Explainable AI-Driven Predictive Maintenance Curriculum for Smart Manufacturing

Dr. Abdullah Al Mamun, Old Dominion University

Dr. Abdullah Al Mamun is an Assistant Professor in the Department of Engineering Technology at Old Dominion University. He earned his Ph.D. in Industrial and Systems Engineering from Mississippi State University in August 2023, following the completion of his M.S. in Manufacturing Engineering at Minnesota State University in December 2015. Dr. Mamun's research focuses on integrating advanced sensor technologies into smart manufacturing systems to enhance the cyber-physical resiliency of additive manufacturing, enable real-time machinery prognostics and anomaly detection, and optimize energy consumption for intelligent and sustainable production environments.

Dr. Murat Kuzlu, Old Dominion University

Murat Kuzlu (Senior Member, IEEE) is an Associate Professor in the Department of Engineering Technology at Old Dominion University in Norfolk, VA, USA. He earned his B.Sc., M.Sc., and Ph.D. degrees in Electronics and Communication Engineering in 2001, 2004, and 2010, respectively. Before joining ODU, he served as a Senior Researcher at the Scientific and Technological Research Council of Turkey (TUBITAK) and as a Research Assistant Professor at the Advanced Research Institute of Virginia Tech. His research interests include smart systems, trustworthy AI, explainable AI, agentic AI, and next-generation networks.

Dr. Katherine Smith, Old Dominion University

Katherine Smith received Ph.D. in Modeling and Simulation, M.S. in Applied and Computational Mathematics, and B.S. degrees in applied mathematics and mechanical engineering from Old Dominion University. Dr. Smith is an Assistant Professor in the Engineering Technology Department at Old Dominion University. Her research interests include data science, machine learning, and artificial intelligence, systems modeling and network analysis with applications in supply chain and cyber physical systems, data-driven implementation of industry 4.0 technologies, and serious game development.

Dr. Vukica M. Jovanovic, Old Dominion University

Dr. Vukica Jovanovic is the Chair of the Engineering Technology Department and a Full Professor of Engineering Technology. She holds a Ph.D. in Mechanical Engineering Technology from Purdue University, with a focus on Digital Manufacturing. In addition, she earned a Magister degree (equivalent to Ph.D. candidacy in the United States) in Industrial Engineering and Engineering Management, specializing in Production Systems Design, as well as a Dipl.-Ing. (Master's) degree in Industrial Engineering and Engineering Management with a concentration in Mechatronics, Robotics, and Automation. Dr. Jovanovic brings extensive industry experience to academia. She worked as an Industrial Engineer in the manufacturing of aircraft wings and nacelles for commercial aircraft models ERJ-145 and ERJ-147 (Embraer) at Gamesa Aeronáutica in Vitoria/Bilbao, Spain. Prior to her university studies, she was a certified master electrician, completing apprenticeships in the power industry and an aluminum mill. Her professional foundation was further shaped through hands-on work in her family-owned manufacturing company, which specialized in sheet metal forming machines and metalworking services. Before pursuing an academic career, she served on the faculty at the University of Novi Sad in Serbia.

Dr. Dalya Ismael, Old Dominion University

Dr. Dalya Ismael holds a Doctor of Philosophy degree in Civil & Environmental Engineering from Virginia Tech. She is currently an Assistant Professor of Civil Engineering Technology (CET) in the Department of Engineering Technology at Old Dominion University, where her research focuses on encouraging sustainable design and construction practices through behavioral interventions and visualization techniques. Dr. Ismael has more than nine years of teaching and industry experience, and is a LEED Green Associate and an Envision Specialist.

Adel El-Shahat, Old Dominion University
Angela Sicaja, Old Dominion University
Md Hedayetul Islam Chy, Old Dominion University

Explainable AI-Driven Predictive Maintenance Curriculum for Smart Manufacturing

Abstract

Nowadays, Industry 4.0 has transformed manufacturing industries into a data-rich system driven by IoT, automation, and Artificial Intelligence (AI). Within this context, Predictive Maintenance (PdM) provides a proactive strategy that leverages heterogeneous sensor data such as vibration, acoustic, electrical, and visual signals along with historical performance and advanced analytics to forecast equipment failures before they occur. Usually, AI-driven PdM (AI-PdM) enhances this capability by integrating AI-based sensor analytics to automate fault prediction and optimize system reliability. However, traditional AI-PdM often functions as a "black box," providing limited interpretability of its decision-making process and posing challenges for trust, validation, and human oversight in critical manufacturing environments. To address these challenges, explainable AI-based PdM (XAI-PdM) extends beyond prediction accuracy by leveraging explainability through transparent-based analysis into intelligent maintenance decision-making. XAI-PdM-based framework integrates heterogeneous sensor data with interpretable AI models that not only predict equipment degradation but also clarify which sensors and features contribute to each prediction. Such transparency bridges the gaps between complex AI algorithms and engineering expertise, fostering trustworthy AI, traceability, and actionable insights for real-time maintenance. Identifying this need, it is essential to transfer XAI-PdM concepts to future educators, empowering teachers to prepare students for the demands of smart manufacturing environments. This paper provides condense guidelines for utilizing cutting-edge XAI tools with real industrial datasets, enabling students to enhance their problem-solving skills and workforce readiness for next-generation advanced manufacturing. The XAI-PdM-driven context integrates both theoretical and experiential learning to equip students with cross-disciplinary competencies in AI and data analytics. Ultimately, these guidelines provide to prepare students who can design, interpret, and deploy data-driven intelligent PdM systems that foster sustainability and innovations in the manufacturing industry.

Keywords: Predictive Maintenance (PdM), Explainable AI, Smart Manufacturing, Heterogeneous Sensor, Engineering Education.

