



ML-Driven Compliance Auditing in Manufacturing with Digital Twin Integration

Mallesham Goli

Independent Researcher

mallesham.goli.research01@gmail.com

Abstract

Compliance auditing aims to assess adherence to regulations or standards, which is critical for managing operational risk. Regulatory drivers, including public expectation of sustainable practices in manufacturing, create urgency to introduce integrated compliance auditing into the industry. Although innovation in compliance auditing could significantly reduce the operational risk and increase the efficiency of businesses involved, the approach has received little attention in the literature. The current study investigates how predictive models of compliance outcomes can support auditing tasks and integrate into a digital-twin framework to provide forward-looking assurance. Research questions explore how predictive models can be developed on a specific data architecture to support compliance auditing and how a digital-twin model can provide the data needed to implement these predictive models.

The study applies a mixture of methodologies: prediction models capture compliance outcomes for a specific manufacturing plant, while the auditing process integrates machine-learning predictions into a digital-twin framework. Predictive models have been developed to support the specific industry context and inform compliance auditing. These models address both compliance auditing components of risk and capability. Risk-scoring models indicate violations of risk thresholds identified by audit procedures, and capability-scoring models indicate process capability for the assessed compliance aspect.

Keywords : Predictive compliance, Compliance auditing automation, Manufacturing compliance monitoring, Machine learning for compliance, Digital twin manufacturing, Real-time audit analytics, Anomaly detection in production, Risk-based compliance management, Quality assurance intelligence, Regulatory compliance forecasting, Smart factory governance, Process deviation detection, Industrial IoT (IIoT) data integration, Model-based compliance validation, Explainable AI (XAI) for audits.

1. Introduction

A wide spectrum of regulations, acts, standards, and codes of conducts govern the safety, security, and sustainability of products and manufacturing processes. The key aim of compliance auditing is to ensure that processes and products in production conform to these requirements. Compliance auditing of the manufacturing industry is largely triggered by regulators or industry authorities and is often a costly undertaking for the organizations involved. Major nonconformances triggered by these compliance audits can lead to serious operational implications for organizations, such as production halts, penalties, reputational damage, and loss of market share.

Reducing the number of minor nonconformances is essential for improving the efficiency of many compliance audits. Furthermore, predicting future compliance deviations enables proactive corrective measures that reduce regulatory risks during the next audit. Therefore, compliance audits would be more effective and efficient with predictive analytics capable of



identifying both near- and long-term deterioration in compliance performance. Industrialization of the twin concept alongside new data-driven technologies creates new opportunities to establish predictive compliance auditing in the manufacturing domain.

1.1. Background and Significance

Compliance auditing aims to assess the proper operation of industrial processes and the conformity of their execution to internal and external regulations. Fulfilling such objectives requires knowledge of regulation content, data preparation to cover all auditing dimensions, and the accurate mapping of judging criteria to data features. Drivers for compliance audits come from internal risk management, external inspectors, and important customers. Thus, predictive compliance auditing is a formal method to provide a foresight view on the results of internal and external audits, showing possible audit deviations before the audit moment. Predictive auditing is not meant to replace the audit operation. It rather reduces the risk of bad outcomes before the operation and prepares organizational readiness to face the audit.

Although predictive analysis is a common practice in many areas, its application within industry compliance auditing is still poorly explored. Two key factors stand in the way: the difficulty of identifying the critical data impacting audit results and the lack of technology for real-time compliance data monitoring. The first gap can be closed by machine learning techniques during the preparation, modeling, and training phases, while the second gap can be solved by using digital twins, which guarantee both structured real-time data generation and an integrated view of the entire manufacturing process lifecycle. Adoption of digital twin technology permits the creation of a digital thread with connectivity through the entire lifecycle and between involved parties.




Fig 1: Impact of Digital Twins on Real Practices

1.2. Research design

Research questions specify the problem to be tackled via the stated hypotheses, and the combination of data sources and methodology blend chosen for the research provide direction for expected contributions. With aspects relating to data architecture design and preparation, model selection and evaluation metrics, predictive audit workflow, and the role of a digital twin cross-referenced and consolidated in subsequent sections, the work's focus shifts to predictive modelling of manufacturing compliance prospects. Predictive compliance models are a type of forecasting tool; given the definition of audit as an evaluation of statutory compliance at a specific time, predictive compliance models may be interpreted as proxies for forward-looking audits.



The methods central to forward-looking audit procedures comprise the identification of signals indicating an unfavourable decision outcome, the scoring of metrics' capacity to ensure regulatory adherence, the analysis of process characteristics affecting potential future breach of regulations, and the determination of reasons for materialised breaches. Typically used in finance and insurance, forecasting models serve to identify unusual behaviour negatively influencing future normality definitions. They outline risk of non-adherence at some positive future time horizon and are defined through a series of well-structured phases, including: (1) engineering of a feature set that captures the behaviour under investigation, (2) choice of an appropriate machine-learning model, (3) validation of the model using hold-out data to ensure good generalisation, (4) use of the model in production, typically under a thresholding logic producing an alert if the forecasted output belongs to an abnormal region, and (5) incorporation of an alerting scheme into the wider organisational setup supporting timely investigation of the alert reason. As noteworthy for predictive processes, a separate why-terrible's-happening question is tackled through capacity and root-cause analysis procedures.

