International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

Digital Twin-Based Audit Engine for Explainable Machine Learning in Automated Manufacturing Compliance

Ghatoth Mishra

Independent Researcher, India

ghatothmishra@gmail.com

Dr. Daniel Moreau

Research Support Assistant, École Internationale de Lyon, France

danielmoreauu@outlook.com

Publication History: Received: 10.05.2026; Revised: 26.06.2026; Accepted: 01.07.2026; Published: 14.07.2026

ABSTRACT: Automated manufacturing is increasingly reliant on machine-learning models to support operations and decision-making, at the same becoming highly regulated and monitored by assurance frameworks that seek to enhance customer trust. Responding to safety, cyber-security, and sustainability requirements creates a burden represented by manual audits consuming part of the operational budget. Furthermore, machine-learning models are black-boxes, even for the processes governed by them. Knowledge gaps prevent the construction of risk-averse systems capable proactively adapt to noncompliance thresholds. Indeed, a technology is needed to match the demand of industry and comply with regulations in a trustworthy manner. Digital twins can ingest data from a physical manufacturing operation and, with sufficient fidelity, can be used for onside auditing of a dedicated or industry-wide audit engine hosted on the cloud

Artificial Intelligence and in particular Machine Learning technologies have a growing role in society and industry. In particular, Machine Learning is used in automated systems that can be examined both for the decision support they provide and for their own decision-making processes. Explainability is seen as a necessity to guarantee the correct operation of a system, especially when humans are responsible for operating the machine or the decisions taken therein. A growing number of frameworks, regulations, and directives support the implementation of explaining techniques for Artificial Intelligence used in product safety, data protection, and sustainability contexts. Nevertheless, these efforts are often disconnected from the rapidly evolving technologies and their requirements of respect for human factors in auditing, acting, and interpreting.

KEYWORDS: Explainable Artificial Intelligence (XAI), Proactive Compliance Monitoring, Automated Manufacturing Systems, Digital Twin Technology, AI-driven Compliance Auditing, Interpretable Machine Learning Models, Industrial Cyber-Physical Systems, Real-time Regulatory Conformance, Model Transparency and Trust, Smart Factory Governance, Anomaly Detection for Compliance, Predictive Risk Assessment, Human-in-the-Loop Decision Support, Industry 4.0 Compliance Frameworks, Audit Automation and Traceability.

I. INTRODUCTION

Manufacturing industries are increasingly adopting fully automated production environments with minimal human involvement. In these systems, the conventional role of the human operator is progressively replaced by algorithms and artificial agents that monitor the execution of the workflow, make decisions, and take actions determined by the algorithmic logic. The majority of those decisions are being made by machine-learning (ML) models trained on salient features of historical production data that are aggregated to reflect specific stages of the workflow or the whole system. These decisions are often seen as "intelligent" because the ML models are intrinsically capable of identifying complex data patterns and data relationships. However, unlike the traditional operation of a production workflow where humans apply their own experience, knowledge, understanding, and gut feelings to assess the act of making decisions, the ML models are mere black-box function approximators that can neither reason nor offer any explanations for their decisions.

International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

Such black-box behavior poses high challenges for the automated systems since the operations industry need to become accountable for the decisions taken by nonhuman agents, being exposed to the same kind of regulatory frameworks that govern any action executed within a human-driven workflow. Therefore, the explanation of machine-learning models and, more generally, the concept of Explainable Artificial Intelligence (XAI) are emerging as hot areas of research, both in academia and in industry. Aligned regulation and supervision need to be present along the entire production workflow and its dynamics in order to safeguard quality, safety, and environment. Proactive compliance indicates that production operations should move beyond the mere detection of problems and should anticipate the occurrence of potential risks. Hence, the decision processes must not only provide accurate predictive models but also predictive models with the adequate level of explainability.

1.1. Problem Statement and Motivation: Regulatory compliance is traditionally ensured by scheduling regular audits of human-operated facilities for procedural adherence and resource and product quality. In automated production systems with sufficient data logs, regulation-driven auditability can be realized, but no structured framework currently exists. The lack of an explanation for critical ML-driven decisions taken during production constitutes a major roadblock to adoption and transferability to safety-critical industries. The inability to anticipate and mitigate compliance-related risks risks hindering the adoption of such systems. Aligning interventions with regulatory thresholds and detecting potential violations well in advance are crucial for both operational and strategic reasons.

Preventive interventions, along with human verification, guarantee compliance and protect against potential liability-triggering incidents. An audit engine, equipped with a suite of predictive, prescriptive, and explainable ML models, integrates with production Data Twin and external rule engines to support risk detection and mitigation, while a dedicated audit-data pipeline manages quality and traceability. A modular, backed component connects internal audit logic with regulatory rule sets. Potential gap detection and guidance, supported by interpretability tools, form a first step toward proactive compliance in risk-critical systems. Emerging Explainable AI methods and frameworks are reviewed, with a focus on preventive compliance.

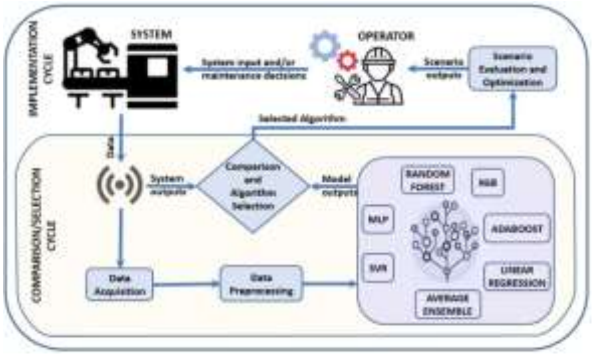

Fig 1: Data-Driven Digital Twin Framework for Predictive Maintenance

1.2. Scope and Objectives: Proactive compliance and risk anticipation in automated manufacturing systems have been defined and scoped with respect to auditability concerns. Regulatory-driven frameworks for auditability and governance—addressing the timing and governance of automated interventions—are mapped to the process of making ML-driven decisions and actions more transparent, verifiable, trustworthy, and easy to explain.

Development proceeds from a set of assets and enablers embodied in a Digital Twin that ingests, simulates, and generates data crucial for testing and validating—under the perspective of crowdsourcing models and actions—the ML-driven subsystem and its expected contribution to proactive compliance and risk anticipation. Gaps in reflexive models of explainable AI that would allow audit-as-a-service or compliance-as-a-service for ML have been identified, setting the stage for a broader research agenda on the explainability of AI in critical decision support applications.

The research introduces a structured approach to identifying and evaluating auditability gaps associated with automated decision support systems. ISO IEC 27001, covering the requirements for information security management systems, can serve as a reference standard for cyber-risk analysis and risk-based decision support in manufacturing. The coverage of digital twin and data-centric AI frameworks establishes a foundation for testing the resilience of ML-based decision support under edge conditions. Techniques from the emerging field of Explainable AI can be exploited to offer a

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

service—either internal or external—that is business-driven and directly connected to the risk domains defined in the cybersecurity framework.

1.3. Relevance to Regulatory Compliance and Operational Efficiency: Proactive compliance with regulatory, safety, and quality standards presents significant challenges in automated manufacturing. Automated decision making is particularly difficult to audit, especially when explainability is missing or insufficient. Tooling and parts suppliers are experiencing increased regulatory pressure, with an increasing number of directives and the requirement for comprehensive production records, including machine learning decision making. Recent Supply Chain Act requirements demand proper attention across all business operations and layers of responsibility in the supply chain. However, an audit-driven or retrospective compliance approach remains a considerable effort, often targeted only at critical control points instead of all production decisions. Explainable approaches must not only offer process support for generating explainability artifacts; they should also encompass audit-driven methods that anticipate Nonconformities according to regulatory boundaries.