1 Introduction

1.1 Background

Smart manufacturing under Industry 4.0 is rapidly transitioning factory environments into highly connected, data-rich systems through the integration of the Industrial Internet of Things (IIoT), cyber-physical systems, advanced automation, and artificial intelligence (AI) [1,2]. Manufacturing assets such as motors, pumps, CNC machines, conveyors, gearboxes, and robotic systems produce continuous streams of operational data through diverse sensors and control systems to measure vibration, temperature, current, sound, and many other signals [1,3]. These data streams enable a shift from traditional maintenance strategies, which fix machines after failure, to data-driven strategies or predictive maintenance (PdM). PdM supports condition-based interventions instead of fixed-interval service, which can improve reliability, reduce unplanned downtime, and optimize maintenance resources. PdM represents this shift by using historical and real-time monitoring data to estimate machine health. It also forecasts failures before they occur [3]. This type of maintenance typically relies on multimodal measurements (vibration, acoustic, current/voltage, temperature, pressure, speed, and other signals) combined with operational context (load, duty cycle) to identify degradation patterns. On the other hand, AI-driven PdM enhances performance by leveraging machine learning (ML) and deep learning (DL) models that can automatically learn complex patterns from large datasets, assist in detecting early-stage anomalies, classify fault types, and estimate remaining useful life. As a result, AI-based PdM has become more widely used in smart manufacturing systems to seek higher productivity, reduce maintenance costs, and extend asset life sustainably.

In practice, many AI-PdM models are difficult to interpret. Such as, high-performing DL models often act as "black-box" predictors. Typically, the output shows only the fault label or risk score without clearly identifying which signals, features or operating conditions drove the decision [4]. In real manufacturing environments, this limited transparency creates barriers to maintenance engineers and managers. Most of the time engineers hesitate to trust the model outputs. Also, managers require justification before scheduling costly interventions and safety-critical systems often demand traceable reasoning [4,5]. Moreover, due to not having interpretability, it becomes harder to validate models, diagnose errors and detect spurious correlations. To resolve this, explainable AI (XAI) can have a huge impact. XAI addresses these barriers by providing tools and methods that make AI predictions transparent, actionable and interpretable. In an XAI-enabled PdM framework (XAI-PdM), predictions are contained in explanations. For example, feature attribution, sensor contribution ranking, time/frequency region importance. These explanations help to connect AI to physical reasoning. In addition to improving user trust, XAI can support model debugging, data quality checks, root-cause reasoning and human-decision-making. Therefore, XAI-PdM is not only a technical enhancement but also it is a safer, more accountable and adoptable pathway for AI systems in terms of industrial maintenance [4,6].

To achieve the principles of Industry 4.0, one of the most important tasks is PdM, which transforms large volumes of sensor machine data into actionable information regarding system reliability [1,5,8]. According to the recent roadmaps, PdM can be classified into three main types, namely, data-driven, physics-based, and hybrid approaches, and researchers still encounter issues related to model performance in different scenarios, secure processing of confidential data, and integration of knowledge across different operating conditions [8]. Since smart manufacturing is a cyber-physical, data-intensive system that relies on shareable data and high-quality sensing and analytics throughout the entire product life cycle [1,5], education should familiarize students with the entire PdM workflow, starting with sensing and processing and ending with modeling, evaluation, explainability, and decision-making. With deep learning methods advancing PdM system performance, explanation of maintenance actions becomes a necessary step to guarantee transparency and reliability and consistency with domain knowledge in maintenance activities. Previous research on XAI-PdM has concentrated on establishing trustworthiness and safety, explaining sensor patterns that change over time, and providing explanation methods that match the way maintenance teams analyze system alerts [4,9,10]. This is in line with the emphasis of NIST on reliable smart manufacturing, which also incorporates performance, reliability and data traceability [1,5]. The standard XAI tools are limited to LIME/SHAP for the classical ML and Integrated Gradients or saliency methods for deep models which are applicable in XAI-PdM labs [11-13]. In terms of educational frameworks in smart manufacturing and Industry 4.0, prior works addressed course-level instruction, lab-based learning environments, and broader curriculum modernization [14-16]. While these studies support manufacturing-oriented digital and AI education, they do not appear to center explainable predictive maintenance as a structured teaching workflow. To our knowledge, fewer studies have translated these combined ideas into a unified curriculum framework that integrates sensing, preprocessing, modeling, explainability, and maintenance decision documentation for smart manufacturing education.

This study is directly connected to manufacturing because it addresses manufacturing assets, industrial sensor streams, and maintenance decision-making under real production constraints [1-5]. Based on this, the proposed XAI-PdM framework is presented as a smart manufacturing educational approach, but not as an AI educational tool that can be used for any purpose. Although engineering and ET programs are increasing Industry 4.0 content through project-based learning curricula and PdM-focused initiatives [6,17,18]. Many courses still teach model accuracy and tool use with little instruction on how to document data context, prevent leakage and spurious correlations, link explanations to physical faults, and use predictions to make maintenance decisions based on cost, safety, and downtime. The AI-based diagnostics and prognostics are growing fast in advanced manufacturing and therefore ET programs must train students to this workforce reality [2,6,7]. The paper addresses these needs by:

- Advancing PdM education beyond accuracy or tool-usage training to implementation-ready applications (data context, leakage/spurious relationships, and decision boundaries)
- Integrating the end-to-end PdM workflow: sensing → processing → modeling → evaluation → explainability → maintenance action.
- Treating explainability as necessary proof through connecting model outputs to physical fault analysis and justification.
- Using industry-style, assessable deliverables (e.g., dataset/model documentation and maintenance action memos) instead of accuracy-only grading.
- Focusing on realistic testing and maintenance-related assessment to educate students to develop and deliver constraint-based recommendations (cost, safety, downtime).

