memory (LSTM) networks and their transformer-based successors model temporal dependencies across long observation windows. Multimodal deep fusion networks jointly learn from heterogeneous input streams by encoding each modality through modality-specific branches before late or intermediate concatenation and joint inference. These approaches require substantial labelled datasets that are typically available in large device fleets but are challenging to assemble for rare failure modes; transfer learning from related domains can partially address this limitation [13].

5.4. Bayesian and Probabilistic Approaches

Bayesian filtering frameworks Kalman filters, unscented Kalman filters, and particle filters provide principled state estimation under measurement noise and model uncertainty, producing calibrated uncertainty bounds on inferred fault state. Gaussian process (GP) regression enables non-parametric RUL prediction with confidence intervals, which are essential for regulatory communication of remaining useful life estimates. Bayesian approaches naturally accommodate sensor dropout by marginalising over missing observations.

5.5. Physics-Informed and Hybrid Models

Purely data-driven models can fail silently when operating outside their training distribution. Physics-informed neural networks (PINNs) embed known governing equations constitutive material laws, thermal diffusion models, or equivalent-circuit representations into the network loss function, constraining the learned model to physically plausible regions. In MedTech, where device physics are often well characterised, such hybrid models can substantially improve extrapolation accuracy and provide mechanistically interpretable failure narratives required for regulatory root cause analysis [14].

6. Regulatory and Quality System Considerations

The deployment of sensor fusion systems within MedTech quality processes must satisfy a layered regulatory framework spanning international standards and national regulations.

6.1. ISO 13485 – Quality Management Systems

ISO 13485:2016 mandates documented processes for the control of monitoring and measuring equipment, including calibration records, measurement uncertainty assessments, and procedures for handling out-of-specification readings. A sensor fusion system constitutes a monitoring and measuring system in ISO 13485 terms; each constituent sensor must be individually calibrated, and the system-level measurement uncertainty must be characterised and documented. Change control procedures apply when sensor channels are added or algorithms are updated [15].

6.2. IEC 60601-1 and Collateral Standards

IEC 60601-1 defines basic safety and essential performance requirements for medical electrical equipment. Sensor fusion systems embedded within or connected to such equipment must not introduce new hazards or impair essential performance. Risk management under ISO 14971 is required to evaluate failure modes of the fusion system itself—including the consequences of false-positive failure alerts (nuisance alarms) and false-negative misses (undetected failures).

6.3. FDA 21 CFR Part 11 and Electronic Records

Where sensor fusion systems generate or store records that are used in regulatory submissions or failure investigation reports, FDA 21 CFR Part 11 requires audit trails, access controls, and electronic signatures. Raw sensor data, derived features, and algorithmic inference outputs must be stored with timestamps, operator identifiers, and version provenance to support inspection-ready failure investigation documentation.

6.4. EU MDR 2017/745 and Post-Market Surveillance

The EU Medical Device Regulation requires manufacturers to implement post-market surveillance (PMS) systems capable of detecting safety signals from field data. Sensor fusion architectures deployed across device fleets can serve as automated PMS data sources, continuously monitoring for population-level anomaly trends that may indicate systemic design or process failures before individual MDR thresholds are reached.

6.5. Algorithm Transparency and Explainability

Regulatory agencies increasingly expect that AI/ML-based diagnostic tools provide explanations for their outputs, particularly for decisions that may influence patient safety. Techniques such as SHAP (SHapley Additive exPlanations) values, attention maps, and saliency analysis can attribute fusion model outputs to contributing sensor channels and time windows, generating the mechanistic narrative required for credible root cause reports [16].

7. Proposed Reference Architecture

Based on the foregoing analysis, a five-layer reference architecture for high-accuracy failure investigation in MedTech is proposed, as illustrated schematically below.