2. Background and Foundations

Regulatory authorities worldwide have developed technical and organizational frameworks that periodically assess compliance with health, safety, tax, financial, environmental, and other regulations. Typically, these audits are conducted during specific timeframes and by external accreditation entities as part of a mandatory scheme. However, the immense cost and resources needed have led to studies investigating alternative continuous compliance approaches. Traditional auditing has been defined as expensive, mostly retrospective, cumbersome, and unsustainable. Machine learning techniques have shown significant promise in direct mapping of regulatory risk, detection of deviations, and monitoring of processes. Nevertheless, adopting predictive compliance models in practice remains limited.

A Digital Twin (DT) can act as a living audit simulation environment able to mirror the physical counterpart and have bids for certifying compliance during the operational phase. Properly defined and governed, the DT represents a comprehensive environment and integration point to support compliance framework application supporting regulation agencies, manufacturing organizations, and other stakeholders. It can support comprehensible and communicable continuous change. Meanwhile, the DT interacts with conventional content-based repositories supporting transparency, auditability, and comprehension.

Equation 1: Classification evaluation metrics (Accuracy, Precision, Recall, F1, AUC)

In the PDF, I derive (step-by-step from the confusion matrix):

- **Accuracy** = $(TP + TN)/(TP + TN + FP + FN)$
- **Precision** = $TP/(TP + FP)$
- **Recall (TPR/Sensitivity)** = $TP/(TP + FN)$
- **F1** = $2 \cdot (\text{Precision} \cdot \text{Recall})/(\text{Precision} + \text{Recall})$
- **ROC**: TPR vs FPR where **FPR** = $FP/(FP + TN)$



- **AUC** as the (numerical) integral under the ROC curve.

2.1. Compliance Auditing in Manufacturing

Compliance auditing determines if a business conforms to relevant regulations or internal governance policies. The scope is wide in manufacturing, covering product quality, environmental emissions, health and safety, and fiscal responsibility, with numerous standards derived from the ISO family. Auditing frameworks may also originate from regional or governmental initiatives, or from industry-specific de facto regulations (e.g., for food production). Auditors assess compliance through different approaches: first-party (internal audit), second-party (supply chain audit), or third-party (contracted audit), where the latter case most commonly applies.

Compliance auditing significantly reduces risk exposure and the resources required to react to an observed deviation; however, it is usually a resource-demanding, one-time event. Auditors do not explore all process facets in depth, as evidenced by the traditional reporting metrics—compliance status usually quantified as a percentage. These limitations have motivated research on predictive methods, specifically in relation to ML models that simulate a future compliance outcome and integrate deviations into a broader risk analysis framework. Such predictions can reduce the time, effort, and cost of an audit by focusing resources where the business is most exposed.

2.2. Machine Learning in Industrial Settings

A broad range of machine learning algorithms has been successfully applied in manufacturing environments, and the same methods could be employed for predictive compliance auditing. However, specific industrial characteristics affect their applicability. Such challenges, together with the model lifecycle and deployment aspects, are addressed in subsequent paragraphs. Data are particularly important for model performance, and dataset preparation often represents the most demanding phase of model building. Indeed, garbage in, garbage out is even more relevant for machine learning models than for the classical statistical approach. The three main aspects for ensuring a good dataset are integration, quality, and privacy management.

Integration is key in machine learning applications because many different sources capture multiple events during a manufacturing process. A failure in any input feed stream can break the data thread. Internal and external integration dimensions also need to be satisfied. Data stewards and owners should manage the lineage of the data to ensure end-to-end traceability and accountability during auditing. Data quality management is essential in machine learning applications because the algorithms detect patterns, dependencies, and correlations on the input data. Therefore, data quality is critical in modelling processes and feeding the model. Privacy in machine learning models is a traditional concern, and threat mitigation is commonly implemented by adopting security measures such as encryption or anonymization. Thus, controls are able to convert a common risk into an acceptable and manageable condition. These three aspects lay the foundations for a successful application of machine learning algorithms in a manufacturing environment.


2.3. Digital Twin Technology and Its Role in Compliance

Digital twins are virtual representations of physical entities or processes, since these models (sometimes also called simulations) are embedded in real-world systems and continuously updated with real-time data. Such digital counterparts can be used to mirror any aspect of a system, including its geometry, physical behaviour, performance characteristics, and operations. They would thus be widely relied on for monitoring, testing, and auditing purposes across many sectors. Compliance demands real-



world developments to be mirrored in digital twins in a way that allows for accurate suggestions on process decision making, adjustments, and risk mitigation. Such suggestions are fundamental to the predictive capabilities of the proposed auditing solution.

A digital twin need not represent the real-world process at a high level of realism, yet it should display sufficient fidelity for its given application. The complexity and structural organisation thus depend on the intended use, purpose, and resources available for development, maintenance, and testing. Digital twins supporting regulatory compliance auditing should be governed by the respective compliance auditing function, enterprise risk management, or the equivalent organisational unit. Close involvement of the auditing function in the criminal offence digital twin aids with model fidelity, process-engineering insights, understanding of causal relations, and the governing decision-making.



3. Methodological Framework



Supporting digital twin integration for predictive audits in manufacturing requires a robust data architecture and quality management procedures. Data streams originating from multiple information systems and devices are captured, integrated, preprocessed, and assigned lifecycle lineage before nourishing machine learning models and digital twins. Activation of the process twins allows prediction of compliance properties and deviations, with instant visual alerts when critical thresholds are crossed. Historical data of compliance performance and supporting features are made available to machine learning candidates for the estimation of process capability, post-AI risk scoring, process deviation detection, and causal-inference explorations.