1.2 Importance of PdM in engineering technology programs

PdM is an important field for ET education because it teaches students how to use hardware and instruments with data analysis to make maintenance decisions. These are essential skills for smart manufacturing jobs. Through PdM training, students will be able to relate what they learned in basic engineering classes, such as vibration, motion, electronics, and reliability, with what they are teaching with modern, data-based processes. Explainable AI (XAI) reinforces their learning by teaching students to justify their decisions with evidence, rather than blindly trusting the output of models. The XAI-PdM course in Figure 1 covers multi-sensor data collection and metadata tracking up to modeling, explanation and maintenance decision support, and evaluates students with industry-paper based deliverables (dataset cards, feature reports, model cards and maintenance memos) rather than just the prediction accuracy; thus, the curriculum focuses on measurable learning outcomes like:

- Acquiring and managing multi-sensor datasets.
- Applying time, frequency, and time-frequency analysis

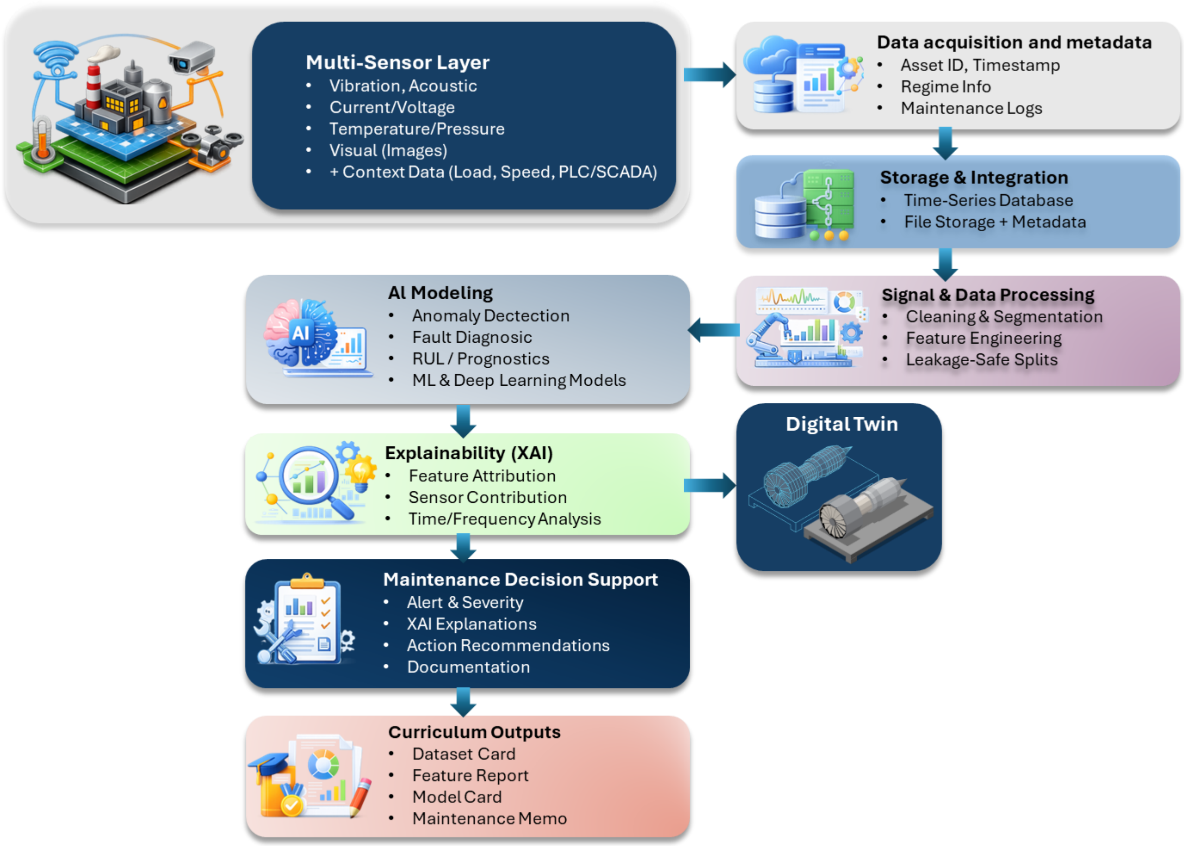

Figure 1. End-to-end XAI-PdM workflow for smart manufacturing education, starts from multi sensor data acquisition to ultimate curriculum outputs.

- Building ML/DL models for fault analysis.
- Evaluating models with appropriate metrics and validation strategies.
- Generating and interpreting explanations.
- Translating predictions into maintenance actions under constraint (Cost, Safety, downtime) .

PdM education also supports workforce readiness by mirroring how maintenance decisions are made in industry. In practice, decisions are not based on model accuracy alone; they depend on trust, interpretability, and documentation that enables validation and accountability. Accordingly, students submit industry-aligned artifacts such as dataset cards, preprocessing logs/feature reports, model cards, explanation reports, and maintenance action memos. These deliverables reinforce clear communication and traceability, both essential for successful adoption of AI in manufacturing organizations. Overall, PdM and XAI-PdM provide a powerful educational platform for preparing students to design, interpret, validate, and deploy data-driven maintenance solutions in smart manufacturing environments through an experiential learning approach grounded in realistic datasets and explainability tools. In this paper, we use term curriculum framework to refer to a modular instructional structure. This structure may be implemented as a course, a lab sequence or a project-based module using public PdM datasets, laboratory generated datasets or institution/industry provided industrial datasets, depending on the data availability and institutional context.

2 Key applications of data-driven AI in PdM

The AI in smart manufacturing systems does not start with a model but with data collected from an operational machine under various loads, speeds, combinations, and environmental conditions. Therefore, the primary use of AI in PdM is the development of a workflow that utilizes raw signals to generate accurate maintenance information. Educators should ensure that students understand that the success of PdM programs depends on pragmatic engineering decisions regarding what to measure, how to store, prepare, and maintain it in service. Students are also taught the proper way of collecting diverse sensor data in a secure and consistent manner and the proper way of organizing it to be reusable and maintain its physical characteristics of failure signatures.

2.1 Data collection and storage: sensor fusion and signal types

PdM uses data from various sources since the machine's health is rarely detected in a single measurement. In a rotating machine, a bearing can fail with the vibration signal being the first indication of the problem, followed by the thermal signal. The change in a certain load will be apparent in the current value of the motor. This is why sensor fusion is primary to PdM. It unifies signals that enhance the health picture and decrease false alarms. In the curriculum, students should be taught to classify following signals in practical terms:

1. High-rate continuous signals (vibration, acoustics): Capture fast mechanical events and are critical for early fault detection; require careful sampling-rate selection, windowing, and data-volume management.
2. Medium-rate condition signals (motor current, temperature): Often reflect load variation, friction, imbalance, and heating; typically, cheaper to instrument and easier to collect at scale.
3. Low-rate process/context tags (PLC/SCADA/MES variables): Speed, torque, recipe, pressure, and operating mode strongly influence sensor patterns; without these context variables, models may confuse normal regime changes with degradation.
4. Image-based data (thermal or visual inspection): Useful for specific defect modes but require different storage formats and preprocessing pipelines than time-series signals.

