

# Machine Learning in Industry 5.0

Pál Varga, and Péter Olaszi

**Abstract**—Industry 5.0 extends the productivity-oriented perspective of Industry 4.0 by placing greater emphasis on human-centricity, resilience, and sustainability in manufacturing. Although deep neural networks dominate much of current AI research, traditional machine learning (ML) methods remain widely used in industrial practice. Their continued relevance stems from their interpretability, data efficiency, and suitability for deployment on resource-constrained hardware. This paper examines their role through the European AIMS5.0 project. More than 40 traditional ML, optimization, and signal-processing methods are mapped to 13 industrial use cases covering machine tools, intralogistics, semiconductor manufacturing, supply chains, and other production settings. We describe the main concepts behind the different method families and summarize the observed selection logic in a decision diagram based on task type, label availability, and the need for operator-facing explanations. Across the use cases, a common human-in-the-loop pattern emerges: expert feedback is used to refine the models, while operator actions connect model outputs back to the physical process. The reported outcomes include an increase in overall equipment effectiveness and reduction in electricity cost, as well. At the same time, the use cases show that strong technical performance does not necessarily translate into scalable deployment. Based on these findings, we derive a research agenda centered on data-efficient learning, controlled adaptation to concept drift, decision-oriented and resource-aware ML, and evaluation criteria that extend beyond predictive accuracy.

**Index Terms**—Industry 5.0, machine learning, industrial AI, anomaly detection, explainable AI, human-in-the-loop, predictive maintenance, AIMS5.0

## I. INTRODUCTION

Artificial Intelligence (AI) and Machine Learning (ML) extend industrial digitalization by enabling systems to infer patterns from data and use them for prediction, classification, anomaly detection, and decision support. These capabilities complement, rather than replace, conventional automation. automation executes predefined logic, whereas ML can derive models from operational data and update decisions when conditions or data distributions change. In Industry 4.0, this is particularly relevant to tasks such as predictive maintenance, data-driven quality assessment, process monitoring, and optimization.

Industry 5.0 broadens the objectives of industrial development beyond productivity and efficiency to include human-centricity, resilience, and sustainability [1]. These objectives are not achieved by AI/ML alone. Instead, AI/ML provides specific analytical and adaptive capabilities within a wider socio-technical system. For example, ML can estimate equipment failures or process deviations before they occur, support

operators with data-driven recommendations, and identify operating parameters associated with lower resource consumption. The resulting actions still depend on automation systems, domain constraints, human decisions, and the surrounding digital infrastructure. This distinction is important when assessing the contribution of AI/ML to Industry 5.0.

This paper is about ML in Industry 4.0 and 5.0, and separates traditional ML solutions from neural-network-based AI approaches. Traditional ML methods rely on structured feature representations and statistical models – such as regression, decision trees, ensemble methods, or support vector machines – to learn predictive relationships from data [2], [3]. These approaches typically require domain-engineered features, are computationally efficient, and offer higher interpretability, making them well suited for many industrial tasks with limited data or strict transparency requirements [4], [5]. In contrast, deep neural network (DNN)-based approaches learn hierarchical feature representations directly from raw data through multilayer architectures, enabling high performance in complex tasks such as image analysis, speech, or time-series modeling. However, DNNs generally demand larger datasets, greater computational resources, and exhibit lower inherent interpretability [6]. Thus, traditional ML and DNN-based methods complement each other: the former excels in structured, explainable industrial contexts, while the latter provides powerful end-to-end modeling for perception and high-dimensional data.

## II. A GENERAL OVERVIEW OF MACHINE LEARNING IN INDUSTRY 5.0

Machine Learning (ML) provides a suite of statistical and algorithmic tools that enable data-driven modeling, prediction, and optimization across industrial processes [4], [7]. While deep neural networks dominate much of current AI research, a wide range of non-neural ML methods remains fundamental in industrial environments, owing to their interpretability, modest data and computational demands, and ease of integration with existing control and monitoring systems [5].

This section first outlines the basic distinctions between supervised, unsupervised, semi-supervised, reinforcement, and evolutionary learning. It then explains the concepts behind the method families employed in the AIMS5.0 use cases (Table I): classification and ensemble methods, clustering and dimensionality reduction, anomaly detection, optimization and search, signal processing and time-series analysis, explainability and causal analysis, and statistical monitoring and simulation. Finally, it surveys typical industrial applications – sustainable production, resilient supply chains, collaborative robotics, and digital twins – and discusses how the data regime and the role of the human operator guide the choice among these methods in an Industry 5.0 setting [8].

P. Varga is with the Department of Telecommunications and Artificial Intelligence, Budapest University of Technology and Economics, Műegyetem rkp. 3., H-1111 Budapest, Hungary.

P. Olaszi is with AITIA International Inc., AI and IIoT Division, 7. Záhony street, 1031 Budapest, Hungary.

TABLE I  
MACHINE LEARNING ALGORITHMS EMPLOYED IN AIMS5.0 USE CASES

| Algorithm / Use case                                                 | 03 | 05 | 07 | 08 | 09 | 11 | 12 | 13 | 15 | 16 | 17 | 19 | 20 |
|----------------------------------------------------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|
| <i>Classification &amp; Ensemble Methods</i>                         |    |    |    |    |    |    |    |    |    |    |    |    |    |
| Decision Trees                                                       | X  |    |    |    |    |    |    |    |    | X  |    |    |    |
| Gradient Boosted Decision Trees (GBDT)                               | X  |    |    |    |    |    | X  |    |    |    |    |    |    |
| Extreme Gradient Boosting (XGBoost)                                  |    |    |    |    |    |    | X  |    |    |    |    |    |    |
| Regression Tree                                                      |    |    |    |    |    |    | X  |    |    |    |    |    |    |
| Random Forest                                                        |    | X  |    |    |    |    |    |    |    |    |    | X  |    |
| K-Nearest Neighbors (KNN)                                            |    | X  |    |    |    |    |    |    |    |    |    |    |    |
| Logistic Regression                                                  |    | X  |    |    |    |    |    |    |    |    |    |    |    |
| Support Vector Machines (SVM)                                        |    | X  |    |    |    |    |    |    |    |    |    |    |    |
| Iterative self-training                                              |    |    | X  |    |    |    |    |    |    |    |    |    |    |
| Binary classification                                                |    |    |    |    |    | X  |    |    |    |    |    |    |    |
| <i>Clustering &amp; Dimensionality Reduction</i>                     |    |    |    |    |    |    |    |    |    |    |    |    |    |
| K-Means Clustering                                                   |    |    |    |    | X  |    |    |    |    |    |    | X  |    |
| Density-Based Spatial Clustering of Applications with Noise (DBSCAN) |    |    |    |    |    |    |    | X  |    |    |    |    | X  |
| Euclidean Clustering                                                 |    | X  |    | X  |    |    |    |    |    |    |    |    |    |
| Principal Component Analysis (PCA)                                   |    |    |    |    | X  |    |    |    |    |    |    |    |    |
| <i>Anomaly Detection</i>                                             |    |    |    |    |    |    |    |    |    |    |    |    |    |
| Isolation Forest                                                     |    |    |    |    |    |    |    | X  | X  |    |    |    | X  |
| <i>Optimization &amp; Search</i>                                     |    |    |    |    |    |    |    |    |    |    |    |    |    |
| Genetic Programming                                                  |    |    |    |    | X  |    |    |    |    |    | X  |    |    |
| Constraint Programming                                               |    |    |    |    | X  |    |    |    |    |    |    |    |    |
| Two-stage Stochastic Programming (2SP)                               |    |    |    |    |    |    |    |    |    |    |    | X  |    |
| Chance-constrained Programming (CCP)                                 |    |    |    |    |    |    |    |    |    |    |    | X  |    |
| Simulated Annealing                                                  |    |    |    |    |    | X  |    |    |    |    |    |    |    |
| Simultaneous Perturbation Stochastic Approximation (SPSA)            |    |    |    |    |    | X  |    |    |    |    |    |    |    |
| Reinforcement Learning                                               |    |    | X  |    |    |    |    |    |    |    |    |    |    |
| Incremental learning                                                 |    |    | X  |    |    |    |    |    |    |    |    |    |    |
| <i>Signal Processing &amp; Time Series</i>                           |    |    |    |    |    |    |    |    |    |    |    |    |    |
| Time Series Analysis                                                 |    | X  |    |    | X  |    | X  | X  |    |    |    |    |    |
| Fast Fourier Transformation (FFT)                                    |    | X  |    |    |    |    |    | X  |    |    |    |    |    |
| Harmonics Analysis                                                   |    | X  |    |    |    |    |    |    |    |    |    |    |    |
| Motif Discovery                                                      |    | X  |    |    |    |    |    |    |    |    |    |    |    |
| Change Point Detection                                               |    | X  |    |    |    |    |    |    |    |    |    |    |    |
| Symbolic Aggregate Approximation (SAX)                               |    |    |    |    | X  |    |    |    |    |    |    |    |    |
| Kalman Filtering                                                     |    |    |    | X  |    |    |    |    |    |    |    |    |    |
| <i>Other Methods</i>                                                 |    |    |    |    |    |    |    |    |    |    |    |    |    |
| Common Spatial Pattern (CSP)                                         |    |    |    | X  |    |    |    |    |    |    |    |    |    |
| Linear Regression                                                    |    |    |    |    |    |    |    |    |    | X  |    |    |    |
| Shapley Additive Explanations (SHAP)                                 |    |    |    |    |    |    | X  |    |    |    |    |    |    |
| Depth-based Isolation Forest                                         |    |    |    |    |    |    |    |    | X  |    |    |    |    |
| Feature Importance (DIFFI)                                           |    |    |    |    |    |    |    |    | X  |    |    |    |    |
| Accelerated Model Explanations for Anomaly Detection (AcME-AD)       |    |    |    |    |    |    |    |    | X  |    |    |    |    |
| Balanced Anomaly and Latent Inference Framework (Extended B-ALIF)    |    |    |    |    |    |    |    |    | X  |    |    |    |    |
| Causal Rule Mining (CRM)                                             |    |    |    |    |    |    |    |    |    |    |    | X  |    |
| Statistical Process Control (SPC)                                    |    |    |    |    |    |    |    |    |    |    | X  |    |    |
| Discrete-event simulation                                            |    |    | X  |    |    |    |    |    |    |    |    |    |    |
| Markov decision process                                              |    |    | X  |    |    |    |    |    |    |    |    |    |    |
| Semi-supervised learning                                             |    |    | X  |    |    |    |    |    |    |    |    |    |    |

#### A. Basic Distinctions of Machine Learning Methods

In this subsection, let us provide a brief overview of the main distinction of machine learning methods – from supervised learning through unsupervised and semi-supervised training to the classical reinforcement learning and evolutionary methods [4], [7].

Traditional supervised algorithms such as *linear and logistic regression*, *decision trees*, *random forests*, and *support vector machines* (SVM) have long been applied to regression and classification problems in manufacturing, logistics, and quality assurance [4], [5]. Regression models enable parameter esti-

mation and demand forecasting where causal relationships are approximately linear. Decision trees and ensemble methods, including random forests and gradient boosting, are particularly effective for fault detection, product quality classification, and process optimization, where mixed numerical and categorical inputs must be interpreted under uncertainty [9]. SVMs, owing to their robustness to high-dimensional feature spaces, are widely used in vibration-based predictive maintenance, anomaly detection, and condition monitoring tasks [10].

Unsupervised methods provide insight when labeled data are scarce or unavailable. *Clustering* algorithms such as *k*-

means, *Density-Based Spatial Clustering of Applications with Noise (DBSCAN)*, or *hierarchical clustering* are widely used for sensor data segmentation, equipment behavior profiling, and energy-consumption pattern analysis [11]. *Principal Component Analysis (PCA)* and other dimensionality reduction techniques (e.g., Independent Component Analysis, t-SNE) support multivariate process monitoring and fault isolation by identifying latent variables that capture the majority of variance in sensor networks [12]. These approaches have been instrumental in process control and statistical quality monitoring, often serving as lightweight, explainable baselines for more complex hybrid models [7].

Industrial datasets often contain abundant unlabeled samples; hence, *semi-supervised* and *active learning* frameworks are gaining traction [13], [14]. Label propagation, co-training, and self-training approaches leverage limited labeled data to expand model generalization, while active learning strategies query human experts for the most informative samples. Such human-in-the-loop paradigms align strongly with the Industry 5.0 principle of human–AI collaboration, ensuring that machine intelligence complements domain expertise rather than replacing it [8].