Key techniques are selected according to their appropriateness for each predictive task, suitability for classification with rare events, and ability to handle behaviour discontinuities and time-series information. Data privacy and security concerns are addressed by controls around access authorizations, encryption, and incident responsiveness. Scaling to deal with larger processes, combining data from distinct manufacturers, and ensuring interoperability with legacy MES and ERP solutions are also considered.

3.1. Data Architecture and Quality Management

Data Architecture and Quality Management presents data sources, integration, preprocessing, lineage, quality gates, and privacy safeguards. Data constitute one of the four pillars for building ML applications. For predictive compliance auditing, data architecture is of paramount importance since correct results directly depend on the fact, and to some extent also quality, of the underlying data. Untrusted and unstructured data sources can generate ample problems for the efficacy of ML and therefore for predictive compliance auditing. The underlying message is clear: without trustable data, extensive deployments can worsen risk instead of mitigating it.

Nonetheless, the design of a robust data architecture not only prepares organizations for effective ML models but it also builds the basis for the creation of data-driven Digital Twins. A software-represented Digital Twin embeds all data, information, and knowledge generated during the entire lifecycle. As such, it becomes a fundamental platform for novel concepts such as Digital Threads or Digital Shadow, which naturally support the definition of a Digital Twin of an Organization.




Fig 2: Data-Driven Digital Twin Framework for Predictive Maintenance



3.2. Model Selection and Evaluation Metrics

An ensemble of ML methods addresses challenges inherent in compliance auditing, including unbalanced datasets featuring rare negative examples. Supervised learning techniques encompass XGBoost, decision forests, logistic regression, artificial neural networks, and both one-class and isolation forests. In addition to standard performance measures (accuracy, precision, recall, and F1-score), model-selection criteria emphasize AUC, calibration plots, and fairness—discriminatory treatment of different subgroups must be avoided.

Anomaly detection leverages unsupervised models that assign novelty scores to observations. Risk scoring—an interpretable approximation of a comprehensive ML model—combines estimates of process capability indices (Cp, Cpk, Pp, Ppk) for designated features of interest. Causal inference applies perturbation experiments on causal graphs inferred from TABNET representations of the data.

3.3. Digital Twin Integration for Compliance Simulation

The digital twin mirrors the physical manufacturing process and continuously receives operational data from the real system. The twin is conceptually extended with additional modules that simulate the results of changes to the system or predict deviations from the required performance. Support for predictive compliance auditing relies on scenario-testing capabilities incorporated into the digital twin. Such capabilities enable the evaluation of modifiable parameters and an exploration of what-if scenarios with respect to specific aspects of compliance.

Digital twins should be designed to go beyond serving basic monitoring purposes and enable the exploration of scenarios that require changes to processes, workflows, or parameters. One important category of such scenarios involves what-if analyses for investigating potential changes to governance frameworks or regulatory requirements. These investigations typically take place in a controlled environment that lacks real-world implications. Stress-testing scenarios assess the compliance position of a system when subjected to unusual operating conditions that expose its weaknesses.

3.4. Predictive Auditing Workflow

The predictive auditing workflow depicts the progression of data throughout the entire modeling process, encompassing model training, validation, deployment, monitoring, and feedback cycles. It draws from auditing process definitions to establish data origins, archives, privacy measures, and segmentation into training, validation, and test sets. During model training, suitable methods are chosen based on performance metrics. After training, models are validated using always-available and frozen datasets that support statistically sound assumptions. The predicted compliance outcomes are regularly evaluated. Developed for deployment, the models become part of the infrastructure, receiving real-time data throughout the manufacturing process for auditing reasons, and raising alarms if compliance is threatened. The importance of the AI-Bayati metric quantifying the model risks can trigger an early move from surveillance to preventive mode. These components enable an economic justification of the paradigm change.

A feedback loop incorporates labeled data—particularly useful when predictions are wrong—to re-train models. The entire workflow may benefit from a digital twin, which integrates all components and encompasses aspects not supported by historical data. The twin dataset defines the federal risk-assessment audience, and the associated compliance-simulator-and-manager ensures that potential compliance deviation is regularly measured and suggested preventive actions are sent in good time to the twin audience.



4. Predictive Models for Compliance Outcomes

Anomaly detection and deviation forecasting leverage supervised or unsupervised classification or regression models to support compliance auditing decisions. Classification or regression models take process time series, contextual factors, and lagged features related to prior deviations as input, and predict the class labels for the following period of a process or process capability indices. Deviation alerts are generated when/if incoming predictions fall into a deviated class. For classification models, the thresholding strategy can also be a calibrated probability threshold instead of class labels, further reducing false alarms. The approach is thus able to forecast compliance deviations and, through the process capability index, degree of these deviations.

Process capability indices are calculated as risk scores whenever a process or capability is within an interval (e.g., monthly). A priori defined thresholds for the severity degree (risk score interpretation) then guide executives in their prioritization decisions for detailed root cause analysis and appropriate mitigation actions. Threshold and index definitions incorporate operational experience and serve as a collective agreement to aid in decision making. Root cause analysis employs a causal graph of process conditions and inter-stage variables linked to the process capability indices, built with domain knowledge and augmented through perturbation synthetic experiments. A counterfactual reasoning analysis helps to identify the root causes of a specific process deviation in a certain period.