A usable dataset module should teach metadata discipline so students can answer deployment-critical questions (e.g., which asset produced the signal, under what load and time, and what maintenance event followed). Accordingly, datasets should include sensor readings linked to asset IDs, operating regimes, maintenance logs, and quality flags (e.g., missing data, sensor replacement, calibration dates). Storage choices should also be part of engineering decision-making: time-series data can be saved in structured arrays or historian-style exports, while images/audio can be stored as files connected through metadata tables; for class projects, a consistent, well-documented folder and schema is often sufficient.

2.2 Signal/data processing

After data is collected, the subsequent primary task involves converting the raw signals into forms that are stable and comparable, and meaningful to modeling. The industrial signal data transmitted in actual time is quite disorganized. Noise, transient and missing regions must be present. Thus, the PdM course needs to consider preprocessing as a necessary engineering step. A normal preprocessing workflow that is curriculum-friendly can be described in the following way. The first major step is transforming raw signals into stable signals. These are useful when modeling and can be compared. The industrial signals used in real-time industrial applications are not clean. There are always distortions, fleeting and missing parts. Therefore, PdM courses should regard preprocessing as a fundamental element of engineering. A typical preprocessing pipeline that is curriculum-friendly is applied in the following steps.

1. Quality checks and cleaning: identifying missing data, saturation, sensor drift, and abnormal spikes; apply basic filtering or denoising as justified.
2. Segmentation/windowing: define window length and overlap; ensure segmentation aligns with machine dynamics and sampling rates.

3. Normalization and operating-regime awareness: account for changes due to load/speed/mode and avoid treating regime shifts as faults.
4. Feature engineering or representation learning: extract time-, frequency-, and time–frequency-domain features (or build learned representations such as spectrograms).
5. Data splitting to avoid leakage: design deployment-realistic splits (e.g., by time, by asset, or by operating regime) to prevent inflated performance estimates.

3 AI-enabled data-driven decision making in PdM

In smart manufacturing, PdM creates real value only when it supports maintenance decision-making, and AI is most effective when it helps personnel identify whether a machine is degrading, which fault is most likely, the associated risk level, and what action is most urgent. Therefore, an effective XAI-PdM curriculum should emphasize decision-support, not just prediction accuracy, by linking (1) processed data and engineered features, (2) modeling methods that translate patterns into predictions, and (3) operational constraints that determine whether to inspect, repair, continue operation, or keep monitoring. The instructional goal is to develop engineering judgment with data: students should justify their modeling choices, evaluate performance under realistic deployment conditions, and translate model outputs (with explanations) into actionable maintenance recommendations. Figure 2 illustrates the support of XAI explanations to the processed data and features, the predictions of AI/ML, and the combination of operational constraints to generate workable maintenance decisions.

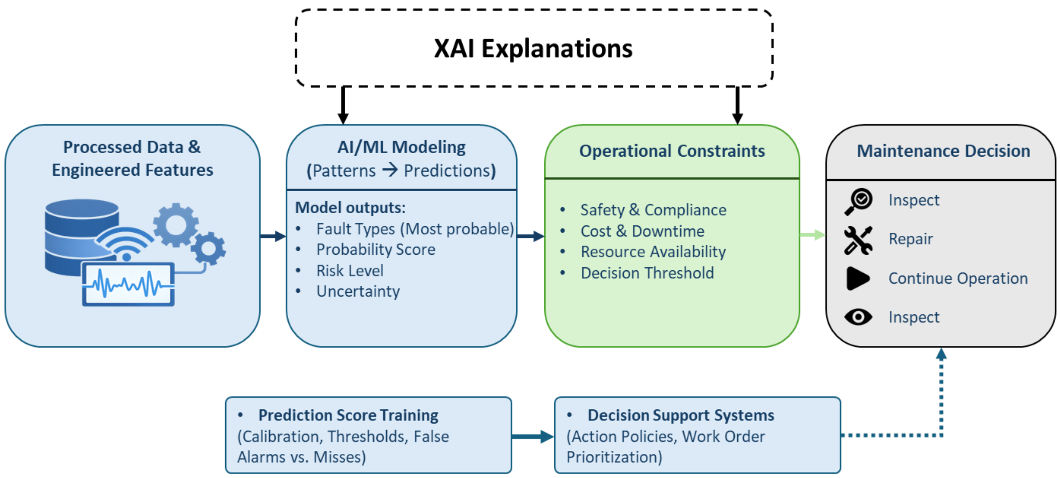

Figure 2. AI-enabled data-driven decision-making workflow for explainable predictive maintenance.

3.1 Implementation of available techniques

In educational settings, students learn most effectively by repeatedly working through a structured project lifecycle that starts with problem definition through the measurement of results. This lifecycle includes building a base version, improving upon it, checking results, drawing conclusions and deciding on a course of action. By introducing the task of understanding the data before progressing to building a complex model, this scaffold allows students to consider the relevance of the data before proceeding with the creation of their model. This aligns well with explainability, as correct evaluation and validation must be carried out for any explanation to be meaningful. In typical courses in Predictive Data Modelling with XAI, students implement models for one or more of the following tasks in Predictive Data Modelling:

1. Anomaly detection: determining whether machine behavior deviates from expected operation.
2. Fault diagnosis (classification): identifying the most likely fault type (e.g., bearing defect, imbalance, misalignment).
3. Prognostics / remaining useful life (RUL): estimating time-to-failure or remaining operational life.

Courses should clearly show students how to pick a method based on practical factors: trade-off between missed failures and false alarms, sensor type and sampling method, quality and quantity of labels, clarity needs, and making decisions in real-time. Student documentation of assumptions and constraints, and their work environment introduces them to methods similar to industrial engineering practice and secure system implementation.