Explainable ML methods and paradigms should support proactive compliance by anticipating risks or knowledge gaps during production and enabling the prior introduction of mitigation actions within the workflows. A twin-based digital audit engine capable of detecting, attributing, and guiding the mitigation of explainable risk Nonconformities appears capable of offering suitable justification and compliance coverage. The timely delivery of suggestions or required modifications is expected to promote a near-zero Nonconformity process without requiring extensive user effort, thus aiding resources, improving client relationships via short timelines, and reducing costs.

II. BACKGROUND AND RELATED WORK

An understanding of current digital twin technology, explainable AI concepts, proactive compliance approaches, and their connections is necessary for the presented work. Automated manufacturing relies on many decisions from computer programs, often based on machine learning models. Therefore, AI explainability is necessary for actionable support in manufacturing systems. Digital twins enable high-fidelity models encapsulating all relevant aspects of the manufacturing process. Adequate fidelity parameters and seamless integration with the rest of the manufacturing system allow the twin to serve as a trusted source of information. Compliance with regulatory constraints is critical; a proactive approach allows regulatory risks to be managed early and effectively. Proactive design allows detection of events in advance or before they reach regulatory risk thresholds, enabling effective action—internal or regulatory—to maintain compliance.

Explainable AI delivers insights on important factors of trained ML decision support. The explainability capacity can be characterized through abstraction capabilities, explanation stability and reliability, and user-centeredness. Relevant explainability techniques are feature attribution, surrogate models, counterfactuals, causal explanation, and experiential explanations. The achievement of actionable explainability and robustness relies on a suitable combination and use of these techniques adapted to the involved stakeholders. Automated systems' regulatory compliance must thus be rigorously managed, powered by appropriate governance structures defining roles and responsibilities in risk anticipation, detection, and mitigation during operation. Automated systems' regulatory compliance must thus be rigorously managed, powered by appropriate governance structures defining roles and responsibilities in risk anticipation, detection, and mitigation during operation.

Equation 1: Core variables and notation (table)

Let a manufacturing system at time t produce sensor/process features:

$$\mathbf{x}_t \in \mathbb{R}^d$$

Digital twin state (latent/estimated "true" system state):

$$\mathbf{s}_t \in \mathbb{R}^k$$

Model suite outputs:

- Predictive model (risk or forecast):

$$\hat{y}_t = f(\mathbf{x}_t)\text{(black-box ML)}$$

- Explainable/surrogate model:

$$\tilde{y}_t = g(\mathbf{x}_t)\text{(interpretable approximation)}$$

International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

Operational "risk / compliance index" (the paper describes continuous monitoring of an index and thresholds
Explainable Machine Learning fo…
):

$$R(t) \in [0,1]$$

Thresholds:

- warning threshold T_w
- breach threshold T_b with $T_w < T_b$

2.1. Explainable AI in Industry 4.0: Six drivers account for growing interest in Explainable AI (XAI) beyond academia: ethics (ensuring fair and accountable algorithms); efficacy (avoiding catastrophic errors of opaque solutions); efficiency (providing human-friendly interfaces that augment rather than replace human capabilities); trust (detecting and mitigating undesirable behaviours); adoption (addressing regulatory and industry certification requirements); and competition (incentivizing organizations to devote resources to model interpretation and critique). Five traditional XAI themes—Feature Attribution, Model-Agnostic Explanations, Counterfactual Explanations, Causal Explanations, and Experiential Explanations—are the focus of growing research effort, particularly in risk averse sectors (e.g., aviation, autonomous vehicles, finance, healthcare).

Production Decision Support, Proactive Engineering, and Manufacturing-as-a-Service represent three types of application. Production Decision Support involves the integration of XAI techniques with predictive analytics for critical decision-making relative to short-term production scheduling, equipment maintenance, and product quality. Proactive Engineering embeds XAI in simulation models to identify potential future problems, enabling timely corrective action. Manufacturing-as-a-Service uses XAI within external control systems that coordinate production systems, potentially across different suppliers. In the case of Manufacturing-as-a-Service, the expert knowledge supplied by the digital twin of the production system is critical—data-driven AI/ML solutions would not be reliable given the high cost of failure, while the expert knowledge embedded in physics-based simulation models is natural to communicate to a human operator through explanations.

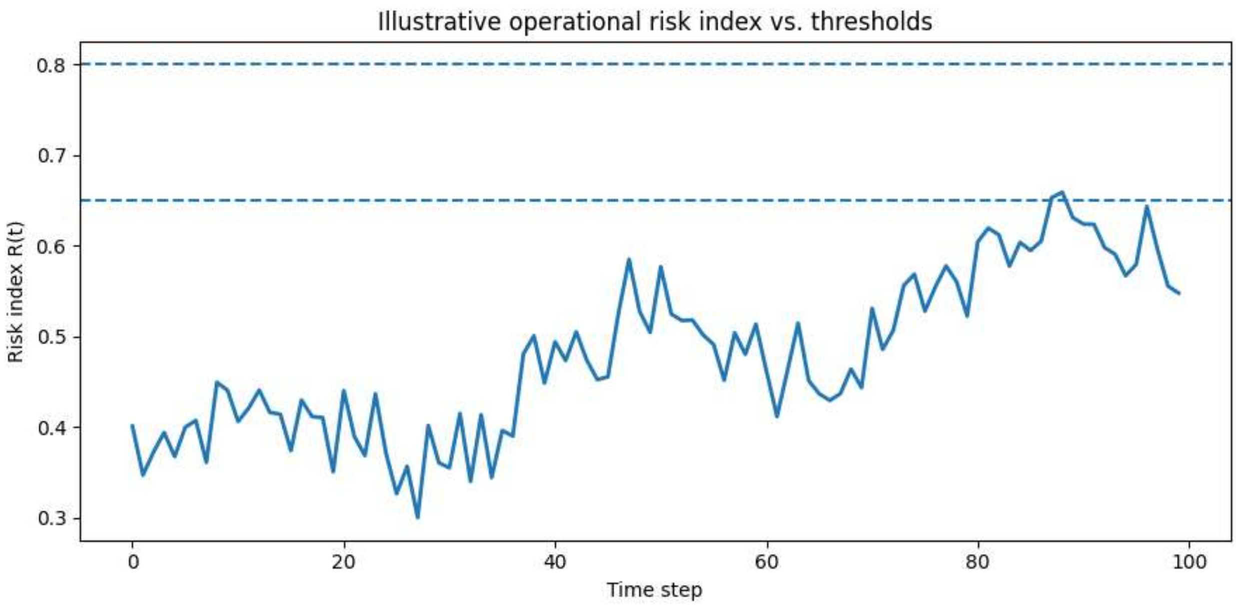

2.2. Digital Twin Technology in Manufacturing: Digital Twin technology bridges the physical and virtual worlds. Driven by data from real-world systems, a Digital Twin becomes a data-rich model—typically captured through simulation technology—of a system behaving in the real world. Within a Digital Twin, data flows into a sink, the model being replicated, where it generates predictive, prescriptive, and explanatory data streams via simulation. The adjacent sink can provide telemetry as the system is operating, or results of "what-if" analyses performed when the system is not running. The Digital Twin creates predictive digital representations of real-world systems. These are used to generate simulated digital models of functionality, efficiency, and any hidden operations.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

Much of the input information that flows into a Digital Twin begins with the system's engineering designers. Data about the environment in which the system is to operate, the materials required, the schedule for operations, and the specification of any functions to be performed are captured in data models that suit the modeling requirements for the simulation techniques being used. In many cases, this Digital Twin is a collection of models generated using different simulation methods that are best suited for parts of the system with different behaviours, for example, fluid dynamics, structural behaviour, and agent-based activity. A Digital Twin containing discrete-event simulation algorithms can receive data streams from the operation of the real system, which are used to monitor performance parameters of the operational system. As the digital models are functioning within the Digital Twin, they can safely be used to assess the functioning of the real system, with the objective of identifying problems early in their development and assessing alternatives for repair or corrective actions.