Classical reinforcement learning methods—such as tabular Q-learning or policy iteration—are applicable to process scheduling, adaptive control, and logistics optimization when the state and action spaces are well defined [15], [16]. Similarly, *genetic algorithms* and *evolutionary strategies* are widely used for parameter tuning, production scheduling, and layout optimization, particularly where the search space is complex but gradient information is unavailable [17], [18].

In the following sections, let us provide an overview of the various ML algorithms with base references. Table I lists a heterogeneous set of methods: classical learners, clustering and anomaly detectors, mathematical optimization techniques, signal-processing tools, and explainability add-ons. The following sections explain the concept behind each group, so that the use cases of Section III can be read without consulting the primary literature. Standard textbooks for the statistical learning methods are [2], [3].

### B. Classification and Ensemble Methods

*Decision trees* split the feature space recursively, choosing at each node the feature and threshold that best separates the target values, e.g., by minimizing Gini impurity for classification or squared error for regression. When the target is continuous, the same procedure yields a *regression tree*; the Classification and Regression Trees (CART) framework covers both [19]. A single tree is easy to read but unstable, which motivates ensembles. A *random forest* trains many trees on bootstrap samples, with a random subset of features considered at each split, and averages their votes; this reduces variance while keeping the method largely parameter-free [20]. *Gradient boosted decision trees (GBDT)* instead build shallow trees sequentially, each one fitted to the residual errors (the negative gradient of the loss) of the current ensemble [21]. *XGBoost* is a scalable GBDT implementation that adds explicit regularization of tree complexity, second-order loss approximation, and sparsity-aware split finding [22].

*K-nearest neighbors (KNN)* stores the training set and labels a new sample by a majority vote of its  $k$  closest neighbors in a chosen distance metric; it needs no training phase, but prediction cost grows with the data size [23]. *Logistic regression* models the probability of a class as a sigmoid function of a linear combination of the features, which makes the fitted coefficients directly interpretable as log-odds contributions. *Support vector machines (SVM)* find the separating hyperplane with maximal margin between classes, and kernel functions extend this to non-linear boundaries [24]. Several of these models are used for *binary classification*, the case where the output is one of two states (e.g., normal/faulty, go/no-go), usually evaluated with precision, recall, and ROC-based metrics.

*Iterative self-training* is a semi-supervised wrapper around any classifier: a model trained on the small labeled set predicts labels for unlabeled samples, the most confident predictions are added to the training set, and the process repeats [13]. It is attractive when annotation is expensive, as in visual assembly monitoring, but the confidence threshold must be chosen carefully to avoid reinforcing early mistakes.

### C. Clustering and Dimensionality Reduction

*K-means* partitions the data into  $k$  groups by alternately assigning each point to the nearest centroid and recomputing the centroids, minimizing the within-cluster sum of squared distances; it is fast but assumes roughly spherical clusters and requires  $k$  in advance [11]. *DBSCAN* defines clusters as regions of high point density: a point with at least a minimum number of neighbors within radius  $\epsilon$  is a core point, clusters grow through connected core points, and isolated points are labeled noise. It therefore finds arbitrarily shaped clusters and outliers without fixing the number of clusters [25]. *Euclidean clustering* is a special case used on 3D point clouds: points closer than a distance threshold are joined into the same object, and a KD-tree index makes the neighbor search efficient [26].

*Principal component analysis (PCA)* projects multivariate data onto orthogonal directions of maximal variance, obtained from the eigenvectors of the covariance matrix. Keeping the first few components compresses correlated sensor channels into a small number of latent variables, which supports both visualization and multivariate process monitoring [27].

### D. Anomaly Detection

Anomaly detection searches for observations that deviate from the expected behavior, often without labeled faults [28]. *Isolation Forest* exploits the fact that anomalies are few and different: it builds random trees by choosing a random feature and a random split value, and anomalies are isolated after fewer splits than normal points. The anomaly score is derived from the average path length over the forest, so no distance or density needs to be computed, and the method scales linearly with the data size [29]. Because an unsupervised score does not always match what domain experts consider anomalous, *Bayesian Active Learning Isolation Forest (B-ALIF)* uses the Isolation Forest output as a prior and updates it with a small number of expert labels, selected by active learning [30].

### E. Optimization and Search

Several use cases rely on mathematical optimization or search rather than on learning in the strict sense. *Genetic programming* (GP) evolves computer programs, typically represented as expression trees, through selection, crossover, and mutation, guided by a fitness function [31]. In scheduling, the evolved trees are priority rules (dispatching rules) that combine job attributes such as due date, weight, and processing time. *Constraint programming* (CP) models a problem as variables with finite domains and declarative constraints, and solves it by combining constraint propagation, which prunes infeasible values, with systematic search; it is especially effective for scheduling with complex resource and timing restrictions [32].

When parameters are uncertain, stochastic programming is used. In *two-stage stochastic programming* (2SP), first-stage decisions (e.g., capacity or order quantities) are made before the uncertainty is revealed, and second-stage recourse decisions correct them for each scenario; the objective minimizes first-stage cost plus expected recourse cost [33]. *Chance-constrained programming* (CCP) instead requires that constraints hold with at least a given probability, e.g., that demand is met in 95% of cases [34].

*Simulated annealing* is a metaheuristic inspired by the cooling of metals: it explores the solution space by random moves, accepting worse solutions with a probability that decreases as a “temperature” parameter is lowered, which helps it escape local optima [35]. *Simultaneous perturbation stochastic approximation* (SPSA) is a gradient-free optimizer for noisy objectives: it perturbs all parameters simultaneously in a random direction and estimates the gradient from only two measurements of the objective, regardless of the number of parameters [36]. This makes it suitable for tuning physical systems, such as antenna configurations, where each evaluation is a measurement.

A *Markov decision process* (MDP) formalizes sequential decision making with states, actions, transition probabilities, and rewards; its solution is a policy that maximizes the expected cumulative reward [37]. *Reinforcement learning* (RL) learns such a policy from interaction or simulation when the transition model is not known explicitly [15]. *Incremental learning* refers to models that are updated sample by sample or batch by batch as new data arrive, without retraining from scratch, which suits non-stationary production environments and resource-constrained devices [38].

### F. Signal Processing and Time Series

*Time series analysis* covers methods that model temporal dependence, trend, and seasonality in sequentially recorded data, such as autoregressive and moving-average models used for forecasting and residual-based monitoring [39]. The *fast Fourier transform* (FFT) computes the discrete Fourier transform in  $O(N \log N)$  time, converting a signal from the time domain to the frequency domain [40]. On rotating machinery, *harmonics analysis* then tracks the amplitudes at multiples of the fundamental rotation or tooth-passing frequency, since faults such as tool breakage, imbalance, or gear damage change these harmonic patterns [41].

*Motif discovery* finds the most similar pairs of subsequences in a long time series, i.e., recurring patterns; the matrix profile computes, for every subsequence, the distance to its nearest neighbor, from which both motifs (repeated behavior) and discords (unusual behavior) can be read [42]. *Change point detection* locates the time instants where the statistical properties of a signal (mean, variance, or distribution) change abruptly, e.g., at the start and end of a machining operation or at a process drift [43]. *Symbolic aggregate approximation* (SAX) first reduces a normalized time series to segment means (piecewise aggregate approximation), then maps each mean to a letter using equiprobable breakpoints of the Gaussian distribution. The resulting strings can be searched and compared with text-based methods, and domain experts can express patterns of interest directly as symbol sequences [44].

The *Kalman filter* recursively estimates the hidden state of a linear dynamic system from noisy measurements. Each step predicts the next state from a motion model and then corrects the prediction with the new measurement, weighting the two by their uncertainties [45]. In tracking, the prediction step bridges short gaps when a sensor loses the target. *Common spatial patterns* (CSP) is a supervised spatial filtering method for multichannel signals, most notably EEG: it finds linear combinations of channels whose variance is maximal for one class and minimal for the other, and the log-variance of the filtered signals serves as a compact feature for motor imagery classification [46], [47].

### G. Explainability and Causal Analysis

Explainable AI (XAI) methods make the output of a model understandable to its users, which is a prerequisite for human oversight in Industry 5.0 [6]. *Shapley additive explanations* (SHAP) assign each feature a contribution to an individual prediction based on Shapley values from cooperative game theory, i.e., the average change of the model output when the feature is added to all possible subsets of the other features. The contributions add up to the difference between the prediction and the average prediction [48]. SHAP is model-agnostic, but efficient exact variants exist for tree ensembles such as XGBoost.

Anomaly detectors need explanations tailored to their output. *Depth-based Isolation Forest Feature Importance* (DIFFI) derives global and local feature importance from the Isolation Forest itself: features that isolate anomalous points at shallow depths, and do so more than for normal points, receive higher importance [49]. *Accelerated model explanations for anomaly detection* (AcME-AD) is a model-agnostic alternative. It builds on AcME, which computes importance scores by perturbing each feature along the quantiles of its empirical distribution [50], and adapts this to anomaly scores, providing fast local feature importance together with a what-if analysis tool [51].

*Causal rule mining* (CRM) goes one step beyond correlation. It first mines association rules (“if condition  $X$  holds, failure  $Y$  is frequent”), then tests whether each rule is causal by comparing matched groups of records that differ only in the presence of  $X$ , in the manner of a retrospective cohort study [52].

### H. Statistical Monitoring and Simulation

*Linear regression* models a continuous output as a weighted sum of input features, fitted by least squares; it remains the default baseline for yield or performance prediction because its coefficients are directly interpretable [2]. *Statistical process control* (SPC) monitors a process with control charts: a statistic such as the sample mean is plotted over time against control limits, typically three standard deviations from the in-control mean, and points outside the limits or non-random patterns signal special causes of variation [53]. Finally, *discrete-event simulation* (DES) models a system as a sequence of events (e.g., job arrival, machine start, machine breakdown) that change its state at discrete time points. It is not a learning method, but it provides the test bed on which schedulers, RL agents, and digital twins are trained and validated before deployment [54].

### I. Typical ML applications for Industry 4.0

The classical ML methods of Section II-A support numerous operational tasks: predictive maintenance, statistical process control, supply chain optimization, product defect detection, and energy-efficiency analytics [4], [55]. In the context of Industry 5.0, their roles extend toward *transparent*, *sustainable*, and *human-centric* decision support systems [8], [56]. Because they are interpretable and computationally efficient, non-neural methods remain a critical component of trustworthy and certifiable AI pipelines that must comply with emerging regulatory frameworks such as the EU AI Act [57].

There are various types of applications already in practice. In perception-oriented tasks, ML methods for image recognition and computer vision enable automated quality control, defect detection, and visual inspection, reducing human error and inspection time [58]. Predictive models applied to sensor and maintenance data drive predictive maintenance, allowing early fault detection and optimized service scheduling that minimize downtime [9]. In robotics, learning algorithms support perception, motion planning, and adaptive control, making human-robot collaboration safer and more flexible [59]. For process optimization and control, supervised and reinforcement learning models adapt production parameters in real time to improve yield, precision, and resource utilization [60].

At a systemic level, supply chain optimization leverages ML for forecasting, risk analysis, and dynamic routing under uncertainty [61], [62], while energy optimization models contribute to sustainable operations through load balancing, peak-demand reduction, and efficiency analytics [63].

1) *ML for Sustainable Production*: Sustainability in manufacturing increasingly relies on ML to optimize energy use, minimize waste, and promote circular production models [64]. Classical regression and ensemble-based predictors are often applied to model energy consumption across machines and process stages [63]. This enables adaptive control of production schedules to reduce peak demand. Reinforcement learning and optimization algorithms support dynamic resource allocation, balancing throughput with environmental performance indicators such as carbon intensity and material utilization. In predictive quality analytics, unsupervised learning methods

detect deviations that lead to rework or scrap, contributing directly to waste reduction.

Furthermore, life-cycle data combined with clustering and probabilistic models help forecast maintenance or recycling needs, advancing closed-loop material flows. Through these data-driven insights, ML transforms sustainability from a static compliance target into a continuously optimized operational objective aligned with Industry 5.0's environmental and social pillars.