Equation 2: Process capability indices (C_p , C_{pk} , P_p , P_{pk}) used for compliance/risk scoring

Predictive Compliance Auditing ...

In the PDF, I formalize:

- Predict probability $\hat{p} = P(y = 1midx)$
- Alert rule: trigger if $\hat{p}getau$
- And show how changing τ changes FP/FN tradeoffs (linked to ROC/PR behavior).

4.1. Anomaly Detection and Deviation Forecasting

Predictive models supporting compliance auditing encompass anomaly detection, deviation forecasting, scenario testing, root cause analysis, risk scoring, and capability assessment. Anomaly detection identifies deviations requiring investigation; deviation forecasting indicates likely future breaches of compliance limits based on current trends. Both model types share the characteristic that positive predictions merit scrutiny, even if not all true anomalies can be discovered and some false alarms invariably arise. First, detection techniques for multivariate data are outlined, followed by the approaches for forecasting deviations from acceptable regions.

Multivariate anomaly detection relies on appropriate feature selection, transformation, and visualisation as well as on the chosen detection technique. Moreover, the determination of a threshold informing the alerting mechanism is required. A hybrid approach employs a standard machine learning model for detection in the feature space before validating findings through a different class of method. Random Forests identify non-compliant windows based on specialist features, while a recurrent neural network with long short-term memory units leverages both historical temporal information and lagged variables. Deviation forecasting is



performed by a recurrent neural network with attention mechanism, producing multivariate predictions supplemented by prediction intervals for interpretability.

4.2. Process Capability and Risk Scoring

Capability indices assess a process's compliance performance over time, while risk-scoring algorithms quantify compliance risk for prioritized action. Calibration failure signals to decision makers that the risk score lacks meaningful interpretation. Grounded on causal graph analysis, risk-underlying factors are detected through perturbation experiments and countefactual reasoning.

Performance monitoring compares real data against statistical distributions derived from historical records. Significantly different behaviour (deviations) raises alerts that, for manufactured components, indicate non-compliance risk. Failed alerts point to a non-fulfilling capability beyond temporal boundaries. Process capability is computed with established capability-indices for C_p , C_{pk} , P_p , and P_{pk} . Risk representation rests on a support vector machine trained with a time-window history of the calculated indices labelled as 'RISK' (1) or 'NO RISK' (0). Calibration during deployment needs supervision, which can exploit information on ground-truth risk.

The capability-index framework allows manufacturing companies to compute a risk score at any moment. This score suddenly becomes a priority for action whenever it reaches a specific threshold. Audit results inform the risk score, functioning as a control loop that defines preventive or corrective actions. Causal-graph analysis reveals risk-underlying factors, identified through perturbation experiments and counterfactual reasoning for intent-focusing actions.

4.3. Root Cause Analysis and Causal Inference

Causal graphs outline putative relationships between compliance-performance drivers and outcomes, aiding identification of crucial predictors. Perturbation tests gauge models' responsiveness and help establish relevance of significant influences. Counterfactual reasoning tests model behavior and evaluates the accuracy of underlying assumptions, serving to identify important interaction effects.

Mining models highlight changes in operational settings, features, or models that raise compliance-risk scores. The multidimensional aspect of risk emphasizes the importance of monitoring all three scores and alerts users to any of these deviations. The resulting causal-graph visualization can help users comprehend the root causes of observed variations in the compliance-risk score.

Empirical validation demonstrates predictive-auditing capabilities using ML models explicitly designed for audit activities in manufacturing processes. Focused on critical quality nonconformities—a particular type of compliance deviation with a large negative impact on product quality—three key predictive-auditing tasks are considered: anomaly detection, deviation forecasting, and risk scoring.

5. Digital Twin as a Compliance Assurance Platform

Digital Twin as a Compliance Assurance Platform



A digital twin supports risk-based compliance auditing by monitoring process parameters, construction of a factual digital thread, what-if scenario analyses, and other means. It accepts time-series data streams from the physical system, providing real-time monitoring capabilities as well as traceability throughout an end-to-end production lifecycle. In addition to alerting toward potential risks or noncompliance, it enables testing of process and regulatory changes with associated what-if analyses.

The digital twin also supports machine learning model implementation for predictive auditing. The combination of a digital twin and predictive compliance auditing represents a significant innovation for manufacturing risk management. It tests the feasibility of predictive compliance auditing in a multiple-batch context and provides an initial step toward its application for accreditation processes that occur frequently and demand substantial effort. The approach can address traditional compliance auditing techniques' limitations in predictive capability, efficiency, and acceptance by auditors and auditees.




Fig 3: Digital Twin as a Compliance Assurance Platform

5.1. Real-Time Monitoring and Digital Thread Construction

Life cycle traceability of compliance auditing relies on continuous data monitoring throughout the whole life cycle of physical items—covering, for example, the production of materials, manufacturing, logistics, maintenance, and end-of-life phases—with a detailed digital thread consisting of the generated data and their provenance. A transparent analysis of the process and data traceability allows replicability and auditability. A digital twin monitors processes in real-time, guaranteeing effective compliance testing of the current status, while a digital thread provides proof of adherence over time through the archiving of data and the relations among data sources. It is thus possible to automatically detect deviations in the process, check for compliance with regulations and standards, and deliver an audit report containing supporting evidence. To comply with regulations using machine learning, the data should represent a process or item that has evolved over time so that its behavior can be clearly understood. Every data point should bear sufficient context to allow for clear identification of and explanation for a deviation from the normal operation of the process or item.