3.1.1 Machine Learning model

The classical ML models offer an advantage for introductory PdM lessons since they train learners to consider features, physical sense and model interpretability. Feature-based mode ls can still be useful in manufacturing environments where they can be easily verified and are not as complex as DL models and hence may be used in training the workforce. A curriculum-friendly progression is:

1) Interpretable baselines: establish physical intuition before introducing higher-complexity models.
   a) Threshold rules based on domain health indicators, b) Linear or logistic regression as transparent baselines.
2) Stronger classical ML methods: Compare performance improvements, and document trade-offs using stronger traditional ML methods. For example:
   a) Decision trees and random forests (easy to explain and robust)
   b) Gradient boosting methods (often high accuracy with engineered features)
   c) SVM (support vector machine) for structured feature sets
3) Unsupervised or weakly supervised models: These models imitate industrial conditions where fault labels are not available or expensive to obtain. Examples include one-class based approaches and the Isolation Forest to identify irregularities.

Students can conveniently use XAI in ML training with global and local explanation techniques to understand which sensor attributes are responsible for the model on the whole and why this definite case was chosen. Representative explainability tools suitable for the proposed curriculum include SHAP and LIME for classical ML models, and Integrated Gradients or saliency-based methods for deep learning models. Clarifications are useful in debugging. When the explanations repeatedly mention the operational regime variables instead of the health-related indicators, the students may notice that the model is probably responding to regime changes rather than degradation, which is one of the important failure modes in real-world implementations.

3.1.2 Deep Learning model

When sensor data is complicated and multilayered and cannot be easily described manually, DL can be helpful in preventive diagnostics of systems (PdM). In DL it is increasingly applied to the analysis of vibration and acoustic data, spectrograms, thermal images, and multi-channel time-series data in the smart manufacturing sector. Educationally the following typical DL methods are appropriate:

a) 1D CNNs over the time-series raw data (e.g. vibration, motor current).
b) 2D CNNs schemes of time-frequency data (e.g. spectrograms).
c) Sequential damage and time relationship LSTM/GRU models.
d) Transformer/attention-based models when long-range dependencies and context integration are critical.

ML develops rationality through feature value and comparatively simple decision borders, and DL focuses on representation learning and the ability to process complex sensor data. Hence, XAI tools are essential to relate the DL output to engineering knowledge. In a DL-based XAI-PdM module, students can create interpretable models which point out critical time intervals, and frequency bands and (in case of multi-channel models) sensor-contribution indices, which support fault diagnosis and maintenance decision-making.

3.1.3 Evaluation metrics

XAI-PdM evaluation should be taught in terms of both predictive performance and maintenance impact, emphasizing deployment-realistic validation rather than misleading practices like random window splitting, which can inflate results because adjacent time windows share similar signal patterns and operating regimes. Students should use realistic splits (by time, asset, or regime) and report task-aligned, cost-aware metrics for fault diagnosis, anomaly detection, and RUL prediction, interpreting each metric through its maintenance consequences (summarized in Table 1). Explainability should serve as a second validation lens: even high-scoring models can be unreliable if they are “right for the wrong reasons” (e.g., responding to regime changes instead of degradation), so evaluation should assess both predictive accuracy and whether XAI provides useful, trustworthy evidence to support maintenance decisions.

Table 1. Recommended evaluation of metrics and their maintenance relevance for PdM tasks

Area	Metrics	Why it matters

Classification (Fault diagnosis)	• Precision, Recall, F1- score • Confusion matrix • PR-AUC	• Fault data are often imbalanced; these metrics are more informative than accuracy. • Reveals which fault types are confused (e.g. imbalance vs misalignment), supporting targeted improvement. • Often more It is often more informative than ROC-AUC when the number of faults is small and the cost of false alarms is high. full than ROC-AUC when faults are rare
Anomaly detection	• Detection delay • Threshold selection • False alarm rate	• This is of some benefit when the message is received early enough to allow the operator to act. • Thresholds must be based on the trade-offs in operations at the level, not random tuning. • High false-alarm burden reduces trust and operational adoption.
RUL/ prognostics	• MAE / RMSE • Early vs late cost • Uncertainty awareness (prediction intervals/ confidence)	• Indicative of the error in estimating the remaining useful life. • Predicting too early wastes useful life; predicting too late risks failure. students Should discuss this trade-off. • A plan requires value and confidence in the value.

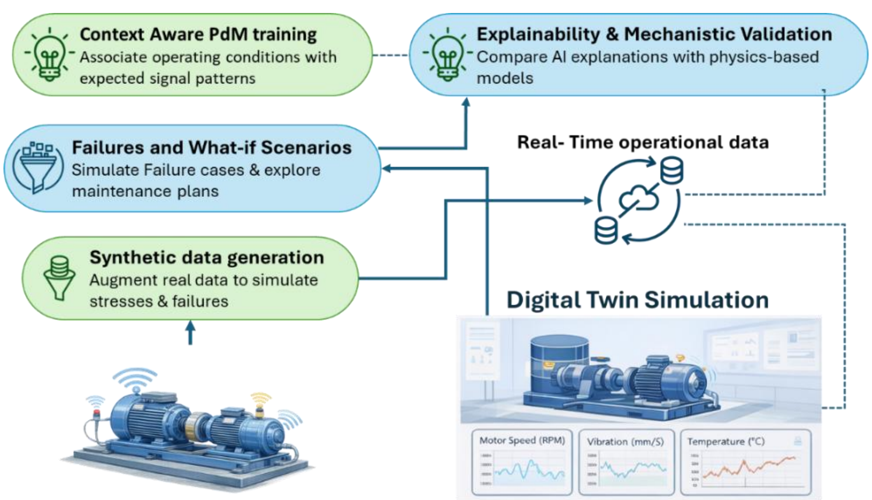

Figure 3. Digital twin simulation updated by real-time sensor data to support explainable PdM.