2) *ML for Resilient and adaptive supply chains*: Modern industrial supply chains must respond to disruptions ranging from component shortages to geopolitical shocks – this requires predictive and adaptive intelligence. ML methods can enhance resilience by enabling probabilistic demand forecasting, anomaly detection in logistics data, and adaptive routing through reinforcement learning [61]. Ensemble models and Bayesian networks quantify uncertainty, allowing planners to make risk-aware decisions rather than deterministic forecasts [62]. Clustering and outlier analysis further identify supplier reliability patterns or hidden dependencies that may threaten robustness. ML-driven supply chains transition from reactive logistics systems toward proactive, adaptive ecosystems – a core requirement of Industry 5.0 resilience.

3) *Collaborative robotics*: Collaborative robots (cobots) are one of the physical embodiments of human-machine cooperation envisioned in Industry 5.0. Machine learning algorithms underpin perception, intent recognition, and adaptive control that allow cobots to safely share workspaces with humans [59]. Supervised and unsupervised learning methods are applied in cobots to process multimodal sensory data—vision, depth, force, and gesture. With these classification of human states and real time prediction of motion trajectories are also possible. Reinforcement learning and imitation learning frameworks optimize task execution policies through experience or human demonstrations while respecting safety and ergonomic constraints [65]. Online learning further enables cobots to adapt to operator preferences or evolving assembly tasks without complete reprogramming. These ML-enabled mechanisms transform robots from deterministic executors into context-aware collaborators, improving flexibility, safety, and worker well-being in hybrid production cells.

4) *ML for Digital Twins in Industry 5.0*: Digital twins (DTs) serve as virtual counterparts of physical assets, integrating multi-source data for monitoring, simulation, and decision support [66]. Machine learning enhances DT fidelity by learning mappings between sensor inputs and system states, estimating unobservable parameters, and forecasting future behaviors [67]. Regression and state-space models support real-time condition estimation, while clustering and anomaly-detection methods reveal deviations between predicted and actual dynamics. Surrogate ML models, such as Gaussian processes [68] or gradient boosting regressors [21], approximate complex physical simulations, enabling faster what-if analyses and control optimization.

When integrated with reinforcement learning, DTs evolve into self-optimizing entities that can test adaptive control strategies before deployment. In Industry 5.0, DTs augmented

with ML not only reflect operational efficiency but can also incorporate human, social, and environmental dimensions.

5) *Selecting Among the Method Families*: It is already clear from the algorithm use vs. use case mapping in Table I that use cases' method choice is driven less by predictive accuracy alone than by the data regime and by the role of the human. Where labeled faults are rare, unsupervised detectors (Isolation Forest, DBSCAN, PCA) provide a starting point, and weakly supervised or active-learning extensions such as B-ALIF and self-training let experts refine them with few labels. Where decisions must be justified to operators or auditors, inherently interpretable models (linear and logistic regression, shallow trees, SPC charts, SAX patterns defined by experts) or post-hoc explainers (SHAP, DIFFI, AcME-AD) are preferred over opaque predictors. Where the task is a decision rather than a prediction, as in scheduling or supply planning, the optimization family (CP, GP, stochastic programming, RL on an MDP) is combined with discrete-event simulation to test policies before deployment. Finally, signal-processing methods (FFT, harmonics analysis, change point detection, Kalman filtering, CSP) mostly act as feature extractors in front of the learners, and their physical interpretability is often what makes the overall pipeline acceptable in a human-centric, Industry 5.0 setting.

Fig. 1 summarizes this reasoning as a decision diagram.

### III. AIMS5.0 USE CASES

The AIMS5.0 project (*Artificial Intelligence in Manufacturing leading to Sustainability and Industry5.0*) is a European Innovation Action funded under the Key Digital Technologies Joint Undertaking, running from May 2023 to September 2026 with 53 participants from industry, research organizations, and academia [69]. Its aim is to strengthen European digital sovereignty by adopting, extending, and implementing AI-enabled hardware and software components across the whole industrial value chain. The results are validated in 20 use cases spanning 9 industrial domains, including semiconductor manufacturing, machine tools, lighting, food production, logistics, and supply chain management. The use cases therefore offer a broad, practice-driven view of how AI and ML are deployed in real production environments on the way from Industry 4.0 toward Industry 5.0.

This section presents the 13 use cases in which traditional ML methods play a substantial role; the algorithms they employ are summarized in Table I, and their concepts are explained in Section II. Each use case is described in terms of its industrial problem, the ML methods applied and the reasons for choosing them, and the results or lessons learned. The use cases are ordered by their project numbering, from edge-level condition monitoring (UC03, UC05) through human-centric assembly, robotics, and scheduling (UC07–UC09), to semiconductor fab operations, contamination management, and supply chain and ontology-based traceability (UC12–UC20). Taken together, they illustrate the recurring pattern discussed in Section II: interpretable and computationally lightweight methods are preferred where data are scarce, decisions must be explained to operators, or models must run on constrained

devices. Most of the use cases share the human-in-the-loop pipeline structure shown in Fig. 2.

#### A. UC03: Servitization and service delivery enabled by Zero administration

In industrial environments – especially within manufacturing, energy, and transportation sectors – there is a growing demand for real-time machine monitoring and predictive maintenance solutions. Traditional approaches that rely on collecting all sensor data for analysis are often infeasible due to bandwidth constraints, latency requirements, and data privacy regulations. This use case addresses the need for autonomous, intelligent analytics capabilities at the edge. The use case made progress in predictive maintenance by analysing data to find correlations and identifying error codes, mainly related to battery issues. Various machine learning models were tested; a time-based dimension was introduced using a *gradient-boosted decision tree (GBDT)* model to predict which lawnmowers might fail and when the failure might occur.

The *Stream Analyse platform* [70] enables running analytics and machine learning models directly on edge devices. The platform supports a variety of lightweight and interpretable algorithms, optimized for constrained environments. Typical algorithms include *decision trees* for real-time classification tasks, and *rule-based models* for threshold monitoring, alerting, and expert system-like decision-making, often tailored to domain-specific knowledge. The solution works even offline or with intermittent connection.

The following best practices are used for deploying machine learning models on the Stream Analyse platform: (1) Machine learning models are selected based on hardware constraints. Decision trees and compact neural networks are preferred for their low resource requirements and high inference speed. (2) Preprocessing (cleaning, normalizing, and transforming sensor data) is implemented directly on the device, reducing the need for external computation. (3) Feedback mechanisms allow operators or connected systems to label and validate predictions, thereby improving model performance over time through supervised learning.

While this approach brings benefits, some challenges must be mitigated. Variations in hardware capabilities across different industrial devices require specific configurations and may limit the uniform deployment of models. Industrial equipment, even of the same type, may exhibit different behaviour under changing operational conditions, making it difficult to deploy a single model across all instances without fine-tuning. Edge devices also often have limited memory, CPU, and power availability, constraining the size and complexity of deployable models. Systems must ensure inference occurs within tight time constraints, especially in real-time or safety-critical environments. Over time, changes in equipment behaviour or environmental conditions can lead to model performance degradation. Updating models without cloud connectivity or significant downtime is a concern. The above methods are primarily applied for predictive maintenance, monitoring asset value, and estimating equipment lifetime in the servitization business model.




Fig. 1. Decision diagram for selecting among the traditional ML method families used in the AIMS5.0 use cases. The three questions are the task type, the availability of labels, and whether operators need an explanation of the output. Each leaf lists the method families and the use cases (UC) that applied them (cf. Table I).

often executed on edge devices close to the machine (UC03 Stream Analyze, UC05 Arrowhead, UC13)




Fig. 2. Typical human-in-the-loop ML pipeline shared by the AIMS5.0 use cases. Sensor data are turned into features by signal processing, a traditional ML learner produces a prediction or anomaly score, and an explainer makes the result understandable to the operator. Two feedback loops close the cycle: expert labels refine the model, and operator actions change the physical process.

The use case also employs TinyML neural networks [71] and incremental learning techniques [38], which fall in the scope of neural networks within the realm of machine learning.

The deployment achieved quantifiable results in the Smart Lawn Service pilot in Germany, managing 500 robotic mow-

ers across 130 contracts. The analytical models deployed in Databricks demonstrated an 82.5% reduction in administrative costs, with one full-time equivalent now capable of managing 200 contracts compared to the baseline of 50 contracts. The AI model development cycle was reduced from

40 working days to less than one week, representing an over 80% time reduction. Multiple analytical models were successfully deployed for real-time monitoring, including Service Level Agreement (SLA) compliance tracking, site-related error scoring, customer churn risk assessment, and preventive maintenance monitoring. Edge AI models for lawn quality scoring and blade sharpness analysis (using audio sensors) were deployed to research mowers equipped with the Stream Analyze Engine. The solution advanced to Technology Readiness Level (TRL) 6-7 through side-by-side validation with actual service delivery operations, demonstrating scalability from 500 to potentially 10,000 products. This real-world validation confirmed the effectiveness of deploying lightweight decision trees and GBDT models on resource-constrained edge devices while maintaining continuous connectivity with cloud-based analytics infrastructure through the Databricks platform.

#### B. UC05: AI techniques at different layers in Machine Tool Domain

Large part heavy-duty precision machining is a long-term milling processes, when the machine works without operator supervision. Once the machine is set up and programmed with the necessary instructions, it can operate for hours without human intervention. However, errors may arise during the milling process, potentially affecting the quality of the final product. These errors may require parts of the workpiece to be re-machined, or, in the most severe cases, the workpiece needs to be discarded.

One of the most common problems in milling processes is tool breakage. Automatic tool breakage detection in milling operations is a problem of interest in the manufacturing industry, as a tool breakage can significantly affect product quality and increase production costs. To address this issue, implementing a system capable of detecting early tool breakage would be highly beneficial for reducing production costs by preventing defective parts and minimizing wasted time. Machine learning algorithms can analyse data to identify patterns indicative of faults, failures, or abnormal machine behaviour. These techniques have the potential to enhance the efficiency and reliability of milling processes.

To understand the behaviour of the data when the machine operates with a broken tool, experimental tests were conducted on a milling machine in IDEKO's laboratory. These tests involved performing machining passes under predefined cutting conditions using known tools while collecting data from accelerometers. Crucially, the tool's condition (broken or not) was predetermined, providing labelled data for tool breakage – an often challenging dataset to obtain.

The aim of this analysis is to detect tool breakage from vibration signals during the milling machining process. Therefore, an analysis based on the vibration signal has been performed. The first step in processing the vibration signal is to select a part of the signal, the part where the machine is actually machining. Then, the *Fast Fourier Transform* is applied. Applying the FFT decomposes the signal into its individual spectral components and thus provides information about its composition. In the next step, the values of the most

relevant harmonics are extracted. The number of peaks or harmonics is directly related to the number of teeth of the tool. Analysing these features extracted from the FFT signals, the pattern that was observed is that when the tool is broken the harmonic of the step frequency is smaller than that of the rotation or some multiple, and, conversely, when the tool is well, the amplitude of the harmonic of the step frequency is the largest. As the aim is to classify whether the test was carried out with a broken tool or not, several binary classification models have been adjusted to these data. More specifically, binary classification algorithms used to solve this problem have been *Random Forest*, *K-Nearest Neighbours*, *Logistic Regression* and *Support Vector Machines*. After optimising the hyperparameters and fitting the model to data, the accuracy obtained was greater than 0.95 for the four models, concluding that these models are able to correctly predict the class to which each observation belongs.

The next step involves implementing this approach under real production conditions. In a real production environment, raw vibration signals are not continuously available throughout the milling process. However, we have successfully extracted key features from these vibration signals. Consequently, the data available in real production consist of the extracted features rather than the raw signals. A significant challenge in this scenario is the absence of labelled data, as we lack examples of tool breakage events with corresponding labels. To overcome this limitation, the initial step in this approach is to segment the machining process, isolating the time series segments corresponding to each operation performed by the machine. This segmentation enables the comparison of repeated executions of the same process. By analysing these executions, we can detect any anomalies or deviations in machine behaviour during the process, thereby identifying potential issues.

The most challenging tasks included automatic data segmentation in the presence of noise and highly variable conditions. We also had limited labelled data: signals were not annotated with failure information. Ensuring model generalization across different tools, cutting conditions, and part types also proved to be difficult, as was handling multivariate time series from various sensors with potentially missing data, and adapting to different sampling frequencies and data formats in heterogeneous environments. Applying *Motif Discovery* helped automatically segmenting repetitive time series data from machining operations, which is traditionally done manually. The used algorithms include *Change Point Detection* and *Distance-Based Clustering*.