The audit follows a user-defined procedure, generating standard reports as well as deviation alerts. Alerts provide stakeholders with timely information for detecting anomalies requiring immediate attention; an automated warning can be sent through the




appropriate internal communication channel. In addition to informing users of current monitoring results, the model periodically runs on the most recent set of data points, updating the view of process capability indices, risk scores, and root-cause analysis.

5.2. Scenario Testing and What-If Analyses

The digital twin's ability to simulate process behavior, with a fundamental and vertical level of fidelity, enables valuable testing of various business process scenarios. For compliance auditing, what-if analyses cover regulatory changes, process capability and performance improvements, and risk scenario testing. The virtual representation reflects an updated and controlled visual state of the business process, grounded in historical data and informed about current steers, thresholds, and operating conditions. Therefore, the organization can execute a virtual audit with a compliance assurance officer to determine the capability of the process to satisfy the policy and any deviations expected and scheduled to occur.

The model governing the business process acts as a vehicle to validate new models or procedures resulting from the evolution of regulations or standards. By perturbing the governing model, simulations allow for the identification of any expected non-compliance, compliance deviation, or possible improvement. Testing any variation of the model with the underlying process-focused governing model supports an understanding of the risk related to abnormal operating conditions.

Digital twins encompass more than simple historic reporting and serve as a compliance assurance layer by enabling a digital thread. All pieces of information regarding the desired properties of the product/process/service and the performance of the business process itself are continuously monitored and verified by the virtual representation. Therefore, compliance assurance officers can assure stakeholders that the process is or will be compliant, and these guarantees can be provided with an enhanced level of quality found in the confidence factor.



5.3. Governance, Transparency, and Explainability

Model explainability enhances trust and engagement. Transparency empowers stakeholders, providing an overview of decision-making processes. Internal auditors rely on audit records to validate studies and assess application suitability. Clear communication aids misunderstandings, clarifies outcomes, and exploits risk score drivers for targeted awareness-raising and training.

The meta-model implements the explainable AI methodology. For classification tasks, SHAP extracts the influence of individual features, while global interpretations employ LIME, decision trees, or rule concentration. Limitations and justifications accompany control interventions. Predictive models perceptible to the auditors facilitate validation and instill confidence. A digital thread combines data lineage and metadata provenance, supporting scrambling and encryption. Audit trails monitor sensitive data queries. Data subject requests are traceable, systematically fulfilled, and disclosed when affected by model retraining. Alerts reporting deviations, non-compliance, and incident responses ensure timely areas of concern and improvement.



6. Implementation Considerations in Manufacturing Contexts

Considering data privacy, security, and access controls is a necessary prerequisite for predictive compliance auditing in manufacturing. Sensitive data must be shielded from unwanted access, use, and loss. Confidentiality threats can originate either externally, from cybercriminals motivated by financial gain, or internally, from authorized users with malicious intentions. Protection can be enacted through strong policies and secure technological solutions. Information security relies on techniques such as data segmentation, encryption, access control, and incident response management; the latter seeks to minimize the probability of occurrence, damage, and recovery time of security breaches. Data security establishes safeguards against unauthorized information disclosure—such as using fine-oriented role-based access control, which incorporates the concept of ownership; access to conflicting or related information is denied to different individuals. Implementing role-oriented access rights would limit the number of users with access to sensitive information or allow the information to be encrypted and decrypted using separate key encryption schemes (protected by distinct passwords).

Predictive compliance auditing considerations must also include adequate scalability, interoperability, and the ability to follow data management standards. The demand for predictive analysis in the manufacturing industry calls for scalable technological solutions that can manage and support analysis of very large sets of data; such scalability can be achieved by distributing data into cloud environments. A crucial consideration when developing a cloud-based system that will be used in the production environment is whether the solution is compatible with data standards. In addition, for predictive models to be useful in a practical setting, models have to be integrated into manufacturing execution systems or enterprise resource planning packages that are widely adopted and used in the industry. Consequently, predictive compliance models have to be packaged and exposed as services so they can be accessed correctly from the back-end system within which they are integrated.

Equation 3: Process capability indices (C_p , C_{pk} , P_p , P_{pk}) used for compliance/risk scoring

In the PDF, I derive:

- $C_p = \frac{USL - LSL}{6\sigma}$
- $C_{pk} = \min\left(\frac{USL - \mu}{3\sigma}, \frac{\mu - LSL}{3\sigma}\right)$
- $P_p = \frac{USL - LSL}{6s}$, $P_{pk} = \min\left(\frac{USL - \mu}{3s}, \frac{\mu - LSL}{3s}\right)$

...and I include:

- A **distribution plot** with LSL/USL and mean (visual intuition of C_p vs C_{pk})
- A **bar chart over months** showing how capability indices can be tracked and thresholded operationally.

6.1. Data Privacy, Security, and Access Controls

Data privacy, security, and access controls comprise essential facets of any data-driven initiative, yet their specifications are often relegated to footnotes after priority design decisions are made. In this application, however, these aspects not only steer the project design but also preserve the essence of a predictive compliance model without undermining the privacy requirements of sensitive data. Manufacturing often involves multifaceted data from various sources, including data from sensors, embedded



chips, operation systems, and connected devices. Because of its complex data composition, for most end users, the compliance of model-related processes is categorized as “not ready” or “avoiding.” The predictive compliance auditing components interrogate compliance, risk assessments, and process monitoring and implement warning systems accordingly.