3.1.4 Digital Twin as an extension for XAI-PdM

Digital twin can enhance the training of XAI-PdM by providing a structured and context-sensitive model of the actual asset and updating it with operational data [19]. In PdM the twin is used to associate sensor reading with the state of operation and expected future behavior which facilitates learning and decision making. Figure 3 depicts the way in which a digital twin simulation is used to support XAI-PdM training through context-aware PdM learning, generation of data, analysis of failures scenario and explainability with mechanistic validation. A digital twin in the classroom allows several high-value educational operations. Initially, it facilitates context-aware PdM in that students are able to associate the operating conditions, e.g., speed, load, duty cycle, and mode, with the expected signal patterns, which reduces the number of the false alarms due to the normal production variation. Secondly, it supports the production of synthetic data and rare-failure learning such that simple physics-based or simulation-informed models can be used to simulate stress conditions and signatures of failures in cases where there is little real failure data. Third, it can improve explainability by being a separate standard: the students are able to compare the AI derived explanations (e.g., sensor/feature importance) to the physics-based expectations of the twin, and the idea that credible explanations should be based on engineering mechanisms is strengthened. Lastly, the twin is useful in maintenance planning since the students get to try what-if scenarios such as continued operation or planned maintenance and see their impact on risk and downtime. Thus, digital twins are not just a representation of an object, they are a strong teaching tool that bridges data-based understanding with physical reasoning and operational outcomes.

3.2 Decision-making

This section is important to an XAI-PdM course since it associates model outputs with actual maintenance actions. PdM is not just about predicting failures, but also about selecting the proper response at the proper time, given production time constraints, safety requirements, inventory, and downtime costs. The students acquire this skill by doing a structured decision record for every alert, which is depicted in Table 2 as a maintenance action memo, that what was detected, severity and confidence, what global/local explanation evidence, physically plausible degradation explanation, and what feasible response.

4 Discussion

XAI-PdM curricula involve students in evidence-based maintenance-decision learning and not model-based learning. In standard AI-PdM courses, the primary concern of students is the development of highly performing models and accuracy reports. Human maintenance decisions do not rely solely on accuracy. Responsible decision makers, including the supervisors, engineers and safety personnel, have to provide clear justification before the allocation of the resources or the stoppage of production or the issue of the order. They must comprehend the reason for an alert and the evidence and the level of confidence of the system in the alert under certain circumstances. Explainable AI (XAI) takes a core position in PdM training since it is a practical tool that establishes the connection between data-driven models and engineering criticality, through transparency, traceability and actionable reasoning.

Table 2. Decision-making template for XAI-PdM: mapping alerts and explanations for maintenance actions.

Step	Component	What it includes
Alert formation	• Detection • Severity • Confidence	• What was detected (anomaly, fault, or RUL threshold). • Severity level (risk score or urgency level). • Confidence score, prediction interval, or uncertainty notes.
Evidence and explanation	• Global context • Local evidence • Engineering interpretation	• What typically drives the model. • Explanation of why this particular instance was detected (local explanation). • Conversion of XAI evidence to physical reasoning (why these sensors/features suggest degradation).
Action recommendation	• Response type • Troubleshooting priority • Parts ordering	• Immediate check or planned service or continued monitoring • What to check first. • Whether to pre-order parts according to lead time and confidence
Decision documentation	• Maintenance memo (deliverable) • Assumptions/limits • Follow-up plan	• A brief memo indicating recommendation, evidence and impact expected. • Operating regime assumptions, limits known and/or failure modes. • What other information or tests proves the diagnosis and how result data.

One of the primary benefits of this research is that it employs deliverable-based assessment, which is similar to the practice of industrial documentation, as opposed to assessing students based on predictive performance only. Figure 4 (curriculum Deliverables Workflow) implements the entire PdM pipeline by means of a set of artifacts prepared by students. The figure demonstrates that reliable PdM starts upstream before the training of the models with careful construction of datasets and proper fabrication of decisions regarding metadata management and preprocessing. In the data collection and storage phase, students describe sensor specification, sampling and synchronization, marking protocol referencing maintenance logs, and data schemas. These artifacts are used to directly solve common deployment errors such as the missing context, inconsistent labeling, and untraceable data provenance. A processing log that can be replicated and that contains cleaning, segmentation, regime-aware normalization, parameters of representation (FFT/STFT or time-frequency features), leakage-safe splits are required at the signal/data processing stage. The curriculum requires these specifics as graded deliverables and thus supports the idea that claims of performance are valid only when data preparation is correct and documented and repeatable.