2.3. Compliance Frameworks and Auditability: Industry 4.0 promises massive increases in efficiency and productivity. Connected autonomous manufacturing systems automatically make decisions that optimize production. Auditing and compliance requirements are expected to increase, creating a paradox for the operational efficiency of Internal Revenue Service processes. Compliance problems are already apparent in cases such as the 737 MAX. Companies must adopt a proactive stance toward risk and compliance by anticipating events and applying cures before passing regulatory thresholds or suffering other serious consequences. Incorporating regulatory considerations into digital twins can help make such a proactive stance possible.

Simply put, a digital twin of a digital twin would follow digital twin principles but add regulation compliance and risk management as design objectives. Such a twin would ingest real-world data, apply predictive models, detect risk, explain the risk and its origin, and present actionable advice. Auditing, compliance, and risk management for transparency in process and compliance—all such tasks usually performed by an internal or external organization—would be implemented as software modules for Active Compliance Control. A failure to meet various conditions would trigger corrective actions: processes (production or internal) would be advised and attempts made to fill knowledge gaps. In short, the digital twin of the digital twin would automate and operationalize compliance and risk detection.

III. THEORETICAL FOUNDATIONS

Explainability methods are classified into feature attribution, surrogate models, counterfactual explanations, causal explanations, and experiential explanations. The relevant metrics include fidelity, stability, and human-centeredness. Proactive paradigms are defined with respect to risk anticipation, timing of risk intervention, and governance framework for automated decision-making systems.

Digital Twin modeling and simulation theory is examined, with attention to data model aspects and requirements for synchronization, data quality, and latency guarantees from production systems to the Digital Twin.

Explainability methods. Explainability brings together a spectrum of methods to render ML-driven decisions understandable to humans within specific contexts and for specific reasons. Feature attribution assigns a contribution score to each feature used in an ML decision, indicating its relative importance. Surrogates are interpretable models trained to approximate complex ML models, allowing insights on the underlying decision mechanism. Counterfactuals express how an input observation should differ to yield a different decision, supporting risk mitigation actions. Causal models identify and reason with the causal structure among variables to predict consequences of changes in particular variables or actions. Experiential explanations describe the influence of inputs based on a human-like experience with ML decisions. The choice of method depends on context, use case, and audience.

Metrics for explainability enumerate and elaborate on desirable properties of explanations. Fidelity measures the degree to which the explainability method preserves correctness and correctness of the original decision. Stability quantifies how similar explanations are for similar input observations. HCI principles define human-centeredness: providing explanations that are tailored to the explainee, aligned with their mental model, and usable for their intended purpose. The last perspective emphasizes that explanations should not exist in isolation, but instead support decision-making and subsequent action-taking.

Proactive compliance. Digital Twins increasingly act like self-aware entities by anticipating, detecting, reporting, and even mitigating undesirable events or risks. Close-to-real-time risk anticipation enables proactive risk intervention. It is particularly important to control the level of risk and ensure regulatory compliance in safety-critical production

International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

environments. However, the moments when intervention is most effective often cannot be treated deterministically, especially when automated systems are involved; nuanced strategic decision-making before triggering the intervention can lead to significant improvement.

Establishing a governance framework is essential to ensure joint decision-making between the human and the automated production decision systems, articulating the responsibility for sensor-based monitoring of the production system and the Digital Twin, as well as ultimately deciding when to activate the production system or intervene in a mitigative way, considering results from the Digital Twin or Cyber-Physical Systems Audit Engine.

3.1. Explainability Methods and Metrics: Explainability research is typically driven by one of two perspectives: either the aim of increasing the trust of humans in algorithmic predictions, or the instrumentation of algorithms that can be readily audited by human overseers (or both). Consequently, the principal categories of methods that address these objectives can be summarized as follows: feature attribution methods (e.g. LIME, SHAP, filters in deep neural networks), surrogate models (e.g. non-deep-ML decision trees), counterfactual perturbations (e.g. Smallwood et al. 2023) and causal analysis (e.g. Pearl and Mackenzie 2018). *Human-centered stability*—that is, the interdependence of predicted output and input variations, including noise, around those inputs—drives user trust in decision systems, especially in the case of deep learning models. *Trust-centering* methodologies perform data-driven intervention on training samples and seek to expose key input-output dependencies.

Experience-driven methods elicit users' preferred decision-making rationales and subsequently train new learned models or design explicit rules that rationally match those desired experiences; training samples satisfying the desired experiences are sought for both LIME and experience-driven designs. *Task-interpretability* identifies the ML design configuration that best exploits training outputs, and such quality measures as utility, causality, and outcome disparity are proposed to directly evaluate task-interpretability. Enables a *definition-and-process* methodology which clarifies a model's unambiguous operational context establishes a designated real-world operating process for the model and thereof defines suitable regulatory requirements governing the deployed model, ensuring that it is indeed suitable for the role it plays in a human-in-the-loop context.

Audit reliability is assured through the concepts of *fidelity* and *stability*. *Fidelity* examines the match between the right-natured decisions of an interpretable model and those of its more complexity-laden counterpart across an interpreter-selected region of interest. *Stability* measures the degree of clustering of predicted decisions of a learned model, for a user-defined proximity relation on the input feature space, in the neighborhood of any point.

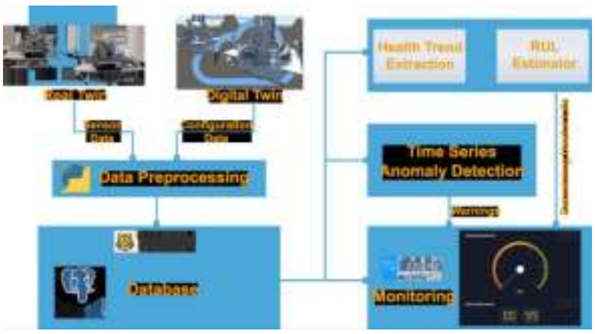

Fig 2: Explainability Methods and Metrics

3.2. Proactive Compliance Paradigms: Anticipating risks and timing interventions to safeguard operational integrity requires a governance structure to evaluate rectify the setup, execution, and results of decision-making processes. Proactive compliance ensures the compliance posture and operational integrity of regulatory systems are safeguarded in accordance with defined rules and thresholds, involving detection, attribution, and mitigation of regulatory risks before the regulated system breaches regulatory framing by a predefined margin. The deviation from compliance can be anticipated and corrected before the NMFR exceeds the defined permissible limit. Such monitoring during the operational cycle shifts focus from ex-post evaluations to ex-ante reflections of the process model, contributes to efficient auditing of high-stakes systems, retains transparency, and reduces the cost of auditability.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

Emerging technologies around AI, IoT, and digital twins are extensively used to monitor manufacturing systems addressing quality defects. While quality and operational defects are monitored by established methods, Non-Malicious Field Record (NMFR) violations remain a herculean task. Manufacturing decision-making systems rely on automated decision-making using data-driven and other trained models. These decisions can save time and improve the efficiency of the manufacturing workflow. However, noncompliance poses a hindrance to the adoption of such AI- and ML-driven decision-making systems. Auditability remains a major concern in these manufacturing systems, as these are part of a highly regulated industry. Compliance is thus an integral aspect of these systems.