Privacy and security policies are governed by legal issues (GDPR consent and related laws). Data management goes through restriction mechanisms. Data are scientifically collected but cryptographically encrypted, and model management follows an authorization approach. Role-based access control (RBAC) is implemented by restricting user activity according to the user’s role in the organization. RBAC enables the organization to efficiently administer the specification of permissions for users of different roles in an organization. Using RBAC, only the data administrators have all the permissions to an organization’s data after the system has been implemented, whereas other users have only a subset of the permissions determined by their roles. In addition, the entire system implements a full incident response capability to coprocess incidents involving data management; for example, when the account of a user is hacked and that user has access to data that is not supposed to be shared, privacy and security can be implemented through damage limitation and damage control.

6.2. Scalability, Interoperability, and Standards

Scaling operations to meet demand fluctuations while maintaining optimum performance requires a production architecture that supports parallelisation of resources. Considerations include load-balancing among digital twins of functional units in system, pattern detection in real-time data for demand forecasting, configuration setup on Area Management System of SCADA, and rollup into informative visual dashboard. Incremental or radical design changes in plant layout call for assessment of scale and delivery to accommodate predicted alterations in product mix. Issues of data quality quality require addressing prior to machine-learning modelling.

Achieving tangible benefits of predictive compliance auditing depends on the architecture’s ability to measure data, information, and model quality-as represented by the technology- and governance-spanning concept of a digital thread. Demonstration of the approach across several manufacturing domains will create the knowledge repository needed to mitigate existing fragmentation and develop the necessary standards for generative semantic interoperability between digital twins and Model Exchange Libraries in subsequent implementations.

6.3. Change Management and Organizational Readiness

A multifaceted approach is necessary for stakeholder support and readiness for predictive compliance auditing. Training and coaching increase knowledge and skills for practical system use, and stakeholder feedback shapes requirements for effective development. Resistance to change is mitigated by involving stakeholders in project definition and collaborating on measurable impact metrics.

Predictive compliance auditing constitutes a fundamental departure from traditional approaches. As with any other pervasive process or technology transformation, a careful change management strategy is needed. Stakeholders with varying technical backgrounds, motivations, and objectives must collaborate to design, develop, and validate systems that meet their needs. Appropriate training, instruction, coaching, and practical exercises help increase the readiness of system users in their particular areas. To maximize the success and facilitate acceptance of predictive compliance governance systems, involvement, and feedback from the various stakeholders must inform all phases of system construction.



Stakeholder acceptance is fostered by linking predictive compliance governance capabilities to measurable operational performance. For example, the automotive industry has shifted toward predictive quality control and process-in-control assurance. Such initiatives reduce production costs and, in turn, improve profitability. Adopting predictive compliance auditing requires change management to promote stakeholder acceptance, underpin innovation design choices, and define quantitative metrics to facilitate adoption.

7. Validation, Evaluation, and Benchmarking

Validation, Evaluation, and Benchmarking: Validation and empirical evaluation are pursued through multiple case studies that assess the effectiveness of models and the integration of predictive compliance auditing technology in a manufacturing environment. Other analyses compare predictive auditing against traditional approaches in terms of perceived effectiveness, efficiency, and public acceptability. The validation and case studies are complemented by a summary of limitations, sources of uncertainty, and strategies for confidence quantification associated with predictive compliance auditing.

Case Studies: Analyses employed representative case studies—digital fingerprints and alarm threshold setting in an air quality compliance context, continuous S function scoring of reheating furnaces, and spare-part breakdown prediction on a company-owned wind farm. Additionally, beyond the empirical outcomes of case studies, the practical use of predictive compliance auditing was explored by comparing it with traditional auditing and risk-assessment methods based on auditing experts assessing compliance risk on the selected processes. The focus was on the three attributes of perceived compliance auditing methods, the Operational Excellence Triangle (effectiveness, efficiency, and public acceptability), and the trade-offs monitored recommended for determining risk assessment methods.




Fig 4: Digital Twin for test and validation

7.1. Case Studies and Empirical Results

A range of case studies demonstrates the practical applicability of the proposed predictive audit framework. Various datasets cover different processes in diverse industrial sectors such as food manufacturing, transportation, and finance; industrial maintenance; and software development and delivery. Examples of predictive auditors include models predicting process capability indices, risk scores, warning signals, and deviations from a priori established tolerances. Internal stakeholders acted as



end users throughout the projects by providing input data, selecting relevant features, establishing thresholds, defining target audiences, and reviewing interpretation reports.

A first evidence-based study sheds light on compliance assurance during the development and release of software updates. GDPR obliges establishments to inform customers promptly when breaches occur. A privacy risk score quantifies reputation loss: the higher the score, the higher the information flow risk associated with software updates. A second study targets long-term structural risk management of physical infrastructures, such as roads. A set of predictors estimates structural condition index values during the life cycle of transportation infrastructures. Estimation values above defined thresholds identify physical roads at risk of noncompliance with patency standards. A third analysis evaluates conditions for internal process compliance with software quality standards in a bank. The ability of a decision tree classifier to detect anomalies in internal software engineering processes—integrating requirements, development, testing, deployment, and maintenance stages—opens up a preventive perspective for internal software quality audits.