The tool breakage detection model achieved an accuracy exceeding 95% in laboratory experiments, successfully identifying breakages based on vibration-based tool harmonics analysis. The system processes high-frequency data at 2kHz with continuous KPI extraction at 1-second intervals, enabling real-time monitoring of machining operations. This infrastructure, integrated through the Arrowhead framework, allows for both continuous medium-frequency monitoring and on-demand high-frequency analysis. The automated data segmentation approach using the Motif Discovery Toolkit significantly reduces the time required for NC program segmentation

compared to manual methods. The complete solution is currently being validated in GOIMEK's shopfloor environment. It demonstrates the combination of traditional machine learning techniques with edge computing architectures for industrial predictive maintenance applications.

### C. UC07: Digital Automated Luminaire Manufacturing Platform

This use case develops an adaptive scheduler for 3D printing hubs at Signify. The goal is to automate and optimize the scheduling process, which is currently done manually. The scheduler aims to anticipate and handle various anomalies, such as disruptions in material availability, and sudden unavailability of workforce. This enables the 3D printing hubs to operate in a more efficient and effective manner.

Signify is collaborating with Eindhoven University of Technology (TU/e) on this project. To validate the performance of the scheduler algorithm, a *discrete-event simulation* model of one of the printing hubs was built using Siemens Plant Simulation. This model allows for testing and evaluating the scheduler's performance in a controlled environment. The simulation model takes into account various factors such as the number of printers, their availability, the set of orders, the work schedule of the operators, and the manufacturing rules. The discrete-event simulation model successfully validated scheduler performance under various disruption scenarios, including material shortages and workforce unavailability.

The collaboration with TU/e involves the exploration of different approaches for the scheduler. The approaches being explored include transformer-based scheduling using sequence-to-sequence models, *Markov decision process with reinforcement learning*, and *genetic algorithms*. The goal is to find the most promising approach that can handle real-time large-scale scheduling problems. In the long term, the collaboration is expected to result in the development of an intelligent scheduler that is adaptable to different requirements and scales of different printing farms. The scheduler will optimize the assignment of print jobs to printers, considering factors such as just-in-time printing, filament reel utilization, and co-production of parts belonging to the same leading order. Overall, the use of AI algorithms and the simulation model developed in collaboration with TU/e will help Signify improve efficiency and productivity in their 3D printing operations.

In addition to scheduling, the use case also addresses Assembly Line Monitoring, which focuses on using Computer Vision (CV) to monitor the execution of manual assembly work. In the context of this assembly monitoring—specifically for object detection and tracking, Automated Data Annotations were explored. This involves *semi-supervised learning*, where small labeled datasets are combined with unlabeled data; pseudo-labeling uses an *iterative self-training* approach where the model generates its own training data; and *incremental learning* gradually improves of the model through multiple training cycles. These methods improve e.g. the *Mean Intersection over Union* (Mean IoU) metric from 0.4 to nearly 0.8 for the automated annotation pipeline, demonstrating the effectiveness of these approaches. However, further analysis

revealed that the effort required to scale this solution – specifically the installation of hardware (cameras, computers) and the extensive time still needed to label and train models for the diverse assembly activities of many product variants across different workstations – did not sufficiently outweigh the benefits. Consequently, full production deployment was halted, exemplifying a "pilot-learn fast-fail fast" approach valuable for industrial AI adoption.

Use Case 7 also uses modern deep learning approaches, which are out of scope for the current chapter: VideoMAE and Florence2 models are both transformer-based deep learning architectures; the semi-supervised learning pipeline uses deep learning models for pseudo-labeling; the reinforcement learning approach for scheduling likely uses neural networks (as most modern RL implementations do); finally, "early exploration" utilizes transformer-based scheduling. In industrial setting it is often necessary to combine them with traditional machine learning techniques.

### D. UC08: AI supporting and strengthening human and manufacturing cycles

Use case 8 concentrates on the use of AI to support and strengthen human workers and manufacturing processes through material intralogistics. It demonstrates an Autonomous Mobile Robot (AMR) task and traffic management system that enables safe and energy-efficient intralogistics within a factory environment. The AMR perception system is based on sensor fusion, allowing object detection, human presence identification, and recognition of human gestures for manual AMR control when required. An active exoskeleton supports factory workers with disabilities, assisting those who would otherwise be unable, or would struggle to perform their tasks. Materials delivered to the workstation by AMRs are handled by the human worker using an AI-supported upper-body exoskeleton based on brain-signal analysis.

The use case employs a few traditional machine learning methods. *Kalman filtering* is used in the human detection and tracking system, helping to maintain continuous tracking even when sensor data is temporarily unavailable or unreliable. By predicting positions based on previous movement patterns, it fills tracking gaps and enables autonomous mobile robots (AMRs) to anticipate human movements during brief occlusions, significantly enhancing tracking robustness in dynamic factory environments.

In the exoskeleton control system, *Common Spatial Pattern* (CSP) is applied to EEG signals to extract meaningful features that distinguish different motor imagery patterns – such as imagined movements. This preprocessing step improves the classification accuracy of subsequent algorithms by highlighting the most discriminative aspects of brain activity patterns. This results in more reliable control of the exoskeleton through brain signals.

In the human detection system, *KD-tree-based Euclidean clustering* processes 3D point cloud data from sensors like time-of-flight (ToF) cameras to separate and identify distinct objects in the environment. By grouping points that likely belong to the same physical entity, the algorithm helps

AMRs distinguish between humans, objects, obstacles, and other environmental features. This segmentation is a crucial preprocessing step that transforms raw point cloud data into meaningful object representations that can be further classified and tracked [26].

Before switching to ToF cameras, a Depth-based clustering method was employed for LiDARs, in which the LiDAR 3D point cloud was segmented according to depth differences computed from LiDAR beam triangulation [72]. However, due to the high cost and fragility of LiDAR sensors, ToF cameras are preferred and employed instead.

The use case achieved 80% human presence detection accuracy within 5 meters, 85% gesture recognition accuracy across five gesture classes, and 75% 3D bounding box quality (IoU). The traditional ML techniques help more complex deep learning models (LSTM, BiLSTM, GRU, RNN, TCN) to operate effectively in the factory environment.

#### E. UC09: Energy-aware Scheduling of Batch Processing Machines in Wafer Fabs

The main goal of UC09 is to understand the energy flows of equipment in the semiconductor industry. A typical semiconductor factory needs as much energy as a small city does, because a lot of electrical power is used to maintain the clean room conditions, and for hundreds of equipment to produce the integrated circuits on the wafer. The teams already installed first sensor systems to measure the power consumption at some points within the factory. The main energy flows are known and first correlations have been found using historical and current weather data. A strong correlation between the power consumption, and the outside heat and humidity is evident from the observations.

Use case 9 employs several traditional machine learning algorithms. *Genetic Programming (GP)* is implemented to automatically discover efficient dispatching rules for energy-aware scheduling of batch processing machines in semiconductor manufacturing. The approach constructs tree structures representing priority indices that determine job scheduling order based on factors like job weights, ready times, and due dates. Initially developed for single tools, the system was progressively extended to handle tool groups and flexible flow shops with parallel identical batch processing machines. The framework incorporates time window decomposition to account for different job ready times and a decision theory heuristic to determine optimal idle time insertion in schedules.

The implementation uses the *Galib framework* with C++ programming and follows an iterative learning approach with training and validation phases. When tested on real-world problem instances from IFD's wafer fabrication facilities in Dresden, the GP approach demonstrated average reductions of 16% in total weighted tardiness and 11% in total electricity costs compared to benchmark approaches using only manual scheduling processes. Field tests with productive lots validated these improvements in the front-end wafer facility. This scheduling approach is applied on diffusion furnaces and other high-energy-consuming equipment, so as to balance traditional production metrics with sustainability considerations under time-of-use energy tariffs.

*Constraint Programming* was developed as a complementary approach to the genetic programming methods for energy-aware scheduling in semiconductor manufacturing. It was designed to schedule jobs in parallel machine environments while accounting for grid operator's load shifting requests. This approach represents a key innovation in the project's second year, enabling the semiconductor fabrication facility to respond dynamically to external energy grid conditions and requests.

The implementation forms part of the broader strategy to enhance the digital twin that accounts for energy consumption, demand, renewable generation and grid procurement. This constraint-based scheduling system helps predict the fab's energy consumption patterns and enables automated decision-making in response to load adjustment requests from grid operators. When load thresholds are requested by grid operators during periods of energy scarcity or grid stress, the constraint programming system can rapidly decide whether to follow these requested thresholds based on production constraints and optimization objectives, contributing to both production efficiency and grid stability. This capability represents a critical advancement in enabling semiconductor manufacturing facilities to participate in demand response programs while maintaining production schedules.

*K-Means Clustering* is deployed for anomaly detection in the complex cooling systems of semiconductor fabrication facilities. The method analyzes approximately four months of time series data collected from cooling equipment sensors, comprising hundreds of machine tags measuring parameters like temperatures and flow rates. For each sensor tag, feature vectors are computed over 5-minute intervals, incorporating statistical measures including mean, minimum, maximum values, variance, and the slope of the regression line, which are then clustered to identify deviations from normal patterns.

*Symbolic Aggregate Approximation (SAX)* represents an innovative approach to analyze time series data from semiconductor facility management systems. The technique transforms continuous sensor readings like temperature or flow data into symbolic representations by dividing the time series into segments and mapping each segment to a symbol based on its statistical characteristics. This dimensionality reduction technique creates new feature vectors that can be used for efficient clustering and anomaly detection in the complex cooling systems operating within the cleanroom environments.

At IFAT's Villach site, SAX was applied to explore the frequency and timing of specific critical events in cooling system operation, such as temperature increases exceeding 3 degrees within 15-minute intervals. By leveraging domain knowledge from facility management experts, the SAX-based analysis proved particularly effective at detecting human-defined anomalies that might otherwise be missed in the vast datasets generated by the monitoring systems. This application demonstrates how SAX provides a complementary approach to traditional statistical methods, offering new insights for predictive maintenance and optimization of energy-intensive cooling infrastructure essential to semiconductor manufacturing.

*Principal Component Analysis (PCA)* serves as a dimensionality reduction technique, enabling the visualization and

analysis of highly complex multidimensional data from semiconductor manufacturing environments. Applied specifically to facility management data comprising over 700 tags and approximately 20GB of information collected over 3.5 months from December 2023 to March 2024, PCA transforms this high-dimensional feature space into a comprehensible two-dimensional representation that highlights the most significant variations in the dataset. This visualization revealed that the smallest cluster, containing only 66 observations out of thousands, harbored the most significant anomalies - including instances of actual equipment damage that had occurred. The approach was integrated into an interactive browser-based tool, combining AI-supported anomaly detection with human expertise for optimal facility management. The PCA transformation made it possible to understand relationships between multiple parameters simultaneously, such as cold water valve positions, bypass control valves, and outlet temperatures, providing insights that would be impossible to detect through traditional monitoring approaches. The application demonstrates how targeted machine learning can effectively monitor performance degradation and identify potential failures in critical manufacturing infrastructure without requiring comprehensive sensor coverage.

*Time Series Analysis* forms the foundation of energy consumption prediction and pattern recognition in Use Case 9, enabling the correlation of energy usage with production parameters, facility loading, and environmental factors. Applied to data from smart power meters installed at strategic locations throughout IFD's 200mm and 300mm wafer fabrication facilities, these methods identified strong relationships between weather conditions and energy demand, revealing that higher humidity and temperature significantly increase energy consumption due to the requirement for dry air in cleanroom environments.

The implementation at IFD includes forecasting models that predict energy consumption 24 hours in advance by combining historical energy patterns with weather data and production schedules. These analyses showed clear weekly patterns and seasonal variations, with box plot visualizations demonstrating the distribution of energy consumption across the weeks of the year using data from 2018-2023. By understanding these temporal patterns, the facility management team can implement peak shaving strategies, adjusting production schedules to reduce energy consumption during high-cost periods or when electrical power shortages might occur. The insights gained from time series analysis directly inform the "golden tool strategy," which aims to optimize energy efficiency using minimal smart metering infrastructure. During high-energy peaks, the system demonstrated potential savings of more than 750 kWh per day through tool scheduling and load shifting strategies.

The use case also employs neural networks for energy forecasting, but traditional ML approaches form the major part of the overall machine learning strategy. The implementation uses Jupyter notebooks, TensorFlow, scikit-learn, Bokeh, Pandas, and HDF5, among others.