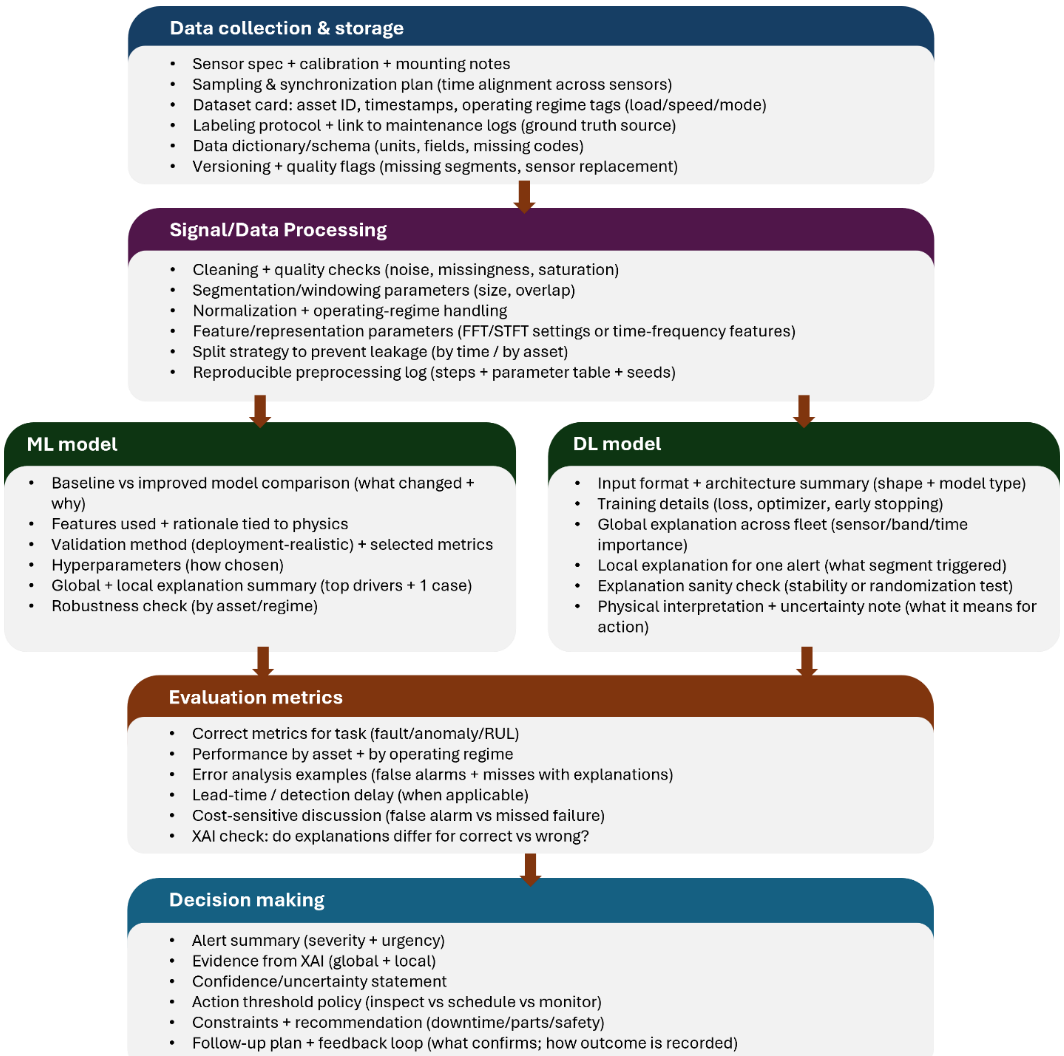

Figure 4. Curriculum deliverable workflow for XAI-PdM. Student artifacts from multisensory data collection to maintenance decision documentation

The illustration also demonstrates the ML and DL routes, maintaining a standard requirement of interpretability. In ML models, students evaluate baselines in the context of enhanced models and justify their selection of features with information based on physics and provide global/local explanations. As for DL models, students are required to report the architecture and training choices as well as explanation reports, which pinpoint the considerable time segments, frequency regions, and sensor contributions. Most significantly, the structure of the deliverable precludes 'black-box acceptance' because of the fact that explanations are used as part of the engineering evidence that will be used to support decisions but not as supplementary visualizations. The succeeding stage reinforces the realism of the deployment by necessitating performance reporting, error analysis, and cost-sensitive discussions of false alarms as opposed to missed failures. The end-use document in the decision-making stage is a maintenance work order that combines severity/urgency, evidence from XAI (global + local), confidence/uncertainty, action threshold policies (inspect vs. schedule vs. monitor), operational constraints, and a follow-up plan with a feedback loop. This last artifact has been deliberately designed to fit the communication and validation of PdM recommendations in manufacturing firms.

Table 3 summarizes how this curriculum approach supports learning across four critical dimensions: educational viewpoint, realistic data, human-centered decision-making, and transferability in constraints. Educationally, the deliverable workflow eliminates accuracy-only thinking since it requires students to interpret model performance and support choices with evidence. In data realism, the workflow employs unreliable industrial data as an expected condition and educates students to experiment with regimes, address imbalance, and test splits to prevent leakage. Human-centered decision-making is promoted by converting predictions into a formal decision template that provides evidence and constraints, which enables defensible actions as opposed to label-only results. Finally, transferability indicates that many courses are run under limited hardware, computing, or instrumentation, and the curriculum can still be implemented using scalable baselines and selective DL, but explanations must be provided to support decisions under actual constraints.

Table 3. Discussion summary of XAI-PdM curriculum focus areas: educational value, data realism, human-centered decision-making, and transferability.

Focuses	Why matters in XAI-PdM	What students do	What XAI reveals
Educational perspective	It prevents from “accuracy only” thinking and forces to interpret model behavior.	Train baseline model, generate explanation for all cases, Compare fault analysis data	The model learns the true failure-related signal vs shortcuts.
Data realism	Real PdM data are imperfect. Students must debug generalization.	Test across operating regimes, check imbalance and verify splits and then run stress tests	Machine leakage, sensor-dominance, unstable explanations
Decision making with human intelligence with AI	Providing just labels for maintenance does not enough, they need defensible action	Convert predictions into decision template (risk/cost/ constraint)	What triggered the alarm and what to inspect first.
Transferability and constraints	Some courses must have limited computing or hardware	Compare lightweight baseline vs DL, choose tools based on resources	Whether explainability still supports decision making under constraints (tradeoff)

Overall, the curriculum deliverable workflow strengthens workforce readiness by teaching students not only to build predictive models but also to produce auditable, explainable, and action-oriented maintenance recommendations. By aligning course outcomes with industrial adoption requirements, transparency, traceability, and decision utility, this approach supports trustworthy AI deployment in smart manufacturing settings. In addition to increasing technical ability, the framework being presented will build students awareness of ethical use of AI by developing critical analysis skills when reviewing "black box" model outputs while also documenting assumptions and limitations, identifying and evaluating questionable relationships and providing transparent, traceable and accountable reasoning on how and when to maintain AI models.

5 Conclusion and future outlook

This paper argues that XAI-PdM is a timely and effective educational theme for smart manufacturing and ET curriculum. The approach suggested harmonizes model training and accuracy-based learning with total industrial PdM processes and required documentation practices to ensure reliable adoption. The training activities undertaken by the students are as follows:

1. Data collection and organization with multiple sensor inputs,
2. Signal processing to produce significant physical representations,
3. ML and DL model training for PdM applications,
4. Performance evaluation based on real-world validation methods and maintenance-based metrics,
5. Explanation generation to summarize model reasoning processes,
6. Conversion of predictions and explanations into justifiable maintenance actions.

Educationally, this paper provides a curriculum that is ready to teach as a course which combines theoretical knowledge with practical implementation. Students gain expertise in sensing technologies and signal analysis and AI model development and operational decision making through creating industry standard documents such as dataset cards and preprocessing logs and feature reports and model cards and explanation reports and maintenance action memos. The students gain practical experience with essential skills for AI manufacturing implementation which include explaining AI decisions and building user trust and identifying system errors and debugging and maintaining detailed records and presenting information clearly to stakeholders. The XAI-PdM curriculum prepares students to create and evaluate and implement PdM systems that function as transparent and traceable and operational in manufacturing settings.