7.2. Benchmarking Against Traditional Auditing Methods

Empirical results demonstrate how predictive compliance auditing leads to improvements in effectiveness, efficiency, and acceptance compared with traditional auditing methods. The studies involve both quantitative validation and qualitative assessment within an action-research framework. Predictive models incorporate machine-learning algorithms that consider the temporal dynamics of deviations and their contextual factors, while digital twins enable the construction of impact, capability, and risk scores. A large-temperature-controlled-regulated-laboratory case study illustrates the concept.

Internal regulatory compliance needs to be checked beyond the conventional routine audit strategy, which relies on historical non-compliance events to detect deviations and assess business risk, thus missing alerts in theory but never in practice. Predictive auditing combines digital-twin technology and machine-learning modelling to identify future non-compliance events, their underlying causes, and the effectiveness of potential preventive actions—an ambitious objective. In an action-research setup, a real-use case application suggests how predictive auditing may, over time, tend towards a comprehensive replacement of conventional internal audits.

7.3. Limitations and Uncertainty Quantification

Predictive compliance auditing contributes to a risk-aware culture where regulated conformity is crucial for stakeholder safety and sustainability. Nevertheless, automatic compliance assurance remains elusive, and empirical studies are scarce. Validation, evaluation, and benchmarking embrace these challenges and establish conditions for industry uptake. Limitations of the proposed solutions, along with the burden of latent uncertainty, are explicitly acknowledged. Analyzing results from multiple case studies with diverse data patterns adds depth to the investigation and redresses previous criticisms about the absence of real-life data. Inadequate data quality still poses a challenge that may inhibit generalization and model transferability. However, hidden deviations, data retentiveness, and risk scoring attainment are now feasible without conventional and time-consuming manual labor.

All candidates for the predictive auditing models are validated, with precision, recall, area under the curve, calibration, and fairness incorporated in the selection criteria. The simulation-based assessment of machine learning models is finally complemented by noise injection and reinforcement learning for hyperparameter calibration. Anomalies and deviations are detected and forecasted accurately and reliably.



8. Conclusion

Classification, detection, and risk scoring models utilize standard data inputs and can be deployed within a compliance-as-a-service model, while root-cause analysis relies on custom architectures for each process. Digital twins serve as a compliance-assurance platform. Automated alerts, risk scores, compliance scores, and key risk indicators leverage real-time data. The twin simulates processes to allow what-if analyses and scenario testing for regulatory change, risk management, and process improvement. Transparent models enable users to understand deviation drivers. Incorporating machine learning into predictive compliance auditing offers substantial efficiency gains over traditional manual methods.

Future trends point toward digital twins that support full-life-cycle management and embrace questions of trust, robustness, and accountability. Their governance will evolve from a spare-mode flagging of time slices that require attention toward a full-time service that scans products for accessibility and compliance, augments them with features that are relevant for multiple stakeholders, and helps to develop compliance scores that permit a more client-friendly, tailored service. The anticipated technological developments and resulting impact on transparency will establish the foundation for automated compliance. In practice, the roadmap for adoption includes auditing services, digital-twin capacities, an ecosystem of accessible processes, and autonomous compliance platforms capable of full-life-cycle management.

8.1. Future Trends

Future trends in compliance auditing will see the increased adoption of advanced technologies and more focus on governance and risk. As generative AI tools mature, models will become more interpretable, better exposed to the fundamental governing the processes, and able to incorporate business logic. Regulatory changes will shape the direction of modelling efforts, placing emphasis on ‘what-if’ analyses. Models that are consumed across the organisation will also gain attention, either regulatory-related governance around model explainability or translation into business friendly language.

The predicted compliance models may also act as a first detection mechanism that serve not only the external auditors but internal quality functions as well. Finally, the advances offered by machine learning and supporting technologies in areas such as context understanding, explainability, predictability, and data capture are likely to converge towards predictive auditing, increasing trust in machine-assisted supervision. The path towards such an integration can be facilitated by addressing the aforementioned challenges and ensuring the developments take a comply-and-explain approach.

9. References

- [1] Alcaraz, J. L. G., Macías, A. A., & Buxmann, P. (2020). Digital supply chain transformation: A literature review and practical insights. *International Journal of Production Research*, 58(15), 4574–4591.



- [2] Birihanu, E., Hammar, K., Främling, K., & Holmström, J. (2025). Explainable correlation-based anomaly detection for industrial control systems. *Frontiers in Artificial Intelligence*, 7, Article 1508821.
- [3] Carvalho, T. P., Soares, F. A. A. M. N., Vita, R., Francisco, R. P., Basto, J. P., & Alcalá, S. G. S. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering*, 137, 106024.
- [4] Chen, Y., Yang, O., Sampigethaya, K., & Poovendran, R. (2020). Digital twin security: A survey. *IEEE Access*, 8, 215948–215971.
- [5] Choi, H., Lee, H., Kim, H., & Lee, J. (2022). Explainable anomaly detection framework for predictive maintenance in a manufacturing system. *Applied Soft Computing*, 118, 108460.
- [6] Cimino, C., Negri, E., & Fumagalli, L. (2019). Review of digital twin applications in manufacturing. *Computers in Industry*, 113, 103130.
- [7] Dai, D., Huang, Z., & Xu, Y. (2025). Generative and predictive AI for digital twin systems in manufacturing. *Frontiers in Artificial Intelligence*, 8, Article 1655470.
- [8] Dogan, O., Hacıoglu, U., & Eren, E. (2024). Empowering manufacturing environments with process mining and statistical process control integration. *Machines*, 12(6), 411.
- [9] Eirinaki, M., Patel, D., & Joshi, N. (2020). A survey of anomaly detection in industrial IoT. *IEEE Communications Surveys & Tutorials*, 22(3), 1853–1881.
- [10] Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952–108971.