7.1. Layer 1 – Sensing and Signal Conditioning

Heterogeneous sensors are grouped into co-located clusters based on physical proximity and causal relevance. Each cluster includes an analog front-end (AFE) providing gain, anti-aliasing filtering, and analog-to-digital conversion at an appropriate resolution (16–24 bits) and sampling rate. Hardware triggering synchronises acquisition across channels within a cluster to sub-microsecond alignment. Sensor health is continuously monitored via self-test excitation and cross-channel consistency checks; channels failing health checks are flagged and their quality-of-information weights are reduced accordingly.

7.2. Layer 2 – Local Feature Extraction

Per-cluster microcontrollers or digital signal processors (DSPs) execute modality-specific pre-processing pipelines. Vibration channels compute PSD, envelope spectrum, and cepstrum. AE channels apply threshold-based hit detection and extract hit descriptors (peak amplitude, rise time, duration, energy, counts). Thermal channels compute spatial gradient maps and rate-of-change estimates. Impedance channels fit equivalent circuit parameters via nonlinear least squares. Output: compact, time-stamped feature vectors per cluster, transmitted over a secure internal bus to Layer 3.

7.3. Layer 3 – Intermediate Fusion and Anomaly Scoring

A device-level embedded processor performs feature-level fusion. An autoencoder or one-class SVM trained on a population of healthy device feature vectors computes a per-observation anomaly score. Multivariate control charts (Hotelling T^2 and Q-statistic) provide threshold-based alarm generation with configurable false-alarm rate targets. A locally deployable gradient-boosted classifier maps fused feature vectors to a taxonomy of known failure modes, outputting class posteriors and SHAP-based feature attributions.

7.4. Layer 4 – Fleet-Level Aggregation and Root Cause Analytics

Encrypted, de-identified anomaly scores and feature vectors are streamed to a cloud or on-premises server aggregating data across the deployed fleet. Bayesian hierarchical models estimate population-level degradation trajectories, enabling early detection of systemic failure modes not apparent from individual device observations. Digital twin simulations, parameterised by fleet-wide statistics, provide what-if analysis to distinguish design-space failures from manufacturing process excursions. Automated report generation produces structured failure investigation narratives aligned with FDA and ISO 13485 documentation requirements.

7.5. Layer 5 – Human Review and Regulatory Documentation

All algorithmic outputs are presented to qualified engineers and quality personnel via a structured review dashboard. Explainability outputs (SHAP waterfall charts, time-series attention overlays) accompany each alarm to support human judgement. Confirmed root causes are recorded in the complaint-handling and CAPA (corrective and preventive action) system, closing the quality loop. Audit trails capture every access, review, and override event in a 21 CFR Part 11-compliant record store.

8. Deployment Challenges

8.1. Calibration and Sensor Drift

Long-term sensor drift is a fundamental threat to fusion system accuracy. Systematic drift shifts individual sensor channels, corrupting features derived from absolute magnitudes. Strategies include: (i) periodic in-situ automated calibration using embedded reference stimuli; (ii) differential measurement between redundant sensors; (iii) adaptive normalisation that tracks sensor baselines using robust statistics; and (iv) drift-aware retraining that updates model parameters as sensor characteristics evolve within defined limits.

8.2. Real-Time Latency Requirements

Failure investigations on in-service devices may require near-real-time anomaly detection to prevent patient harm. Latency budgets from sensor event to actionable alarm vary by failure mode: mechanical overloads may demand sub-second response while fatigue-crack propagation monitoring may tolerate multi-minute aggregation cycles. Architecture design must partition processing across edge, device, and cloud tiers to meet the most stringent latency requirements with the available computational budget.

8.3. Data Scarcity and Class Imbalance

Failure events in well-designed medical devices are rare by design. Supervised learning algorithms trained on heavily imbalanced datasets hundreds of thousands of nominal observations and tens of failure examples tend to be biased toward the majority class. Mitigation strategies include: one-class and anomaly-based learning formulations; synthetic minority oversampling (SMOTE) and generative adversarial network (GAN)-based augmentation; and transfer learning from accelerated-life-test datasets to field deployment contexts [17].