#### *F. UC11: AI-supported Industrial IoT for Indoor Food Production*

The aim of UC11 is to enhance plant product lifecycle automation. The proposed system utilizes sensors to monitor the conditions of each plant and analyzes the data in real-time with AI-based decision-making algorithms. The framework employs IoT devices, automation, and digital infrastructure in vertical farming systems. The wireless connectivity enhancement component developed by GUT employs *simulated annealing* to improve signal quality in indoor farming environments. This algorithm controls the Electronically Steerable Parasitic Array Radiator (ESPAR) antenna's beam switching capabilities to maximize signal strength and reduce interference in complex indoor environments - filled with dense equipment, plants, and other obstacles. The simulated annealing approach enables adaptive beam switching without requiring environmental information like signal source positions, allowing the antenna to dynamically adjust its radiation pattern from 4,096 possible configurations to maintain reliable connectivity even when in case of interference or potential jamming.

The ML optimization models for the ESPAR antenna's radiation pattern adaptation prioritize *binary classification* with adjustable thresholds, achieving approximately 400% improvement in Signal-to-Noise and Interference Ratio (SNIR) - a 6 dB gain over the omnidirectional baseline in approximately 20 iterations. This approach is implemented on TTTech's Nerve computing unit which serves as the central processing unit for wireless connectivity in the indoor farming setup. The classification model processes signal quality indicators from the wireless communication modules and determines the optimal antenna configuration. This methodology enables the system to quickly identify the best antenna pattern without exhaustively scanning each configuration, providing efficient and reliable wireless connectivity for real-time monitoring and control of farming parameters.

*Simultaneous Perturbation Stochastic Approximation (SPSA)* algorithms are being investigated as an alternative approach for minimizing loss functions in antenna steering vector optimization for the wireless communication system. These algorithms focus on optimizing the antenna's radiation pattern by approximating gradient information to find optimal configurations without exhaustively testing all possibilities. In the context of the use case UC11, the SPSA approach represents an established methodology that serves as a baseline for comparing the performance of newer ML-based approaches being developed. The team is evaluating whether these traditional algorithms can be enhanced or replaced by more advanced machine learning techniques to improve adaptation speed and reliability in the dynamic indoor farming environment.

#### *G. UC12: Human-Centered AI for the Optimization of Robust and Competitive Semiconductor Manufacturing Networks*

Use case 12 aims to optimize semiconductor production and decision-making through a more efficient and sustainable use of resources, transportation, and energy - lowering the CO<sub>2</sub>

emissions per product. The focus lies on the holistic visualization and realization of tool and process related machine state scenarios to increase transparency of the processes. The use case achieved a 33.8% increase in Overall Equipment Effectiveness (OEE) from a baseline of 42.9% (December 2023) to 57.4% (May 2025), exceeding the 10% target; potential to reduce cycle time by 23% through 50% reduction in repair variability; and a 8.3% difference in AI acceptance between experiment groups measured through weight-on-advice metrics.

*Extreme Gradient Boosting (XGBoost)* was implemented for work in progress (WIP) bubble prediction in semiconductor manufacturing. When training on one year of production data (2023) for the oven work center at IFD, XGBoost consistently proved to be the most effective model. The implementation uses machine, process, and product data (51 variables) to predict future WIP situations, helping to identify bottlenecks in advance and contributing directly to the achieved OEE improvements.

While *Shapley Additive Explanations (SHAP)* is not a predictive algorithm itself, it was implemented as an explainable AI approach to interpret the results of the traditional ML models. It was specifically used to analyze the feature importance of XGBoost models, increasing transparency and helping human decision-makers understand which variables most influence the WIP predictions. This implementation supports the human-centered aspect of the use case by making ML model outputs more interpretable. Experiments comparing black box versus transparent AI presentations showed that transparent AI achieved a weight-on-advice score of 0.58 compared to 0.63 for black box representations, demonstrating an 8.3% difference in how decision makers accept and trust AI predictions.

In the use case XGBoost emerged as the primary implemented solution based on its superior performance. Some other traditional ML algorithms have been considered or were in testing stages, but were not fully implemented: *Lasso regression*, *Ridge regression*, and *Elastic Net* regularization regression techniques were included in the planned algorithm testing suite – particularly for managing the high-dimensional feature space of the data –, but were not implemented eventually. *Support Vector regression (SVR)* and *Optimal Decision Tree* were also among the candidate algorithms. *Seasonal Autoregressive Integrated Moving Average with Exogenous Regressors (SARIMAX)* was considered as a comparison benchmark against XGBoost for time series prediction capabilities.

#### H. UC13: Intelligent Sensors improving Robustness of Automated Wafer Transportation and Storage Systems

*Isolation Forest* is used for anomaly detection in the audio signals collected from Overhead Transportation (OHT) shuttles in the workshop. The algorithm analyzes sound patterns from Micro-Electro-Mechanical Systems (MEMS) microphones installed along the rail system to detect abnormal acoustic signatures that might indicate mechanical issues before catastrophic failures occur. This approach enables predictive maintenance by identifying shuttles that deviate

from normal sound profiles. The implementation at Infineon Dresden leverages edge devices to process these audio files through feature extraction before applying the Isolation Forest algorithm, which effectively "isolates" anomalous sound patterns by partitioning the data space.

*Density-Based Spatial Clustering of Applications with Noise (DBSCAN)* is applied as part of the unsupervised anomaly detection strategy for monitoring the health status of automated material handling system components. This algorithm groups similar sensor readings together and identifies outliers in the acceleration data collected from the sensor platform housed in Front Opening Unified Pods (FOUPs). The density-based approach is particularly valuable in this use case as it doesn't require predefined knowledge about how many clusters exist in the data, making it suitable for detecting emerging failure patterns in the transportation systems where labeled failure data is scarce.

Traditional *spectral analysis* and *feature extraction* are implemented to process audio signals from the MEMS microphones monitoring the OHT shuttles, which operate on the rail system at Infineon Dresden. These methods transform time-domain sound data into frequency-domain representations that reveal characteristic patterns of both normal operation and potential mechanical failures. This enables the maintenance team to identify specific frequency signatures associated with bearing damage by comparing the frequency spectrum of good and bad bearings.

Traditional *time series analysis* techniques are applied to shuttle acceleration data from the multi-sensor platform (vibration, sound, and distance sensors). These algorithms process the data to detect shocks (peaks) and resonance frequencies that might indicate excessive forces on wafers or mechanical wear in the transportation system components. The implementation includes shock analysis algorithms that evaluate acceleration patterns to identify potentially damaging events during wafer transport, allowing for early intervention before components fail. This approach helped reduce maintenance efforts and saved significant downtime costs.

#### I. UC15: Productionized Machine Learning models at scale

Use Case 15 addresses the integration and deployment of scalable, interpretable, and production-ready AI/ML models for anomaly detection and predictive maintenance within semiconductor manufacturing processes. The use case is split to five complementary sub-activities for anomaly detection, defect classification, maintenance status classification, OEE prediction, and defect density image classification. Out of these, UC15-subA, Anomaly Detection and Virtual Inspection employs non-neural AI methods.

*Isolation Forest* is the core algorithm used in UC15-subA for unsupervised anomaly detection in semiconductor manufacturing. LFoundry chose this algorithm because it has been shown to be one of the best-performing anomaly detection approaches across many settings. It's well-suited for identifying anomalies in multivariate time-series sensor data to detect early-stage warnings of faults in sputtering deposition and metal etch processes.

*Depth-based Isolation Forest Feature Importance (DIFFI)* was implemented as an explainability layer on top of the Isolation Forest algorithm in UC15-subA. Given that operators need to take actions based on anomaly detection results, an approach was needed to make the black-box anomaly detection interpretable. DIFFI provides both local and global feature importance values, helping operators understand which sensor measurements contributed most to detecting an anomaly. This explainability component addresses the "human-in-the-loop" challenge by building trust in the ML solution by making its decisions transparent to users.

*Accelerated Model Explanations for Anomaly Detection (AcME-AD)* [51] was developed and implemented in UC15-subA as a model-agnostic explainability framework for anomaly detection in tabular data. It provides fast and interpretable insights into detected anomalies through local feature importance scores and a what-if analysis tool. This allows semiconductor manufacturing operators to understand the factors behind each detected anomaly without the computational overhead of more resource-intensive explainability techniques. This framework enabled a 55% reduction in time for monitoring industrial processes across four sputtering machines.

The *Balanced Anomaly and Latent Inference Framework (Extended B-ALIF)* algorithm extends the Isolation Forest model by incrementally incorporating expert feedback based on Bayesian Active Learning principles. It was implemented to improve detection performance in situations where labeled data is scarce, which is common in semiconductor manufacturing where fault examples are rare. The framework refines anomaly identification iteratively by leveraging limited human input, making it valuable in the Decision Support Systems developed for LFoundry, where tagged data for anomalies is limited.

#### J. UC16: AI enhanced production and performance of RF transceivers for civilian SATCOM and radar applications

The main aim of UC16 is to use AI solutions to improve the manufacturing throughput and sustainability of future radio frequency (RF) components and systems. AI-based methodologies are developed for design and facility optimization from radar applications through avionics to next generation high power satellite communication.

*Linear Regression* is utilized in the radio frequency (RF) component design and fabrication optimization processes. The teams from THALES and FORTH apply this algorithm to establish relationships between fabrication parameters and device performance characteristics. For example, during the manufacturing of RF Micro-Electro-Mechanical Systems (MEMS) switches and surface acoustic wave (SAW) temperature sensors, linear regression models help predict how changes in fabrication parameters (like dielectric film thickness or air gap spacing) affect performance metrics such as switching time and isolation. This approach helps understanding cause-effect relationships in the manufacturing process, where input features from the ISO quality control measurements can be directly correlated with output performance metrics.

*Decision Trees* are implemented as part of the supervised learning approach for quality control and fabrication process

optimization. Decision trees are used to classify potential defects or predict whether devices will meet specifications based on measurements taken during the ISO quality control process after each fabrication step. This is useful for identifying decision points in the manufacturing process that might lead to device failure or suboptimal performance. For example, when fabricating the mini-MEMS RF leakage sensors, decision trees help determine which combinations of process parameters are most likely to achieve the target specifications of 1 microsecond switching time and -10 dB isolation. Both algorithms form part of the use case's strategy to "utilize commonly adopted open-access algorithms" rather than developing new ones, with the focus being on proper feature and label selection during the training process.

#### K. UC17: Intelligent Contamination Management & Prescriptive Control Plans and Process Flows

Use case 17 primarily focuses on sensor networks, data collection infrastructure, and machine learning tools for contamination monitoring in semiconductor manufacturing facilities. The use case employs cost-effective static sensors utilizing meander structures, replacing expensive legacy systems. The sensors feed into a *data lake* architecture that consolidates data from multiple sources.

IFD designed a heuristic scheduling framework for flow shops based on *genetic programming*, applied to scheduling tasks related to Front Opening Unified Pod (FOUP) cleaning and monitoring processes. This helps optimize the sequence and timing of contamination control activities in the cleanroom environment, improving overall efficiency and contamination prevention.

Pfeiffer Vacuum improved its contamination control equipment with edge computing capabilities for real-time *anomaly detection*. The system analyzes equipment data (including measurement data, tool status, and analyzer status) every 5 seconds – a form of statistical anomaly detection that enables maintenance personnel to take immediate corrective action when issues arise, reducing downtime and preventing contamination events.

Fraunhofer IISB connected analytical instruments such as Inductively Coupled Plasma Mass Spectrometer (ICP-MS) with *machine learning models* for prediction and root cause analysis of contamination events. The implementation includes automated data transfer and ML-based heatmap visualization for metal contamination control measurements, reducing the need for time-intensive offline analysis.

While not always categorized as machine learning, *Statistical Process Control (SPC)* is being used in this use case for monitoring Airborne Molecular Contamination data from cleaned FOUPs. The documents show SPC data being visualized through tools like Tableau and Zeppelin. This statistical approach helps maintain control over the cleaning processes by identifying when measurements fall outside acceptable limits, allowing for process adjustments before contamination problems occur.

#### L. UC19: Extension of Digital Reference Ontology for Smart Sustainable Business Planning, Quality Traceability and Incident Handling

Use Case 19 addresses supply chain planning and quality management. One major driver of the recent semiconductor crisis is the bullwhip effect along OEM supply chains, which can be mitigated through closer connectivity and collaborative planning. IC miniaturization increases failure risks and complicates root cause analysis, wasting resources. This use case contributes to the elaboration of the Digital Reference Ontology (DRO) across three dimensions: traceability, supply chain management, and sustainability.