- [11] Gruber, S., Matzner, M., & Staudt, M. (2021). Process mining for compliance and auditing: A systematic literature review. *Information Systems*, 98, 101715.
- [12] He, Y., Guo, Y., Zheng, Y., & Zhang, C. (2021). Deep learning for anomaly detection: A review. *IEEE Access*, 9, 113158–113175.
- [13] ISO. (2018). ISO 9001:2015 quality management systems—Requirements. International Organization for Standardization.
- [14] ISO. (2023). ISO 23247-1:2021 automation systems and integration—Digital twin framework for manufacturing—Part 1: Overview and general principles. International Organization for Standardization.
- [15] Jia, Z., Liu, F., & Zhang, C. (2022). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 73, 102229.
- [16] Klessascheck, F., Becker, J., & Kalenkova, A. (2025). Reviewing uses of regulatory compliance monitoring. *Software and Systems Modeling*, 24, 1–25.
- [17] Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016–1022.
- [18] Lee, J., Davari, H., Singh, J., & Pandhare, V. (2018). Industrial AI and cyber–physical systems for smart manufacturing. *Manufacturing Letters*, 18, 20–23.
- [19] Lee, J., Kao, H. A., & Yang, S. (2014). Service innovation and smart analytics for Industry 4.0 and big data environment. *Procedia CIRP*, 16, 3–8.
- [20] Li, C., Mahadevan, S., & Ling, Y. (2021). A physics-informed machine learning framework for digital twin enabled smart manufacturing. *Journal of Manufacturing Systems*, 60, 418–432.



- [21] Liao, Y., Deschamps, F., Loures, E. F. R., & Ramos, L. F. P. (2017). Past, present and future of Industry 4.0—A systematic literature review and research agenda proposal. *International Journal of Production Research*, 55(12), 3609–3629.
- [22] Liu, Z., Meyendorf, N., & Mrad, N. (2018). The role of data fusion in predictive maintenance using digital twin. *International Journal of Prognostics and Health Management*, 9(1), 1–18.
- [23] Lu, Y., Liu, C., Wang, K. I. K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61, 101837.
- [24] Moosavi, S., Jafari, M., & Ghasemaghaei, M. (2024). Explainable AI in manufacturing and industrial cyber-physical systems: A survey. *Electronics*, 13(17), 3497.
- [25] Negri, E., Fumagalli, L., & Macchi, M. (2017). A review of the roles of digital twin in CPS-based production systems. *Procedia Manufacturing*, 11, 939–948.
- [26] Park, G., van der Aalst, W. M. P., & Cho, S. (2025). Operational process monitoring: An object-centric approach. *Computers in Industry*, 155, 103700.
- [27] Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison. *IEEE Access*, 6, 3585–3593.
- [28] Raj, A., Dwivedi, G., Sharma, A., de Sousa Jabbour, A. B. L., & Rajak, S. (2020). Barriers to the adoption of Industry 4.0 technologies in manufacturing: An interpretive structural modeling approach. *Journal of Cleaner Production*, 268, 122078.
- [29] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?” Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.



- [30] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
- [31] Seshia, S. A., Sadigh, D., & Sastry, S. S. (2016). Towards verified artificial intelligence. *arXiv*.
- [32] Shafiq, S. I., Sanin, C., Szczerbicki, E., & Toro, C. (2020). Virtual engineering object / virtual engineering process: A specialized digital twin framework for manufacturing. *Procedia Computer Science*, 180, 687–695.
- [33] Tao, F., Zhang, M., Liu, Y., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
- [34] Tao, F., & Qi, Q. (2019). Make more digital twins. *Nature*, 573(7775), 490–491.
- [35] van der Aalst, W. M. P. (2016). *Process mining: Data science in action* (2nd ed.). Springer.
- [36] Villegas, L. F., Aoyama, K., & Ríos, J. (2025). Digital twins in manufacturing: A unified conceptual framework and research directions. *CIRP Journal of Manufacturing Science and Technology*, 50, 1–20.
- [37] Wan, J., Tang, S., Hua, Q., Li, D., Liu, C., & Lloret, J. (2018). Context-aware cloud robotics for material handling in cognitive industrial internet of things. *IEEE Internet of Things Journal*, 5(4), 2272–2281.
- [38] Wuest, T., Weimer, D., Irgens, C., & Thoben, K.-D. (2016). Machine learning in manufacturing: Advantages, challenges, and applications. *Production & Manufacturing Research*, 4(1), 23–45.
- [39] Xu, X., Lu, Y., Vogel-Heuser, B., & Wang, L. (2021). Industry 4.0 and Industry 5.0—Inception, conception and perception. *Journal of Manufacturing Systems*, 61, 530–535.
- [40] Zhang, H., Wang, J., & Xu, Y. (2022). Data-driven anomaly detection in smart manufacturing: A review. *Journal of Manufacturing Systems*, 62, 468–485.