3.3. Digital Twin Modeling and Simulation Theory: Digital Twin modeling and simulation underpin Digital Twin Audit Engines. That requires careful consideration of the data underlying Digital Twins: their fidelity, the types of data flows, and integration sources. When are non-physically based Digital Twins important? How can they maintain their auditability properties? What are the Digital Twin features required for a Digital Twin Audit Engine?

To support auditability, Digital Twin data models must be suitable for the audit operation and aligned with the Digital Twin Audit Engine purpose. Also, auditability beyond production verification and validation requires attention to Digital Twin data synchronization. What guarantees and quality thresholds ensure that the operating status of a manufacturing system is accurately reflected in a Digital Twin?

Why (Regulation)	**When (Trigger/time window)**	**What (rule/threshold)**	**Who (owner)**	**Where (asset/line)**	**How (data + model + explanation)**	**Evaluation**
ISO/Policy ID	Audit cycle / real-time	$R(t) \geq T_w$, etc.	Role	Machine / cell	sensors → twin → model → SHAP/LIME	Pass/Fail + notes

Table: Compliance Coverage Register (CCR): a usable table template (5W+H)

IV. ARCHITECTURE OF THE DIGITAL TWIN–BASED AUDIT ENGINE

The risk anticipation and mitigation system serves a proactive compliance audit engine, designed as a modular assembly of five components: data flows, predictive, prescriptive, and explainable models, and the audit engine itself. Interfaces external to the system are defined by the interactions and data exchanges with other components.

Data ingestion into the audit engine encompasses pipelines for monitoring data (process and product parameters, quality inspection results, and external factors) and risk data (early warning indicators from the explainable model suite). These flows incorporate stages for preprocessing, quality checks, checks on data lineage and provenance, and versioning.

The model suite consists of predictive, prescriptive, and explainable ML models. Predictive models forecast monitoring data and risk early warning indicators for the period between subsequent audits. Prescriptive models take the monitoring data for the same period as input, produce the anticipated consequences of the manufacturing process, and enable quantifying risks using the audit engine. Explainable models within the suite support risk attribution, counterfactual analysis, and validation of process corrective actions. Other explainable models enable compliance coverage checks, providing auditability support to the production system. The role, inputs, outputs, development ownership, frequency of updates, and governance structure for each model are documented and regularly reviewed.

The audit engine implements rules to determine compliance status at audit time, verifying noncompliance either with a regulation dataset or with quantified risk tolerances. Decisions and published explanations are traceable to the specific sources and models involved. Interaction with all model types enables detection of new risks, as well as early warning signals and their meanings for business continuity, and supports quantifying risks for the entire period between consecutive audits. Interpretability is addressed with specific interfaces tailored to each stakeholder group, offering aligned dashboards with relevant explanations and actionable items.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|
|Volume 11, Issue 4, July-August 2026|
DOI:10.15662/ijetemr.2026.1104003

Equation 2: Digital Twin synchronization (state update) + latency constraint
Step 1 (predict from previous state):

$$\mathbf{s}_{t|t-1} = F\mathbf{s}_{t-1} + \mathbf{u}_{t-1}$$

where:
- F is a transition operator (physics-based or learned)
- $\mathbf{u}_{t-1}$ are known inputs (commands, planned settings)

Step 2 (correct using new observations):

$$\mathbf{s}_t = \mathbf{s}_{t|t-1} + K(\mathbf{x}_t - H\mathbf{s}_{t|t-1})$$

where:
- H maps state $\mathbf{s}$ to expected observations
- K is a gain (like Kalman gain) controlling how strongly new data corrects the twin

This matches the paper's emphasis that the twin ingests real data and must remain aligned/synchronized
Explainable Machine Learning fo…
.

Step-by-step: "minimum latency" / no out-of-order audit data
The paper states an acceptance condition:
"the slack between the latest data acquisition and the auditing step is smaller than the data acquisition time"
Explainable Machine Learning fo…
Let:
- data sampling period Δt_{acq}
- latest acquisition timestamp t_{acq}
- audit execution timestamp t_{audit}
- slack $\Delta t_{\text{slack}} = t_{\text{audit}} - t_{\text{acq}}$

Constraint:

$$\Delta t_{\text{slack}} < \Delta t_{\text{acq}}$$

Why this prevents out-of-order data (step-by-step):
1. New data arrives every Δt_{acq}.
2. If the audit happens less than one acquisition period after the latest sample, then **no newer sample should exist yet** at audit time.

So the audit cannot accidentally read "future" data relative to its audit window.

4.1. System Overview and Components: The system comprises modular components that facilitate explainable decision support for proactive compliance within automated manufacturing. The component types are data, models, audit engine, and interfaces; interactions and data flow across these modules are specified. Data ingestion pipelines, encompassing source selection, quality assurance, preprocessing, lineage capture, and traceability, are described. The administrative model suite includes predictive (risk detection), prescriptive (intervention timing), and explainable (risk attribution, governance) models, along with roles, inputs, outputs, and governance. Audit engine logic, covering rule encoding, decision criteria, and resources for model-driven decisions, is detailed. Finally, interfaces for conveying explanations and information to stakeholders are defined. Components are subsequently elaborated in order.

The system comprises modular components that support explainable decision making related to proactive compliance within automated manufacturing. Four types of components can be distinguished: data, models, audit engine, and interfaces. The following subsections specify component interactions and data flow; describe data ingestion pipelines that encompass source selection, quality assurance, preprocessing, lineage capture, and traceability; catalogue the administrative model suite, which includes predictive (risk detection), prescriptive (intervention timing), and explainable (risk attribution, governance) models; explain the logic of the audit engine, which encodes rules and criteria for decision making related to automated systems; and define interfaces for conveying explanations and information to stakeholders. Components are elaborated in order.

4.2. Data Ingestion, Preprocessing, and Provenance: Realizing the system requires robust, smooth data ingestion with traceability throughout the entire journey from source to model consumption—the foundation for compliance and reliability of the powered decision support. Data ingestion happens through pipelines that ensure model input

International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

requirements are fulfilled. As previously mentioned, the digital twin can ingest various data streams originating from different sources. Given the multitude of data sources involved, the data passed to the model suite is preprocessed along dedicated paths in a functionally redundant pattern. Regardless of the precise pathway, tail end components apply quality checks dictated by the acceptance criteria of the receiving model. Where models accept only certain subsets of data, validation becomes a quality factor of the data itself; failure to conform spawns a provenance warning in the respective dashboard, regardless of model outcome.

Quality checks also guarantee that a model consuming its own prediction receives the actual predicted value, neither older nor newer. Provenance factors above that ensure the highest fidelity of risk analysis, linking detected issues to the definition and implementation of mitigating actions. All models consider time in their specification; when derived from distributed systems with momentary desynchronization, data provenance is paramount. The framework thus implements versioning, triggering warnings when a data set duplicates an existing entry. Repeated matches during the decision “Crisis detections is Human Factors Analysis” generate a human-in-the-loop request for constant supervision, as human moderators possess the highest real-time authoritative power.

4.3. Model Suite: Predictive, Prescriptive, and Explainable Models: A suite of predictive, prescriptive, and explainable models constitutes the core of the digital twin audit engine. Predictive models provide forecasts of states, properties, and events in the system being monitored, while prescriptive models recommend courses of action that anticipate and mitigate the increasing likelihood of future regulatory breaches. Both sets of models draw on system data to predict and prescribe near-term behaviour, typically over the next days or weeks. For example, ML forecasts of machining-tool wear help to avoid equipment malfunction and ensure part quality, while predictive models of scheduled equipment maintenance are used to minimize production downtime, meet delivery requirements, and maintain quality levels. A novel class of explainable models enhances reliability by attributing risk to decision-making via supervisory control in accordance with the system-understanding addressed above.