The *Random Forest* classification model was implemented for quality incident management in the semiconductor manufacturing process – specifically for the wire bonding manufacturing step. This model was used to differentiate between regular and problematic wire bond manufacturing conditions where problematic conditions were created through experiments that mimicked known root causes of failures. The method successfully identified three relevant machine-signal-related features important for classification that could be directly linked to potential IC failures. This approach enhances root cause analysis efficiency and enables proactive detection of failures in one of the most critical assembly steps in semiconductor manufacturing.

*Causal Rule Mining (CRM)* was applied to identify causal relationships between wire bonding and final test IC failures using *knowledge graph*-based traceability data. By adding cluster membership information from the knowledge graph, the team was able to conduct more sophisticated analysis of complex causal relationships that might otherwise remain hidden. CRM helps understand the cause-effect relationships in the manufacturing process, though it is less reliable with noisy or incomplete datasets.

*K-means Clustering* was used to validate observations from the machine signal data analysis. For instance, when analyzing the three relevant machine-signal features discovered by the Random Forest model, k-nearest neighbor validation achieved an accuracy of 0.738 (0.734; 0.742), confirming that clusters of different experimental conditions could be distinguished. However, k-means clustering is limited when dealing with high-dimensional features and imbalanced datasets, which are common in quality incident management scenarios.

For supply chain optimization, *two-stage stochastic programming (2SP)* and *chance-constrained programming (CCP)* models were implemented to handle demand uncertainty more robustly than traditional linear programming approaches. These stochastic optimization models showed superior performance under demand variability while balancing trade-offs among service level, inventory position, and planning stability.

#### M. UC20: Harmonization of Automated Digitized Equipment Control (AuDi-EC)

Use case 20 addresses the lack of a harmonized equipment test strategy for two 300 mm high volume factories, which should work as a "One virtual Fab". Test wafers are used for the permanent check of the tool status to assure the quality

demands of the customer. The ratio of test wafers compared to production wafers is higher and an important cost driver. The main goal is to balance the conflicting goals between quality and costs. Highest priority, especially for health and automotive customers, is to fulfil the quality demands. On the other hand, the processes must be lean and efficient to reach the cost targets. The test processes are very complex even in one of the existing factories. The goal of this use case is to harmonize different processes, tools and software to get a single monitoring strategy for the two different 300 mm production plants in Villach and Dresden.

*Isolation Forests* are used in the anomaly detection component of the test wafer strategy. In semiconductor manufacturing, detecting anomalies in equipment behavior is crucial for maintaining quality standards. Within Use Case 20, Isolation Forests help identify unusual patterns in test wafer data across both the Villach and Dresden facilities, supporting the harmonization efforts by detecting when equipment monitoring results deviate from normal parameters.

*Density-Based Spatial Clustering of Applications with Noise (DBSCAN)* is implemented for unsupervised anomaly detection in the test wafer monitoring process. In UC20, DBSCAN helps process engineers identify clusters of similar equipment behavior across the two manufacturing sites and detect outliers that might indicate equipment issues requiring intervention. This supports the OneVirtualFab concept by providing a consistent approach to anomaly detection across both facilities, contributing to the harmonization of equipment control processes and optimizing the use of test wafers.

## IV. CHALLENGES, RESEARCH GAPS AND FUTURE DIRECTIONS

The use cases presented in Section III demonstrate that traditional machine learning remains highly relevant to industrial applications. It offers practical solutions when computational efficiency, limited data, and interpretability are important, particularly when compatibility with existing industrial systems must also be ensured. At the same time, the use cases make clear that successful industrial ML depends on far more than predictive performance [56]. Persistent challenges include data quality, generalization under changing conditions, and the maintenance of models after deployment. Equally important is the ability to incorporate human expertise and translate model outputs into reliable operational decisions.

Consequently, the transition from Industry 4.0 toward Industry 5.0 [1], [8] calls for a shift from an algorithm-centric view of ML to one that considers the entire system and its lifecycle. The central question is no longer simply whether a model achieves sufficient accuracy on a given dataset. Rather, it is whether that performance can be sustained, understood, validated, and translated into action throughout the lifetime of an industrial system.

### A. Data Availability, Quality, and Heterogeneity

Industrial ML often operates in a data regime that differs fundamentally from that of benchmark applications. Data from normal operation are usually plentiful, whereas the events

of greatest interest – such as failures, quality deviations, or contamination events – occur only rarely. As a result, class imbalance and a shortage of labeled fault examples are common challenges.

This issue appears in several AIMS5.0 use cases. In UC05, labeled tool-breakage data could be generated under laboratory conditions, but equivalent labels were not readily available in real production. UC13 relies substantially on unsupervised anomaly detection, while UC15 combines Isolation Forest with expert feedback to progressively refine the identification of relevant anomalies.

The usefulness of an industrial dataset cannot be judged by its size alone. Sensor calibration, missing measurements, and sampling rates may be just as important, as may the relationship between the recorded measurements and the underlying physical events. Data collected from different machines or factories can also vary in format, resolution, and semantics. These differences make it difficult to reuse models directly across production environments.

A major research gap is therefore the development of *data-efficient industrial learning* methods that can exploit unlabeled process data while requiring only limited expert annotation. Active, semi-supervised, and incremental learning are particularly relevant in this context. Future research should also place greater emphasis on semantic interoperability. Ontologies and standardized asset descriptions, such as those of the ISO 23247 digital twin framework [73], can help distinguish genuine changes in the process from differences introduced by sensing and information infrastructures.

### B. Generalization, Concept Drift, and Model Adaptation

A model that performs well on historical data or on a single machine cannot be expected to perform equally well indefinitely after deployment. Tools wear, materials and recipes change, and sensors may be replaced or recalibrated. Such changes can alter the distribution of the input data or the relationship between inputs and the target variable, a phenomenon known as concept drift [74].

This limitation is visible in UC05, where generalization across tools, cutting conditions, and part types remains challenging despite high accuracy under laboratory conditions. UC03 highlights a related issue, as differences among edge devices constrain uniform deployment. UC20 addresses the problem at a broader scale by considering how monitoring strategies can be harmonized between two fabrication sites.

Future systems therefore require mechanisms for recognizing when current operating conditions have moved beyond those represented in the training data. Monitoring should track predictive performance whenever labels eventually become available, but it should also examine shifts in feature distributions, anomaly scores, and residuals. Once a meaningful change is detected, the system must determine whether recalibration, incremental updating, or complete retraining is the most appropriate response.

Incremental learning provides a promising mechanism for this purpose [38]. However, unrestricted online adaptation brings its own risks, because a model may learn from erroneous measurements, transient disturbances, or incorrect

labels. An important research direction is therefore *controlled adaptation*, in which drift detection, model updating, validation, and human approval are treated as parts of a single lifecycle. At the same time, every operational decision should remain traceable to the model version and the data on which it was based.

### C. From Prototypes to Maintainable and Measurable Systems

Another important gap lies in the transition from successful experiments to sustainable production deployment. Algorithmic performance is only one component of a production ML system. In practice, the surrounding data pipelines, integration mechanisms, and monitoring infrastructure can dominate long-term costs [75]. Machine learning operations (MLOps) practices address this broader lifecycle, but their adaptation to shop-floor and edge environments remains limited [76].

Edge deployments such as UC03 illustrate how memory, processing capacity, and energy consumption can restrict the models that are practical to deploy. UC07 reveals another side of the same challenge. Automated annotation substantially improved the computer-vision pipeline, yet the hardware and labeling effort required across numerous product variants made full deployment economically unattractive. An ML method may therefore be technically effective without being operationally scalable. Such negative results are scientifically valuable and deserve to be reported more systematically.

The use cases also report a diverse range of outcomes, from classification accuracy to energy savings, OEE improvements, and indicators of human acceptance. This diversity makes direct comparison difficult. Future studies should therefore distinguish between three levels of evaluation. *Model-level* metrics describe predictive or detection performance. *System-level* metrics capture factors such as latency, resource consumption, and maintenance effort. Finally, *industrial-impact* metrics measure effects on cost, downtime, quality, and human decision making. Temporal, cross-machine, and cross-site validation should complement random hold-out tests, which primarily demonstrate that a model can interpolate within the distribution represented by its training data.

These considerations support a broader move toward *resource-aware and frugal industrial ML*. In some applications, a slightly less accurate decision tree or linear model may create more industrial value than a complex alternative. If it can be deployed, inspected, updated, and maintained at substantially lower cost, the simpler model may ultimately be the more practical solution.

### D. Explainability, Uncertainty, and Human Oversight

Interpretability becomes especially important when ML outputs influence maintenance, production planning, or human work. Traditional models are often comparatively easy to inspect, whereas post-hoc techniques such as SHAP, DIFFI, and AcME-AD can provide explanations for more complex predictors and anomaly detectors [6], [77]. In the pipeline of Fig. 2, the explainer therefore serves as the link between the learner and the operator.

The AIMS5.0 use cases indicate, however, that simply providing an explanation does not guarantee effective human–AI cooperation. In UC12, decision makers relied less on the transparent presentation of the model than on the black-box version, with a weight-on-advice of 0.58 versus 0.63. UC15 illustrates a different situation, in which anomaly explanations become valuable because they help operators identify the sensor variables responsible for an alert and investigate the underlying process.

This highlights an important gap in current XAI evaluation. Explanations are often assessed as properties of the model interface, whereas Industry 5.0 requires their influence on the *human decision process* to be evaluated [78]. Future studies should examine whether an explanation improves diagnostic accuracy, reduces investigation time, or helps users recognize incorrect predictions. The desired outcome is appropriate reliance: operators should follow the model when its recommendation is sound and override it when it is not [79].

Uncertainty should likewise become an explicit component of decision support. A model should be able to indicate when the available data are insufficient or differ substantially from its training experience. Low-confidence cases can then be escalated to human experts rather than forcing an automated decision. Explainability, uncertainty estimation, and abstention should therefore be studied together. From a governance perspective, traceability, human oversight, and robustness are also becoming increasingly important regulatory requirements [57]. The research challenge is to embed these properties technically within ML pipelines instead of treating compliance as documentation added after development.

#### E. From Prediction to Multi-Objective Decision Support

Many industrial ML applications ultimately revolve around decisions rather than predictions (cf. the *decide* branch of Fig. 1). Predicting a machine failure, energy demand, or WIP level has limited value unless that prediction can be translated into an appropriate operational response.

Several use cases demonstrate that such decisions involve competing objectives. UC09 must balance production performance against electricity cost. UC19 considers service level alongside inventory and planning stability, while UC20 weighs quality assurance against the cost of test wafers. Similar tensions can arise between productivity and other concerns, such as worker safety or equipment lifetime.

Predictive accuracy, however, is typically represented by a single statistical metric. Future work should therefore connect ML more systematically with multi-objective, stochastic, and constraint-based optimization. Predictions should not, in particular, be treated as exact inputs to an optimizer. Their uncertainty should propagate into the decision layer, and models may even be trained directly to improve decision quality rather than to minimize prediction error [80].

#### F. Hybrid Models and ML-Enhanced Digital Twins

Digital twins provide a controlled environment in which alternative decisions can be evaluated before they are applied to the physical process [66], [67]. ML can complement these

environments in several ways. It can estimate parameters that cannot be measured directly, construct fast surrogate models, and detect discrepancies between expected and observed behavior. Within AIMS5.0, UC07 uses discrete-event simulation to evaluate scheduling strategies, while UC09 integrates energy forecasting and scheduling into a digital representation of fab operations.

A particularly relevant research direction is the development of *hybrid industrial models* that combine first-principles knowledge with data-driven learning [81]. Purely data-driven models may struggle when confronted with abnormal conditions they have never observed, whereas purely physical models can be computationally expensive or incomplete. Hybrid approaches offer a middle ground: they can reduce training-data requirements while preserving a physically meaningful structure. Digital twins can also provide safe test beds for reinforcement learning and adaptive scheduling [82]. This potential, however, depends on continuously validating that the twin remains representative of the physical process it is intended to mirror.