Multiple advances can be used to enhance XAI-PdM technology and to make it more suitable in educational context:

1. Large-scale Multimodal PdM: In the future, PdM systems are expected to integrate vibration, acoustic, electrical/current, thermal and process context data within unified multimodal monitoring frameworks. Teaching programs may follow the development towards the combination of several types of data and trial between various types of sensing paths.
2. Uncertainty-based decision-making: Maintenance planning requires confidence and predictions only. Students, learning to measure and communicate uncertainty (such as prediction intervals, calibration or ensemble agreement) will be more prepared for real situations and more inclined to trust the outputs of the models.
3. Edge deployment and real-time constraints: Many PdM models must react rapidly by running on devices close to equipment. Students need to learn lightweight models and speed/accuracy tradeoffs, and explanation techniques that work on edge platforms.
4. Human-centered XAI: Explanations must match the information needs and methods of decision-making of stakeholders. Explanations should fit the user's job and decision process. We need better explanation testing, role-based reporting (operator or reliability engineer or manager), and human checks to ensure XAI helps, not hinders, maintenance work.
5. Hybrid learning models that utilize both physics-based models and data-driven models show better performance in paradigm generalization, model explanation, and model reliability. PdM researchers and students need to learn how to combine AI methods with engineering concepts, and this approach can help them achieve good learning experience.

The proposed XAI-PdM based curriculum provides a structured and potentially effective framework for preparing students who will work in smart manufacturing environments where AI systems must be both accurate and interpretable. The ongoing development of AI technology and industrial data systems creates expanding opportunities for students to build, evaluate, explain, and deploy PdM solutions. A curriculum that emphasizes explainability, traceability, and decision utility can help broaden participation in advanced manufacturing and support innovation in sustainable, reliable industrial operations.

6 References

[1] Barkmeyer, E., & Wallace, E. K. (2016). *Reference architecture for smart manufacturing part 1: Functional models*. Gaithersburg: US Department of Commerce, National Institute of Standards and Technology.

[2] Singh, P. T., Pothier, D. A., Sethupathy, A., Gooding, G., Soukup, D., Reich, J., ... & Warrick, J. S. (2024, November). Industry 4.0 and Modernizing Manufacturing Education. In *ASME International Mechanical Engineering Congress and Exposition* (Vol. 88650, p. V007T09A002). American Society of Mechanical Engineers.

[3] Tsallis, C., Papageorgas, P., Piromalis, D., & Munteanu, R. A. (2025). Application-wise review of Machine Learning-based predictive maintenance: Trends, challenges, and future directions. *Applied Sciences*, *15*(9), 4898.

[4] Cummins, L., Sommers, A., Ramezani, S. B., Mittal, S., Jabour, J., Seale, M., & Rahimi, S. (2024). Explainable predictive maintenance: A survey of current methods, challenges and opportunities. *IEEE access*, *12*, 57574-57602.

[5] National Institute of Standards and Technology (NIST), "Trustworthy Systems, Components, and Data for Smart Manufacturing Program," *Programs and Projects*, NIST. Available: https://www.nist.gov/programs-projects/trustworthy-systems-components-and-data-smart-manufacturing-program. Accessed: Jan., 2026.

[6] Mullett, G. (2023, February). Modifying Engineering Technology Curriculum to Adapt to the Demands of Industry 4.0. In *2023 CIEC*.

[7] Niemirowski, J. A., Cruse, K. C., & Hall, D. (2023, June). Implementation and evaluation of a predictive maintenance course utilizing machine learning. In *2023 ASEE Annual Conference & Exposition*.

[8] Meriem, H., Nora, H., & Samir, O. (2023). Predictive maintenance for smart industrial systems: a roadmap. *Procedia Computer Science*, *220*, 645-650.

[9] Pashami, S., Nowaczyk, S., Fan, Y., Jakubowski, J., Paiva, N., Davari, N., ... & Gama, J. (2023). Explainable predictive maintenance. *arXiv preprint arXiv:2306.05120*.

[10] Amaliah, N. R., Tjahjono, B., & Palade, V. (2025). Human-in-the-Loop XAI for Predictive Maintenance: A Systematic Review of Interactive Systems and Their Effectiveness in Maintenance Decision-Making. *Electronics*, *14*(17), 3384.

[11] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). " Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (p. 1135-1144).

[12] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, *30*.

[13] Sundararajan, M., Taly, A., & Yan, Q. (2017, July). Axiomatic attribution for deep networks. In *International conference on machine learning* (pp. 3319-3328). PMLR.

[14] Laky, Montgomery D., et al. "Smart Manufacturing Course: Proposed and Executed Curriculum Integrating Modern Digital Tools into Chemical Engineering Education." *Systems and Control Transactions* (2025): 2271-2276.

[15] Cuiffi, Joseph Dennis, et al. "Factory 4.0 Toolkit for Smart Manufacturing Training." *2021 ASEE Virtual Annual Conference Content Access*. 2021.

[16] Ralph, Benjamin James, et al. "Implementation of a six-layer smart factory architecture with special focus on transdisciplinary engineering education." Sensors 21.9 (2021): 2944.

[17] Nithyanandam, G., Munguia, J., & Marimuthu, M. (2022). “Digital literacy”: Shaping industry 4.0 engineering curriculums via factory pilot-demonstrators. *Advances in Industrial and Manufacturing Engineering*, *5*, 100092.

[18] Chiou, R. Y., Kavitha, A. K. S., Tseng, T. L. B., & Vasudevan, N. (2025, June). Integration of FPGA-Accelerated AI for Predictive Maintenance Education in Industry 4.0. In *2025 ASEE Annual Conference & Exposition*.

[19] Ma, S., Flanigan, K. A., & Bergés, M. (2024). State-of-the-art review and synthesis: A requirement-based roadmap for standardized predictive maintenance automation using digital twin technologies. *Advanced Engineering Informatics*, *62*, 102800.