Actions taken by the digital twin audit engine are labeled human-readable, meaning that they can be understood by an operator or set of operators responsible for taking a real-world action in response to an advisory, notification, or alert from the audit engine. These actions are mapped to role-based considerations for explanation content, such as priority or severity. Such models can encompass high-fidelity digital twins focused on specific operational processes in the industrial enterprise, with an input-output behaviour sufficiently accurate for the intended application of regulatory compliance in addition to risk management. In general, the audit engine relies on a well-conceived suite of digital twins that serve to manage enterprise risk from an audit-and-certification perspective.

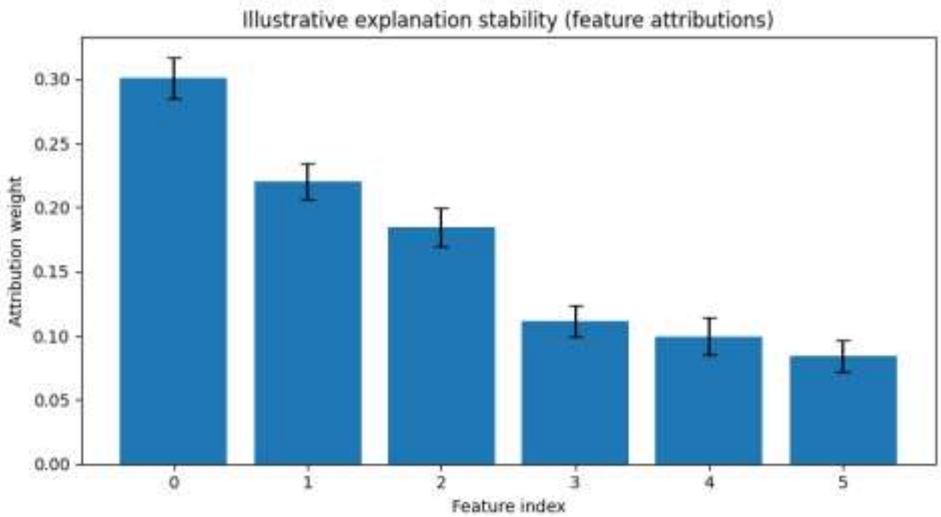

V. PROACTIVE COMPLIANCE THROUGH EXPLAINABLE ML

An audit engine based on Digital Twin technology enables risk detection, attribution, and mitigation workflows adapted to regulatory compliance. Risk alerts and corresponding explanations indicated deviations with potential regulatory repercussions. Anticipated corrective actions are formally integrated into production workflows or flagged for operator

International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

attention based on timing and risk thresholds. Traceability relates response decisions to regulatory thresholds, extracted risk explanations, and historical audit trails.

Built-in risk mitigation prescriptions or operator reminders align the audit engine with a proactive compliance approach. Alerts fire in anticipation of detected risk but maintain operator-in-the-loop control for non-automated interventions. Control system settings support automated actions, enabling operator distraction from routine processes. In this role, a pattern of real-time alerts by simulated hypotheses assists audit and decision-making. Actions on identified risk enable feedback loops to improve risk scan fidelity in the digital twin or physical manufacturing systems.

The detection and mitigation workflows support the temporal element of a proactive compliance governance structure, defined as the ability to go beyond detecting risk and provide guidance on the timing of mitigation actions. Timing is critical for follow-on operations and especially emphasized in manufacturing, where defects often cascade into costly and catastrophic later stages. Deadlines serve regulatory reporting requirements and introduce controls on response times and responsibilities, scaled by risk levels.

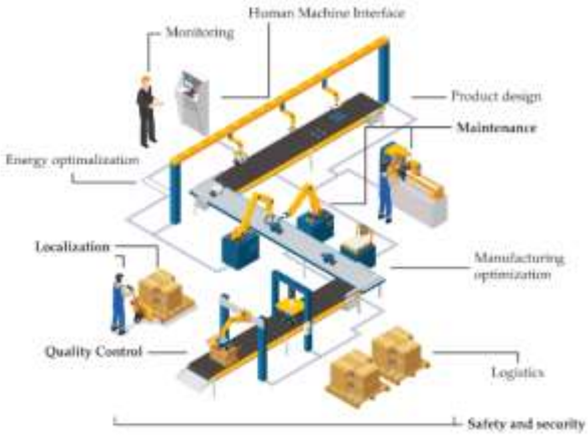

Fig 3: Applied Machine Learning for IIoT and Smart Production

5.1. Risk Detection, Attribution, and Mitigation: Three workflows support proactive compliance operations: detection, attribution, and mitigation. The detection workflow continuously monitors an operational index. When the index exceeds predefined thresholds, the attribution workflow investigates whether the exceeding source can be linked to a source of risk defined in the twin audit engine. When an exceeding condition is detected and attributed to a source of risk, the risk mitigation workflow produces recommendations on how to deal with the event and informs the action owners. These actions are routinely executed when the risk mitigation workflow output is defined to be in full agreement with the production workflow owner. Once executed, the results are analyzed, and the analysis is fed back to the Audit twin. The digital twin is therefore a living model continuously growing, improving, and maintaining the overall production system in a compliant risk-mitigated mode.

These operating procedures naturally establish the link to regulatory compliance without the classifier being task-oriented. The audit twin GPI is able to define specific conditions that, if violated, indicate a potential risk of regulation violation. An exceeding condition thus speaks for itself: when such a condition is detected, it is merely showing that the automated system is approaching a region where the regulations will clearly be violated. The position of such regions is influenced by regulatory limits captured inside the audit engine. An exceeding condition, if not attended to, indicates an anticipated breach of those specific limits.

5.2. Proactive Interventions in Automated Manufacturing: Intervention decisions arising from compliance audits are vital for maintaining efficient audit trails while mitigating breaches through preemptive actions. The following considerations inform intervention logic.

Manufacturing workflows are structured as business process models (BPMs) for automation. Formal definition and semi-automated creation of BPMs enable online disambiguation and integration with an audit engine. Installed BPMs must

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

remain consistent with running systems: a BPM can support automated compliance checking only if runtime parameters/variables match those of the process specified in the model.

Regulatory thresholds for risk detection are encoded in rule-based format within the audit engine. Risk attacks can therefore be detected and attributed using suitable predictive ML models. When risk exceeds a specified threshold for embracing the risk, anticipatory action can be undertaken. Support for such intervention is naturally aligned with a risk-based approach to compliance. Action timing is determined according to a timing model defined in the audit engine. An intervention may be automatically triggered within a workflow, but whether it is auto-executed or flagged for human approval must consider previous action history, especially human-in-the-loop aspects and previous (dis)approvals for similar actions.

5.3. Feedback Loops to Digital Twin and Control Systems: Feedback mechanisms guide proactive interventions in manufactured workflows, shaping corrective actions, timing, automation triggers, and human engagement. Auditing anticipated risks and defining the nature and timing of suitable mitigation measures lay the groundwork for directing compliance, quality, and safety interventions. Noncompliance events or near misses recorded by dedicated monitor models provide feedback to the digital twin, facilitating the risk-specific tuning of underlying predictive and prescriptive models in the model governance process. Model improvements in the digital twin database then complement underlying automation control systems, driving proactive mitigation by controlling systems, operator alerts, or operator directive changes.