#### G. Cross-Site Learning and Data Sovereignty

Many industrial organizations operate fleets of similar machines or maintain multiple production sites. Learning jointly from these environments could improve model robustness, particularly when failures are rare. Yet industrial data cannot always be centralized because of intellectual-property concerns, confidentiality requirements, or organizational boundaries. Federated learning offers a way to train models without transferring raw data [83], although its applicability to strongly heterogeneous industrial sites remains an open research question. For traditional ML, simpler approaches to knowledge reuse may also be valuable. Engineered features, anomaly thresholds, model parameters, or domain rules can potentially be transferred without moving complete datasets. UC20's OneVirtualFab concept illustrates the strategic importance of this type of cross-site harmonization.

#### H. Research Agenda

Based on the challenges discussed above, the following questions define a practical research agenda for machine learning in Industry 5.0:

- 1) **Data-efficient learning:** How can useful models be trained when failures are rare and labels are expensive?
- 2) **Transfer and adaptation:** How can models remain reliable across machines, sites, and changing conditions without complete retraining?
- 3) **Lifecycle-aware ML:** How should drift detection, re-training, and versioning be integrated into maintainable industrial pipelines?
- 4) **Human–ML cooperation:** Which explanations and uncertainty cues help operators recognize when a recommendation should not be followed?
- 5) **Decision-oriented ML:** How can prediction uncertainty be incorporated into scheduling and optimization?

- 6) **Hybrid modeling:** How can physical knowledge and digital twins be combined with data-driven models to improve generalization?
- 7) **Resource-aware ML:** How should model complexity be balanced against latency, energy, and maintainability?
- 8) **Comparable validation:** Which evaluation protocols demonstrate benefits beyond a single dataset or pilot installation?

Overall, the AIMS5.0 use cases suggest that the next stage of industrial ML will depend less on developing ever more complex algorithms and more on building learning systems that are *adaptable, maintainable, interpretable, and operationally measurable*. Traditional ML is well positioned to contribute to this transition, particularly when it is combined with domain knowledge and human expertise. The emphasis therefore shifts from isolated model performance to dependable performance across the full industrial lifecycle.

## V. CONCLUSIONS

This paper examined the role of machine learning in the transition from Industry 4.0 to Industry 5.0. Particular attention was given to the continued relevance of traditional ML methods in industrial environments. Industry 4.0 established the foundations for connected, data-driven, and increasingly automated production. Industry 5.0 builds on these foundations, but shifts the focus toward how technology can support people, sustainability, and resilience.

The paper mapped the traditional ML methods used in 13 AIMS5.0 use cases (Table I), explained the concepts behind each method family, and condensed the observed selection logic and human-in-the-loop structure into Figs. 1 and 2. Section ?? derived a research agenda from the recurring difficulties.

The reviewed literature and the AIMS5.0 use cases show that established ML techniques remain highly relevant. These include regression methods, decision trees, random forests, support vector machines, clustering, anomaly detection, and time-series models. Their strengths become especially clear when industrial data are structured, datasets are limited, or computing resources are constrained. They are also attractive when model behaviour needs to be understandable. In such situations, a more complex model is not automatically a better industrial solution. Table II summarized the abbreviations used throughout the paper.

The use cases also show that ML rarely creates value as an isolated predictive component. Its outputs are most useful when they become part of a broader operational workflow. ML can support predictive maintenance, quality monitoring, production planning, energy management, anomaly detection, and process improvement. In practice, these models are often combined with optimization methods, simulation tools, expert knowledge, or edge-computing platforms. This creates a layered decision process in which different technologies play different roles. ML extracts useful information from data. Other components determine how that information is translated into action.

One of the clearest lessons from the investigated cases is that predictive performance alone is not enough. A model

may achieve excellent accuracy and still be difficult to deploy. Industrial environments introduce challenges that are often absent from benchmark datasets. Labels may be scarce, operating conditions may change, and machines may behave differently across sites. Integration and maintenance can also be costly. These factors can determine whether a technically successful prototype survives beyond the pilot stage.

For this reason, industrial ML should be evaluated more broadly. Accuracy, precision, recall, or similar model-level metrics remain important, but they tell only part of the story. A practical ML system must also be robust, scalable, and maintainable. Its computational demands must fit the deployment environment. Most importantly, it should produce measurable operational value.

The AIMS5.0 use cases reinforce this point. Some applications demonstrated strong technical performance but also revealed economic or deployment constraints. Such outcomes are valuable rather than disappointing. They expose the distance between a successful experiment and a sustainable industrial solution. They also show why large-scale adoption cannot be justified by predictive accuracy alone.

Human involvement is equally important. Industry 5.0 does not simply replace human expertise with automated decisions. Instead, ML should strengthen the ability of operators, engineers, and planners to understand what is happening and decide what to do next. Explainability can support this process, but explanation alone is not enough. More transparency does not automatically lead to greater trust, nor does it guarantee better decisions.

A more useful goal is therefore appropriate human reliance. ML systems should help users understand when a recommendation is well supported and when caution is needed. They should also make uncertainty visible where possible. In high-impact settings, experts should be able to question, validate, and, when necessary, override model outputs. This is where human-machine cooperation becomes more than a design principle: it becomes part of system reliability.

Looking ahead, the future of ML in Industry 5.0 is likely to depend less on increasingly complex standalone models and more on dependable learning systems. These systems will need to work with limited data, adapt to changing conditions, and operate across machines and sites. They will also need to fit within existing industrial infrastructures rather than around them.

Several research directions are especially important. Continual and incremental learning can help models adapt over time. Hybrid approaches can combine data-driven learning with physical knowledge and simulation. Resource-aware methods can support efficient deployment at the edge. Cross-site learning can improve knowledge transfer between factories. Closer integration with optimization can also help convert predictions into more robust decisions.

Traditional machine-learning methods are therefore likely to remain an important part of the Industry 5.0 landscape. Their relative simplicity is not a weakness. In many industrial applications, it is an advantage. These methods can be efficient, interpretable, and well suited to structured process data. They

can also be easier to validate and maintain than more complex alternatives.

The key challenge is no longer simply to build models that perform well in controlled experiments. The real test is whether those models continue to create value once they enter the factory floor. For Industry 5.0, successful ML will be defined not only by what it can predict, but by how reliably it can support people, processes, and decisions under real operating conditions.

#### ACKNOWLEDGMENT

The research leading to these results was conducted in the AIMS5.0 project, which received funding from the Chips Joint Undertaking under grant agreement No. 101112089 and from the partners' national funding programmes.

#### REFERENCES

- [1] A. Renda, S. Schwaag Serger, D. Tataj, A. Morlet, D. Isaksson, F. Martins, M. Mir Roca, C. Hidalgo, A. Huang, S. Dixon-Declève, P.-A. Bolland, F. Bria, C. Charveriat, K. Dunlop, and E. Giovannini, "Industry 5.0, a transformative vision for Europe: Governing systemic transformations towards a sustainable industry," European Commission, Directorate-General for Research and Innovation, Luxembourg, ESIR Policy Brief 3, 2022, catalogue number KI-BE-21-005-EN-N. [Online]. Available: [https://research-and-innovation.ec.europa.eu/knowledge-publications-tools-and-data/publications/all-publications/industry-50-transformative-vision-europe\\_en](https://research-and-innovation.ec.europa.eu/knowledge-publications-tools-and-data/publications/all-publications/industry-50-transformative-vision-europe_en)
- [2] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. New York, NY, USA: Springer, 2009.
- [3] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [4] T. Wuest, D. Weimer, C. Irgens, and K.-D. Thoben, "Machine learning in manufacturing: advantages, challenges, and applications," *Production & Manufacturing Research*, vol. 4, no. 1, pp. 23–45, 2016.
- [5] A. Presciuttini, A. Cantini, F. Costa, and A. Portioli-Staudacher, "Machine learning applications on iot data in manufacturing operations and their interpretability implications: A systematic literature review," *Journal of Manufacturing Systems*, vol. 74, pp. 477–486, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0278612524000773>
- [6] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable ai: A review of machine learning interpretability methods," *Entropy*, vol. 23, no. 1, 2021. [Online]. Available: <https://www.mdpi.com/1099-4300/23/1/18>
- [7] Z. Ge, Z. Song, S. X. Ding, and B. Huang, "Data mining and analytics in the process industry: The role of machine learning," *IEEE Access*, vol. 5, pp. 20 590–20 616, 2017.
- [8] X. Xu, Y. Lu, B. Vogel-Heuser, and L. Wang, "Industry 4.0 and industry 5.0— inception, conception and perception," *Journal of Manufacturing Systems*, vol. 61, pp. 530–535, 2021.
- [9] T. P. Carvalho, F. A. A. M. N. Soares, R. Vita, R. d. P. Francisco, J. P. Basto, and S. G. S. Alcalá, "A systematic literature review of machine learning methods applied to predictive maintenance," *Computers & Industrial Engineering*, vol. 137, p. 106 024, 2019.
- [10] A. Widodo and B.-S. Yang, "Support vector machine in machine condition monitoring and fault diagnosis," *Mechanical Systems and Signal Processing*, vol. 21, no. 6, pp. 2560–2574, 2007.
- [11] A. K. Jain, "Data clustering: 50 years beyond K-means," *Pattern Recognition Letters*, vol. 31, no. 8, pp. 651–666, 2010.
- [12] S. J. Qin, "Survey on data-driven industrial process monitoring and diagnosis," *Annual Reviews in Control*, vol. 36, no. 2, pp. 220–234, 2012.
- [13] J. E. van Engelen and H. H. Hoos, "A survey on semi-supervised learning," *Machine Learning*, vol. 109, no. 2, pp. 373–440, 2020.
- [14] B. Settles, "Active learning literature survey," University of Wisconsin–Madison, Department of Computer Sciences, Computer Sciences Technical Report 1648, 2009.
- [15] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.

TABLE II  
ABBREVIATIONS USED THROUGHOUT THE PAPER.

| Abbreviation | Description                                                                         |
|--------------|-------------------------------------------------------------------------------------|
| 2SP          | Two-Stage Stochastic Programming                                                    |
| AcME-AD      | Accelerated Model Explanations for Anomaly Detection                                |
| AI           | Artificial Intelligence                                                             |
| AIMS5.0      | Artificial Intelligence in Manufacturing leading to Sustainability and Industry 5.0 |
| AMR          | Autonomous Mobile Robot                                                             |
| AuDi-EC      | Automated Digitized Equipment Control                                               |
| B-ALIF       | Bayesian Active Learning Isolation Forest                                           |
| CART         | Classification and Regression Trees                                                 |
| CCP          | Chance-Constrained Programming                                                      |
| CP           | Constraint Programming                                                              |
| CPU          | Central Processing Unit                                                             |
| CRM          | Causal Rule Mining                                                                  |
| CSP          | Common Spatial Pattern                                                              |
| CV           | Computer Vision                                                                     |
| DBSCAN       | Density-Based Spatial Clustering of Applications with Noise                         |
| DES          | Discrete-Event Simulation                                                           |
| DIFFI        | Depth-based Isolation Forest Feature Importance                                     |
| DNN          | Deep Neural Network                                                                 |
| DRO          | Digital Reference Ontology                                                          |
| DT           | Digital Twin                                                                        |
| EEG          | Electroencephalography                                                              |
| ESPAR        | Electronically Steerable Parasitic Array Radiator                                   |
| EU           | European Union                                                                      |
| FFT          | Fast Fourier Transform                                                              |
| FOUP         | Front Opening Unified Pod                                                           |
| GBDT         | Gradient-Boosted Decision Tree                                                      |
| GP           | Genetic Programming                                                                 |
| HDF5         | Hierarchical Data Format Version 5                                                  |
| IC           | Integrated Circuit                                                                  |
| ICP-MS       | Inductively Coupled Plasma Mass Spectrometry                                        |
| IoT          | Internet of Things                                                                  |
| IoU          | Intersection over Union                                                             |
| ISO          | International Organization for Standardization                                      |
| KD-tree      | $k$ -Dimensional Tree                                                               |
| KNN          | $K$ -Nearest Neighbors                                                              |
| KPI          | Key Performance Indicator                                                           |
| LiDAR        | Light Detection and Ranging                                                         |
| MDP          | Markov Decision Process                                                             |
| MEMS         | Micro-Electro-Mechanical Systems                                                    |
| ML           | Machine Learning                                                                    |
| NC           | Numerical Control                                                                   |
| OEE          | Overall Equipment Effectiveness                                                     |
| OEM          | Original Equipment Manufacturer                                                     |
| OHT          | Overhead Transportation                                                             |
| PCA          | Principal Component Analysis                                                        |
| RF           | Radio Frequency                                                                     |
| RL           | Reinforcement Learning                                                              |
| ROC          | Receiver Operating Characteristic                                                   |
| SARIMAX      | Seasonal Autoregressive Integrated Moving Average with Exogenous Regressors         |
| SATCOM       | Satellite Communications                                                            |
| SAW          | Surface Acoustic Wave                                                               |
| SAX          | Symbolic Aggregate Approximation                                                    |
| SHAP         | Shapley Additive Explanations                                                       |
| SLA          | Service Level Agreement                                                             |
| SNIR         | Signal-to-Noise and Interference Ratio                                              |
| SPC          | Statistical Process Control                                                         |
| SPSA         | Simultaneous Perturbation Stochastic Approximation                                  |
| SVM          | Support Vector Machine                                                              |
| SVR          | Support Vector Regression                                                           |
| TinyML       | Tiny Machine Learning                                                               |
| ToF          | Time of Flight                                                                      |
| TRL          | Technology Readiness Level                                                          |
| UC           | Use Case                                                                            |
| WIP          | Work in Progress                                                                    |
| XAI          | Explainable Artificial Intelligence                                                 |
| XGBoost      | Extreme Gradient Boosting                                                           |