8.4. Cybersecurity and Data Integrity

Sensor fusion systems connected to cloud infrastructure represent an expanded attack surface for medical device cybersecurity. Adversarial manipulation of sensor streams could conceal genuine failures or generate spurious alerts. Defences include: cryptographic authentication of sensor data at the point of acquisition; anomaly detection on the sensor data stream itself to identify injection attacks; and air-gapped investigation workstations for post-incident forensic analysis. FDA cybersecurity guidance and NIST frameworks apply to networked medical device components.

8.5. Interoperability and Standardisation

The absence of standardised data interchange formats for sensor fusion in MedTech creates integration barriers when devices from multiple manufacturers must be monitored within a single health system. Emerging standards such as IEEE 11073 (Personal Health Devices), HL7 FHIR device resource profiles, and the Medical Device Communication Standard (ISO/IEEE 11073-10206) provide candidate frameworks for interoperable health data exchange, though specific bindings for sensor fusion metadata remain under development.

9. Future Directions

9.1. Edge AI and Neuromorphic Processors

The emergence of dedicated neural processing units (NPUs) and neuromorphic chips offers the prospect of executing complex multimodal inference at the device edge with milliwatt power budgets enabling always-on fusion without cloud connectivity. Spiking neural networks trained on AE and vibration event streams show promise for ultra-low-power anomaly detection in implantable and wearable device contexts [18].

9.2. Digital Twin Integration

Physics-based digital twins high-fidelity simulations of device behaviour parameterised by design and process data can serve as virtual sensors, predicting unmeasured states and filling sensing gaps. When combined with real sensor fusion systems in a hybrid observer framework, digital twins improve anomaly localisation accuracy and enable prospective failure mode analysis that identifies the most information-rich sensing configurations for future designs.

9.3. Federated Learning for Privacy-Preserving Fleet Analytics

Federated learning enables model training across distributed device fleets without transmitting raw patient or device data to a central server. Each device trains a local model update on its private data; only model parameter gradients are aggregated centrally. This approach addresses the dual challenge of data scarcity (by aggregating across fleets) and privacy (by retaining raw data locally), and is compatible with HIPAA and GDPR constraints on health-related data flows [19].

9.4. Multimodal Large Language Model Integration

Large language models (LLMs) augmented with tool use and multimodal perception capabilities represent a nascent but promising direction for failure investigation: an LLM orchestrator could autonomously direct data collection, interpret fusion model outputs, formulate hypotheses, retrieve relevant failure history from structured databases, and draft investigation reports in natural language. Rigorous validation and human-oversight requirements will be essential before such systems are deployed in regulated MedTech quality processes.

9.5. Standardised Fusion Benchmarks

The MedTech community lacks publicly available, labelled failure datasets that would enable objective comparison of fusion architectures. Collaborative data-sharing consortia analogous to ImageNet for computer vision or PhysioNet for physiological signals—would accelerate research and reduce redundant engineering effort across the industry. Regulatory agencies could catalyse such initiatives by recognising dataset contributions within the quality management framework.

10. Conclusion

This white paper has presented a comprehensive survey of sensor fusion architectures for high-accuracy failure investigation in MedTech. It has been shown that no single fusion paradigm is universally optimal: data-level fusion maximises information preservation but demands tight synchronisation and high bandwidth; feature-level fusion balances information richness with modularity; decision-level fusion prioritises flexibility and resilience; and hierarchical hybrid architectures can achieve all of these properties in proportion to system complexity.

The integration of machine learning from classical statistical process control through deep multimodal networks to Bayesian probabilistic models substantially enhances fault detectability and failure mode discrimination. However, deployment in regulated MedTech environments imposes non-negotiable requirements on data integrity, calibration traceability, algorithm explainability, and cybersecurity that must be treated as first-class design constraints rather than post-hoc additions.