Mitigation model feedback loops strive to minimize the cost of near misses and noncompliance events. Recorded near misses support a practical trade-off between the cost of intervention timing and the benefit of avoiding severe consequences, while corrective actions and their consequences (beneficial or detrimental) inform near-miss detection models. Noncompliance events guide the reassessment of mitigation timing and action feasibility, while their operational consequences direct risk-aware decision-making models.

VI. DIGITAL TWIN REALIZATION FOR AUDIT PURPOSES

Machine Learning (ML) models driving decisions in automated environments are subject to higher-order governance and may need to comply with external rules—formal or informal—that assume they have certain desirable properties. These properties must be checked over time. For a Twin Audit Engine to be valuable, the underlying models would ideally be Correct by Construction (CbC) or Free from Critical Issues (FCI). However, to achieve these ideal conditions without external data or profound lead times would be exceedingly costly. As hard guarantees cannot be retained, some form of semi-automated testing must be arranged to provide a level of conditional trust. The Twin Data Systems Occasion/Temporal Synchronization Plan outlined above hints at the potential support the Twin Digital Model could offer.

The Twin runs a simulation that takes the Digital Twin current states and the Supply Chain Connection’s planned or agreed inputs to produce predictions and suggest actions for every operational or functional model at every data-collection frequency. Some of these predictions may be procured internally and be of the same absolute or relative quality order as the operational data. Others may be irrelevant or need replacement by either regular or ad hoc human inputs. Such ad hoc inputs may relate to higher-frequency operational or functional decisions that are beyond the upper effective capability thresholds of the operational models, both from an absolute and relative standpoint. Machine Learning (ML) models driving decisions in automated environments are subject to higher-order governance and may need to comply with external rules—formal or informal—that assume they have certain desirable properties. These properties must be checked over time. For a Twin Audit Engine to be valuable, the underlying models would ideally be Correct by Construction (CbC) or Free from Critical Issues (FCI). However, to achieve these ideal conditions without external data or profound lead times would be exceedingly costly. As hard guarantees cannot be retained, some form of semi-automated testing must be arranged to provide a level of conditional trust. The Twin Data Systems Occasion/Temporal Synchronization Plan outlined above hints at the potential support the Twin Digital Model could offer.

The Twin runs a simulation that takes the Digital Twin current states and the Supply Chain Connection’s planned or agreed inputs to produce predictions and suggest actions for every operational or functional model at every data-collection frequency. Some of these predictions may be procured internally and be of the same absolute or relative quality order as the operational data. Others may be irrelevant or need replacement by either regular or ad hoc human inputs. Such ad hoc

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|
|Volume 11, Issue 4, July-August 2026|
DOI:10.15662/ijetemr.2026.1104003

inputs may relate to higher-frequency operational or functional decisions that are beyond the upper effective capability thresholds of the operational models, both from an absolute and relative standpoint.

6.1. Twin Data Synchronization and Fidelity: Data from production systems are continuously ingested into the digital twin and aligned with execution commands and control system outputs. Available safeguards are checked against current operating conditions, and the data are subsequently formatted for all models in the model suite. Data preparation for predictive models includes outlier detection and flagging, data imputation, and distribution calibration. Data preparation for explainable models addresses data quality, lineage, and provenance management, including checks to ensure minimum stability and representativeness thresholds for sound interpretation. Data quality and representativeness checks are also performed on data streams feeding the audit engine. Data from completed operations and activities are continuously synchronized in a way that enables genuine twin data testing of all models while ensuring minimum data latency.

Fulfilling these requirements provides compliance with the regulatory need for monitoring twin data quality and (sufficiently) ensuring minimum twin-data-latency guarantees. Constant monitoring reduces the risk that remaining discrepancies, uncertainties, and latency issues within the twin may compromise resilience to regulatory compliance, especially in high-risk situations. Nevertheless, toxicity-related aspects might still be less monitored because their consideration cannot be automated before the expert evaluation.

Equation 3: Proactive compliance index $R(t)$ and threshold exceedance logic
Explainable Machine Learning fo…
Suppose you have m normalized risk indicators (tool wear, defect rate, cyber anomaly score, etc.):

$$r_i(t) \in [0,1], i = 1, \dots, m$$

A standard weighted index:
Step 1 (choose weights):

$$w_i \geq 0, \sum_{i=1}^{m} w_i = 1$$

Step 2 (combine indicators):

$$R(t) = \sum_{i=1}^{m} w_i \, r_i(t)$$

This creates a single monitored operational index consistent with the article's "operational index" concept.
Define warning/breach events:

$$\text{Warn}(t) = \mathbb{1}[R(t) \geq T_w]$$
$$\text{Breach}(t) = \mathbb{1}[R(t) \geq T_b]$$

Where $\mathbb{1}[\cdot]$ is 1 if condition is true else 0.
If a predictive model forecasts the index H steps ahead:

$$\hat{R}(t+h), h = 1, \dots, H$$

Define predicted first breach time:

$$\tau_b(t) = \min\{h \in \{1, \dots, H\}: \hat{R}(t+h) \geq T_b\}$$

If no future breach in horizon, set $\tau_b(t) = \infty$.
This formalizes the paper's "timing of mitigation actions" emphasis
Explainable Machine Learning fo…

6.2. Simulation-Based Verification and Validation: A simulation-based methodology verifies functional requirements and validates the ability to assess compliance with regulatory thresholds. The audit engine runs a Digital Twin application and is tested against risk-relevant scenarios. Data flows are integrated and synchronized at simulation run-time. Audit results are interpreted for the predefined regulations—CMMI level and metal shutter cycle. Four tests focus on regulatory compliance and general audit operation: metal-shutter cycle time and CMMI rating lower bounds (delivered by the

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

predictive model), the shutter position (prescriptively controlled), and shutter cycle time outside the allowed range. In contrast, a stress-test scenario concentrates on the active process of risk detection and mitigation. Here, the lurking risk of a zero CMMI level is addressed by mechanism activation. A smear in the data for quality thickness shows how the audit engine reacts to indirectly controlled systems.

Acceptance conditions evaluate that all essential regulatory properties have been tested at least once. Independently, the slack between the latest data acquisition and the auditing step is smaller than the data acquisition time, thus preventing out-of-order data during audits. All other functional requirements are covered. Further tests examine multiple risk factors (with their limits) simultaneously and under stress or high-intensity conditions.

6.3. Scenario Testing and Stress Scenarios: Proactive compliance hinges on a robust capacity to anticipate risk and recommend timely interventions. While most interventions can be executed autonomously, regulatory thresholds may require a human-in-the-loop decision process; in these cases, the model triggers the relevant operator actions but does not execute them. The context of the underlying Proactive Compliance Paradigm frames the planning requirements, specifying the timing of model-driven actions based on min-max thresholds defined by the control system or the operator.

Stress scenarios ascertain whether the governing models safeguard the production process against edge case malfunctions; these cases are not intended to challenge model fidelity but to disrupt the operational regime and test risk prediction. The stress testing logic examines the regulatory compliance of the analysis and suggests additional tests to trigger under particular circumstances, such as during mode transitions; it adds to the risk anticipation pipeline by investigating non-ML sources of risk exposure and anticipating the emergence of previously undetected risks. The stress testing framework is embedded in the analysis data model, which also orchestrates simulation-based validation and verification; the latter evaluates model fidelity, resilience under non-ML failure modes, and environmental resilience against cyber-physical attacks.

Additional scenario tests assess whether the production system complies with the safety integrity level levels stipulated in ISO 26262 when operating under safety-related conditions.