- [16] M. Panzer and B. Bender, "Deep reinforcement learning in production systems: a systematic literature review," *International Journal of Production Research*, vol. 60, no. 13, pp. 4316–4341, 2022.
- [17] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computing*, 2nd ed. Berlin, Germany: Springer, 2015.
- [18] J. Branke, S. Nguyen, C. W. Pickardt, and M. Zhang, "Automated design of production scheduling heuristics: A review," *IEEE Transactions on Evolutionary Computation*, vol. 20, no. 1, pp. 110–124, 2016.
- [19] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Belmont, CA, USA: Wadsworth, 1984.
- [20] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [21] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *The Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.
- [22] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD)*, 2016, pp. 785–794.
- [23] T. M. Cover and P. E. Hart, "Nearest neighbor pattern classification," *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.
- [24] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [25] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proc. 2nd Int. Conf. Knowledge Discovery and Data Mining (KDD)*, 1996, pp. 226–231.
- [26] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (pcl)," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 1–4.
- [27] I. T. Jolliffe and J. Cadima, "Principal component analysis: a review and recent developments," *Philosophical Transactions of the Royal Society A*, vol. 374, no. 2065, p. 20 150 202, 2016.
- [28] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 15:1–15:58, 2009.
- [29] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *Proc. 8th IEEE Int. Conf. Data Mining (ICDM)*, 2008, pp. 413–422.
- [30] D. Sartor, T. Barbariol, and G. A. Susto, "Bayesian active learning isolation forest (B-ALIF): A weakly supervised strategy for anomaly detection," *Engineering Applications of Artificial Intelligence*, vol. 130, p. 107 671, 2024.
- [31] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA, USA: MIT Press, 1992.
- [32] F. Rossi, P. van Beek, and T. Walsh, Eds., *Handbook of Constraint Programming*, ser. Foundations of Artificial Intelligence. Amsterdam, The Netherlands: Elsevier, 2006.
- [33] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, 2nd ed. New York, NY, USA: Springer, 2011.
- [34] A. Charnes and W. W. Cooper, "Chance-constrained programming," *Management Science*, vol. 6, no. 1, pp. 73–79, 1959.
- [35] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [36] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Transactions on Automatic Control*, vol. 37, no. 3, pp. 332–341, 1992.
- [37] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: Wiley, 1994.
- [38] V. Losing, B. Hammer, and H. Wersing, "Incremental on-line learning: A review and comparison of state of the art algorithms," *Neurocomputing*, vol. 275, pp. 1261–1274, 2018.
- [39] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed. Hoboken, NJ, USA: Wiley, 2015.
- [40] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.
- [41] R. B. Randall, *Vibration-based Condition Monitoring: Industrial, Aerospace and Automotive Applications*. Chichester, UK: Wiley, 2011.
- [42] C.-C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, and E. Keogh, "Matrix profile I: All pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets," in *Proc. 16th IEEE Int. Conf. Data Mining (ICDM)*, 2016, pp. 1317–1322.
- [43] C. Truong, L. Oudre, and N. Vayatis, "Selective review of offline change point detection methods," *Signal Processing*, vol. 167, p. 107 299, 2020.
- [44] J. Lin, E. Keogh, L. Wei, and S. Lonardi, "Experiencing SAX: a novel symbolic representation of time series," *Data Mining and Knowledge Discovery*, vol. 15, no. 2, pp. 107–144, 2007.
- [45] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [46] H. Ramoser, J. Müller-Gerking, and G. Pfurtscheller, "Optimal spatial filtering of single trial EEG during imagined hand movement," *IEEE Transactions on Rehabilitation Engineering*, vol. 8, no. 4, pp. 441–446, 2000.
- [47] B. Blankertz, R. Tomioka, S. Lemm, M. Kawanabe, and K.-R. Müller, "Optimizing spatial filters for robust EEG single-trial analysis," *IEEE Signal Processing Magazine*, vol. 25, no. 1, pp. 41–56, 2008.
- [48] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30 (NeurIPS)*, 2017, pp. 4765–4774.
- [49] M. Carletti, M. Terzi, and G. A. Susto, "Interpretable anomaly detection with DIFFI: Depth-based feature importance of isolation forest," *Engineering Applications of Artificial Intelligence*, vol. 119, p. 105 730, 2023.
- [50] D. Dandolo, C. Masiero, M. Carletti, D. Dalle Pezze, and G. A. Susto, "AcME—accelerated model-agnostic explanations: Fast whitening of the machine-learning black box," *Expert Systems with Applications*, vol. 214, p. 119 115, 2023.
- [51] V. Zaccaria, D. Dandolo, C. Masiero, and G. A. Susto, "Acmead: Accelerated model explanations for anomaly detection," 2024. [Online]. Available: <https://arxiv.org/abs/2403.01245>
- [52] J. Li, T. D. Le, L. Liu, J. Liu, Z. Jin, B. Sun, and S. Ma, "From observational studies to causal rule mining," *ACM Transactions on Intelligent Systems and Technology*, vol. 7, no. 2, pp. 14:1–14:27, 2015.
- [53] D. C. Montgomery, *Introduction to Statistical Quality Control*, 8th ed. Hoboken, NJ, USA: Wiley, 2019.
- [54] A. M. Law, *Simulation Modeling and Analysis*, 5th ed. New York, NY, USA: McGraw-Hill, 2015.
- [55] O. Mypati, A. Mukherjee, D. Mishra, S. K. Pal, P. P. Chakrabarti, and A. Pal, "A critical review on applications of artificial intelligence in manufacturing," *Artificial Intelligence Review*, vol. 56, no. 1, pp. 661–768, Oct 2023. [Online]. Available: <https://doi.org/10.1007/s10462-023-10535-y>
- [56] J. Leng, X. Zhu, Z. Huang, X. Li, P. Zheng, X. Zhou, D. Mourtzis, B. Wang, Q. Qi, H. Shao, J. Wan, X. Chen, L. Wang, and Q. Liu, "Unlocking the power of industrial artificial intelligence towards industry 5.0: Insights, pathways, and challenges," *Journal of Manufacturing Systems*, vol. 73, pp. 349–363, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0278612524000335>
- [57] European Data Protection Supervisor, *AI Act Regulation (EU) 2024/1689 – Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance)*. Publications Office of the European Union, 2025.
- [58] T. Czimmernmann, G. Ciuti, M. Milazzo, M. Chiurazzi, S. Roccella, C. M. Oddo, and P. Dario, "Visual-based defect detection and classification approaches for industrial applications—a SURVEY," *Sensors*, vol. 20, no. 5, p. 1459, 2020.
- [59] F. Semeraro, A. Griffiths, and A. Cangelosi, "Human–robot collaboration and machine learning: A systematic review of recent research," *Robotics and Computer-Integrated Manufacturing*, vol. 79, p. 102 432, 2023.
- [60] D. Weichert, P. Link, A. Stoll, S. Rüping, S. Ihlenfeldt, and S. Wrobel, "A review of machine learning for the optimization of production processes," *The International Journal of Advanced Manufacturing Technology*, vol. 104, pp. 1889–1902, 2019.
- [61] D. Ni, Z. Xiao, and M. K. Lim, "A systematic review of the research trends of machine learning in supply chain management," *International Journal of Machine Learning and Cybernetics*, vol. 11, pp. 1463–1482, 2020.
- [62] G. Baryannis, S. Validi, S. Dani, and G. Antoniou, "Supply chain risk management and artificial intelligence: state of the art and future research directions," *International Journal of Production Research*, vol. 57, no. 7, pp. 2179–2202, 2019.

- [63] D. A. C. Narciso and F. G. Martins, "Application of machine learning tools for energy efficiency in industry: A review," *Energy Reports*, vol. 6, pp. 1181–1199, 2020.
- [64] A. Jamwal, R. Agrawal, and M. Sharma, "Machine learning applications for sustainable manufacturing: a bibliometric-based review for future research," *Journal of Enterprise Information Management*, vol. 35, no. 2, pp. 566–596, 2022.
- [65] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 297–330, 2020.
- [66] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 4, pp. 2405–2415, 2019.
- [67] M. M. Rathore, S. A. Shah, D. Shukla, E. Bentafat, and S. Bakiras, "The role of AI, machine learning, and big data in digital twinning: A systematic literature review, challenges, and opportunities," *IEEE Access*, vol. 9, pp. 32 030–32 052, 2021.
- [68] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [69] European Commission. (2023) Artificial intelligence in manufacturing leading to sustainability and Industry5.0 (AIMS5.0), grant agreement no. 101112089. EU Research Results. HORIZON-KDT-JU-2022-1-IA. Accessed: Sep. 29, 2026. [Online]. Available: <https://aims50.eu/>
- [70] Stream Analyze AB. (2026) Stream analyze platform: The end-to-end edge AI platform. Stream Analyze AB. Uppsala, Sweden. Accessed: Aug. 29, 2026. [Online]. Available: <https://www.streamanalyze.com/>
- [71] P. P. Ray, "A review on TinyML: State-of-the-art and prospects," *Journal of King Saud University – Computer and Information Sciences*, vol. 34, no. 4, pp. 1595–1623, 2022.
- [72] L. Nunes, R. Marcuzzi, J. Behley, and C. Stachniss, "Towards generating realistic 3d semantic training data for autonomous driving," 2026. [Online]. Available: <https://arxiv.org/abs/2503.21449>
- [73] G. Shao et al., "An analysis of the new ISO 23247 series of standards on digital twin framework for manufacturing," *NIST, Tech. Rep.*, 2023. [Online]. Available: <https://www.nist.gov/publications/analysis-new-iso-23247-series-standards-digital-twin-framework-manufacturing>
- [74] J. Gama, I. Žliobait'e, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 44:1–44:37, 2014.
- [75] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, and D. Dennison, "Hidden technical debt in machine learning systems," in *Advances in Neural Information Processing Systems 28 (NIPS)*, 2015, pp. 2503–2511.
- [76] D. Kreuzberger, N. Kühl, and S. Hirschl, "Machine learning operations (MLOps): Overview, definition, and architecture," *IEEE Access*, vol. 11, pp. 31 866–31 879, 2023.
- [77] S. Moosavi, M. Zanjani, R. Razavi-Far, V. Palade, and M. Saif, "Explainable ai in manufacturing and industrial cyber–physical systems: A survey," *Electronics*, vol. 13, p. 3497, 09 2024.
- [78] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," arXiv:1702.08608, 2017.
- [79] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004.
- [80] A. N. Elmachoub and P. Grigas, "Smart "predict, then optimize"," *Management Science*, vol. 68, no. 1, pp. 9–26, 2022.
- [81] L. von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, J. Pfrommer, A. Pick, R. Ramamurthy, M. Walczak, J. Garcke, C. Bauckhage, and J. Schuecker, "Informed machine learning – a taxonomy and survey of integrating prior knowledge into learning systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 1, pp. 614–633, 2023.
- [82] H. Cao, H. Söderlund, Q. Fang, S. Chen, L. Erdal, A. Gubartalla, P. V. Lopes, G. Shao, P. Lönnehed, H. Putto, A. Ahmed, S. Ekered, and B. Johansson, "Towards ai-based sustainable and xr-based human-centric manufacturing: Implementation of iso 23247 for digital twins of production systems," 2025. [Online]. Available: <https://arxiv.org/abs/2508.14580>
- [83] P. Kairouz, H. B. McMahan et al., "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.