The proposed five-layer reference architecture provides a concrete design template that accommodates these constraints while remaining adaptable to diverse device categories, sensing configurations, and regulatory jurisdictions. Open challenges in data scarcity, interoperability, and real-time latency motivate continued research, and emerging technologies edge AI, digital twins, and federated learning offer credible pathways to resolution.

As medical devices grow more complex and post-market surveillance requirements intensify, sensor fusion is expected to transition from an investigative tool applied reactively after failures occur to an always-on patient safety system that detects precursor signatures before clinical harm is realised. Realising this vision requires sustained collaboration among device manufacturers, clinical engineers, regulatory scientists, and the AI research community.

Acknowledgements

Contributions from colleagues in the biomedical reliability engineering and quality systems communities whose discussions shaped the perspectives presented in this work are gratefully acknowledged. No specific experimental datasets are reported in this paper; all claims are drawn from the published literature cited herein.

References

[1] U.S. Food and Drug Administration, "MAUDE – Manufacturer and User Facility Device Experience Database," FDA, Silver Spring, MD, USA. [Online]. Available: https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfmaude/search.cfm

[2] R. Luppold, K. Kimmich, and R. Lupa, "A survey on sensor data fusion in aerospace and automotive structural health monitoring," IEEE Sens. J., vol. 21, no. 4, pp. 5423–5436, 2021.

[3] M. Elforjani and S. Sherif, "Acoustic emission-based condition monitoring and damage growth tracking of slow speed slewing ring bearings," IEEE Trans. Ind. Electron., vol. 63, no. 10, pp. 6401–6410, 2016.
[4] E. Suhir, "Predictive modeling of the dynamic response and the reliability of a solder joint assembly in a PCB system," J. Phys. D: Appl. Phys., vol. 50, no. 43, p. 433001, 2017.
[5] A. K. S. Jardine, D. Lin, and D. Banjevic, "A review on machinery diagnostics and prognostics implementing condition-based maintenance," Mech. Syst. Signal Process., vol. 20, no. 7, pp. 1483–1510, 2006.
[6] C. Grosse and M. Ohtsu, Eds., Acoustic Emission Testing: Basics for Research – Applications in Engineering. Berlin, Germany: Springer, 2008.
[7] H. Wang, K. Ma, and F. Blaabjerg, "Evaluation of temperature cycling profiles for electrolytic capacitors in motor drive applications," in Proc. IEEE APEC, 2012, pp. 2048–2055.
[8] M. Itagaki, S. Suzuki, I. Shitanda, and K. Watanabe, "Electrochemical impedance and complex capacitance to interpret electrochemical capacitor," Electrochemistry, vol. 75, no. 8, pp. 649–655, 2007.
[9] L. Rokach, "Ensemble-based classifiers," Artif. Intell. Rev., vol. 33, no. 1–2, pp. 1–39, 2010.
[10] D. Hall and J. Llinas, "An introduction to multisensor data fusion," Proc. IEEE, vol. 85, no. 1, pp. 6–23, 1997.
[11] J. E. Jackson and G. S. Mudholkar, "Control procedures for residuals associated with principal component analysis," Technometrics, vol. 21, no. 3, pp. 341–349, 1979.
[12] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. ACM SIGKDD, 2016, pp. 785–794.
[13] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015.
[14] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," J. Comput. Phys., vol. 378, pp. 686–707, 2019.
[15] International Organization for Standardization, ISO 13485:2016 – Medical Devices – Quality Management Systems – Requirements for Regulatory Purposes. Geneva, Switzerland: ISO, 2016.
[16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems (NeurIPS), 2017, pp. 4765–4774.
[17] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," J. Artif. Intell. Res., vol. 16, pp. 321–357, 2002.
[18] W. Maass, "Networks of spiking neurons: The third generation of neural network models," Neural Netw., vol. 10, no. 9, pp. 1659–1671, 1997.
[19] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017, pp. 1273–1282.