VII. EVALUATION AND VALIDATION METHODOLOGIES

A framework to assess explainability—encompassing aspects such as evaluation metrics, human-centered design, and auditor usability—is proposed. Coverage of regulatory compliance, including alignment with existing rules and the fulfilment of auditability and traceability needs, is examined. Empirical validation strategies—spanning both digital-twin pilot studies and interviews with industry stakeholders—are outlined.

The explanation of an AI model is of utmost importance, but is often not ensured by the approach itself—the reliability of similarly interpretable machine learning (IML) methods lies solely in the underlying model. The presence of IML does not remove the need for the use case to be well evidenced nor for the economic impact, safety, and ethics of the system to be considered. As IML remains a very active area of research, there are still many questions surrounding the adequacy of different explanation types for different ML uses cases in different contexts. For high-stakes domains, such as automated manufacturing, adequate scrutiny remains necessary. Thus, a method not only asking for IML but why the decision made by the AI model is the right one is needed. Such a meta-evaluation allows users of the system to assess the extent to which the IML provided fulfils their requirements and appropriately addresses their needs.

7.1. Explainability Evaluation Frameworks: At a high level, the proposed Explainability Assessment Framework considers three dimensions that govern the explainability of a model and its utility for an intended stakeholder: required explainability for a specific use (or application) context, achieved explainability defined by the capabilities of the model, and the role of experts/humans in the loop who validate and consume the model explanations. For a specific application context (for example, Compliance Verification) requiring a specific type of explanation (Compliance Coverage), the interaction among achieved explainability, required explainability and usability of the explanations determines the overall explainability assessment. In detail, the ground-up approach first considers how explainability needs for the explained model depend on the intended application context and the relevant performances of the model. Subsequently, it aims to assess the explanation per se, when evaluable, and lastly the benefit of humans/expert in the usability of the explainability layer.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

At a deeper level, Explanation Quality Evaluation Metrics are classified as General Explanation Reliability factors (reputation of the model and model explanation source) and Information Content metrics (Information Gain, Complexity, Calculable & Visualizable, Privacy-Relevance, and External-Consistency). Finally, Auditor-Factor Usability encompasses several Human-explained and Human-factors Consideration metrics, enriched with definition of Training & Testing and Explainability Evaluation datasets, ability to Visualize and Manipulate Explanation Structure, Usability Testing, Type-Specific Qualitative Assessment (for example, Specific Domain Considerations fit for Specific Domains) and Consumer Proficiency.

The Compliance Coverage Assessment Framework utilizes an integrated 5W+H analysis methodology to construct the Compliance Coverage Register (CCR), which links the relevant model-information with the jurisprudential-compliance-sources from the applicable legislation. The CCR considers separate tables for the Why & When (covering regulation-resources source), What, Who, Where & How, and embeds an Evaluation column for on-demand validation of the compliance status.

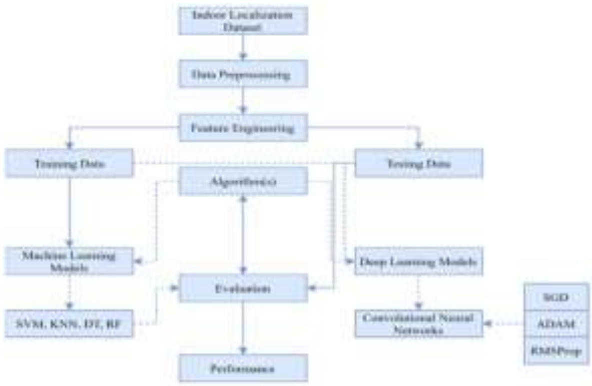

Fig 4: Digital Twin Framework

7.2. Compliance Coverage Assessment: Compliance Coverage Assessment evaluates coverage of the proactive compliance paradigm and provides a method to identify compliance-related questions that could be answered by the audit engine. Proactive compliance supports the anticipation of compliance breaches, detects risk and vulnerability, determines the impact of risk on compliance-related thresholds, and informs the timing of preventive measures or audits. As identified conduct and operational risk management constitutes two specific areas of the proactive compliance paradigm. Auditors require a set of questions that address the risk of the automated systems generated or triggered by the audit engine.

Such questions could be expressed in terms of the underlying production systems and linked to theory-based or empirical models. A question should specify a risk, the associated consequence in a production system, how that consequence would violate an established compliance threshold and when compliance would be breached. The resulting question set can be converted to an audit engine rule set and subsequently evaluated against. Questions not yet covered can be included in the rule set for future iterations of the audit engine.

7.3. Empirical Validation in Pilot Studies: Empirical validation hinges on pilot studies or experiments replicating, or closely mirroring, real-world settings. Sourcing suitable data is critical, driven largely by availability, privacy, and ethical concerns. Automation and digitization fuel growing volumes of operational and non-operational data, including risk-management data generated by increasingly sophisticated incorporated levels of AI and machine-learning technologies. Data originating from sensors monitoring production equipment and associated production control systems reporting production performance and environmental data, and business-information systems supporting the role of specialist departments—such as health-and-safety, quality-control, or surveillance—are particularly relevant.

Enhancing compliance-management processes—detection, management, and mitigation of compliance violations—using production data of these forms remains a priority, especially where decision support relies heavily on machine learning. Formalized data governance in terms of quality, lineage, provenance, and versioning dynamically controls these data sources during exploitation. Requirements are defined for the purpose of detection and mitigation of violations, with clear threshold indicators traced to compliance rules. Product-design data monitoring controlled levels of production

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

resources, for example, should not rely heavily on ML-based predictions and a lack of model transparency. Business considerations typically govern risk-assessment rule definition and support implementation of compliance-management processes, including data selection, validation, and business justification. Statistical significance of these choice decisions can be included in enhanced formal analysis and audit processes.

VIII. CONCLUSION

Machine Learning (ML) markets will grow by 16.70% annually until 2029. Users increasingly rely on ML for decision-making, yet decisions lack transparency and trust. In manufacturing, ML underpins Digital Twins, virtual representations of real production systems with Digital Twin Audits serving as virtual-compliance checks. These audits draw on explainable AI (XAI)—open ML decision-making—to expand the ML-Decision Auditing Engine. Addressing regulatory compliance, a proactive-auditing structure detects, attributes, mitigates, and audits regulatory-risk exceeding ML decisions. A supporting-data model specifies data pipelines, Global Compliance-coverage recommendations stub data stores, model-portfolio descriptions govern ML audit models, audit-engine logic encodes functions, and interfaces deliver decisions, explanations, and actionables for stakeholders.

Emerging trends in manufacturing and risk management are poised to reshape the segments. Ascent to self-governing risk anticipates potentially disruptive thresholds, yielding Need for Speed—auditing along the roadmap. Automated production—control opening up shifting finale interpretation—is Pollution at Source; machines doing the job eliminates unintended side effects, decoupling cost and pollution variance. Previous focus on predictive all-quitting all-impaired underscores the industrial shift to anticipatory. Natural disasters cannot be eliminated but their impact can be negated. Industrialization can manage risk down to design, building tolerance into service-provision systems. Industry will drive compliance developmental cycles, from pilots through standardization. Machine Learning—driven Data-Driven Business—relies on explainable AI from Decision-to-Decision, Acting-to-Outcome, Why-to-What, behind-and-front of Machine Learning, shaping Human-Centered Machine Learning—wrapper forth and interior.

8.1. Emerging Trends: A growing number of commercial Explainable AI solutions—including those embedded in off-the-shelf ML software—are making their way into niche applications in manufacturing (e.g., predictive maintenance, quality prediction, utility consumption) and more broadly in areas such as medicine, autonomous driving, finance, and insurance. Nonetheless, far-sighted companies are considering next-generation systems that automate other parts of the decision process; for instance, they are deploying (or developing) autonomous agents that can continually decide on any combination of resources, suppliers, products, and services.

These systems are able to conceivably address a broader set of goals and business needs, including regulatory and industry compliance. Achieving the same level of explainability for the ML components of those systems by using out-of-the-box solutions proves difficult in practice: they may not match the full requirements set by regulatory bodies beyond residing within an acceptable region of the explainability measure. DST—defined as an extension, within a set of prescribed operational conditions, of a Digital Twin supported by Digital Thread capabilities—provides additional capabilities, layers, and components to allow these decisions to be traced back, and the corresponding prediction and action suites (what-if analysis) to be audited and approved by a human before deployment.

REFERENCES

1. Kannan, S. (2026). Generative AI for Smart Agricultural Equipment Automation and Pharmaceutical Crop Yield Prediction. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 9(3), 1068-1074.
2. 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.
3. Krishnan, M., Nandan, B. P., Rongali, S. K., Meda, R., Kalisetty, S., & Singireddy, J. (2026, June). AI-Driven Data Engineering and Predictive Analytics Framework for Semiconductor Supply Chain Optimization and Digital Infrastructure Modernization. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1-6). IEEE.
4. Pandiri, L. (2026). AI-Powered Predictive Analytics for Climate-Aware Flood and Mobile Home Insurance Risk Management. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 834-847.
5. Kummari, D. N., Kaulwars, P. K., & Gadi, A. L. (2026). Decision-Making in Industrial Robotics and Manufacturing Plants. Artificial Intelligence: Theory and Applications: Proceedings of AITA 2025, Volume 3, 3, 332.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

6. Kumar, S. S., Gadi, A. L., Sheelam, G. K., Kummari, D. N., Koppolu, H. K. R., & Pamisetty, A. (2026, March). AI-Driven Compliance and Audit Framework for Manufacturing Infrastructure in Automotive Connected Services and Financial Ecosystems. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1-6). IEEE.
7. Singireddy, S. (2026). Autonomous Insurance Analytics Using Agentic AI for Personalized Coverage and Predictive Risk Intelligence. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 848-861.
8. Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. IEEE Access, 8, 108952–108971.
9. Hoffmann, R., & Reich, C. (2023). A systematic literature review on artificial intelligence and explainable artificial intelligence for visual quality assurance in manufacturing. Electronics, 12(22), 4572.
10. Bitencourt, J., Wooley, A., & Harris, G. (2025). Verification and validation of digital twins: A systematic literature review for manufacturing applications. International Journal of Production Research, 63(1), 342–370.
11. Kolla, S. H., Nagabhyru, K. C., & Kummari, D. N. (2026). Letter to the Editor re:"CurvAssist: An AI assisted pipeline for penile shaft segmentation/curvature measurement in children with hypospadias". Journal of Pediatric Urology.
12. Pamisetty, V. (2026). Explainable Agentic AI for Secure and Adaptive Supply Chain Decision Intelligence in Food Service and Financial Systems. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 862-876.
13. Karabulut, E., Pileggi, S. F., Groth, P., & Degeler, V. (2023). Ontologies in digital twins: A systematic literature review. Journal of Industrial Information Integration.
14. Inala, R., Sheelam, G. K., Aitha, A. R., Lakshmi, A. U., Nagabhyru, K. C., & Segireddy, A. R. (2026, March). Architecting Hybrid Data Products Using AI/ML and Agentic AI for Group Insurance and Retirement Solution Platforms with Advanced Data Governance. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
15. Kummari, D. N., Aitha, A. R., Kumar, M. V. K., Pandiri, L., Gottimukkala, V. R. R., & Nagubandi, A. R. (2026, March). Real-Time AI-Driven Audit and Compliance Framework for Smart Manufacturing Financial Workflow. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
16. Nandan, B. P. (2026). Cloud-Scale Data Engineering for Real-Time Semiconductor Testing and AI-Powered Chip Diagnostics. International Journal of Computer Technology and Electronics Communication, 9(3), 1058-1070.
17. Sanku, R., Recharla, M., BG, M. B., & Chikop, S. A. (2026, February). Privacy-Preserving Federated Learning Framework with Adaptive Aggregated Gradient Perturbation for IoT Healthcare Systems. In 2026 3rd International Conference on Integrated Intelligence and Communication Systems (ICIICS) (pp. 1-6). IEEE.
18. Vankayalapati, R. K., Polineni, T. N. S., Ahammad, S. H., Pandugula, C., & Selvan, R. S. (2026). IoT-Enabled Augmented Reality for Real-Time Equipment Diagnosis. In Virtual Reality, Real Emergency (pp. 104-121). CRC Press.
19. Moosavi, S., Farajzadeh-Zanjani, M., Razavi-Far, R., Palade, V., & Saif, M. (2024). Explainable AI in manufacturing and industrial cyber-physical systems: A survey. Electronics, 13(17), 3497.
20. Inala, R. (2026). Cloud-Native AI and MDM Framework for Next-Generation Insurance and Retirement Data Products. International Journal of Engineering & Extended Technologies Research (IJEETR), 8(3), 5050-5063.
21. Nagabhyru, K. C., Gadi, A. L., Seenu, A., Davuluri, P. N., Segireddy, A. R., & Pamisetty, V. (2026, March). Towards Automated Financial Risk Scoring in Automotive Financing with Explainable Machine Learning. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
22. Davuluri, P. N., Segireddy, A. R., & Sheelam, G. K. (2026). Comment on" Machine learning-based prediction of CAC-defined cardiovascular risk using routine health examination data: a retrospective cross-sectional study in a Taiwanese population". Journal of the Formosan Medical Association= Taiwan yi zhi, S0929-6646.
23. Vankayalapati, R. K., Polineni, T. N. S., Ahammad, S. H., Pandugula, C., & Selvan, R. S. (2026). IoT-Enabled Augmented Reality for Real-Time Equipment Diagnosis. In Virtual Reality, Real Emergency (pp. 104-121). CRC Press.
24. Garapati, R. S., Aitha, A. R., & Singireddy, S. (2026). Secure Cloud Architecture for Privacy-Preserving Machine Learning on Electronic Health Records with Web-Based Analytics Tools. In International Conference on Intelligent Human Computer Interaction (pp. 407-417). Springer, Cham.
25. Onaji, I., Tiwari, D., Soulatiantork, P., Song, B., & Tiwari, A. (2022). Digital twin in manufacturing: Conceptual framework and case studies. International Journal of Computer Integrated Manufacturing, 35(8), 831–858.
26. Peruthambi, V., Singireddy, S., Kummari, D. N., Nandan, B. P., Pamisetty, V., & Amistapuram, K. (2026, June). Adaptive Security Framework for AI-Driven Risk Assessment and Compliance Automation in Cloud Security and Vulnerability Management within the Insurance Domain. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1-6). IEEE.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

27. Lattanzi, L., & Rossi, M. (2021). Digital twin for smart manufacturing: A review of concepts towards a practical industrial implementation. International Journal of Computer Integrated Manufacturing, 34(6), 567–597.
28. Genga, L., & Winter, K. (2025). Artificial intelligence in conformance checking: State of the art and research agenda. Process Science, 2, Article 9.
29. Liu, S., Zheng, P., & Bao, J. (2024). Digital twin-based manufacturing system: A survey based on a novel reference model. Journal of Intelligent Manufacturing, 35(6), 2517–2546.
30. 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.
31. Nandan, B. P., & Tummoju, S. T. (2026, February). Attention-Enhanced Image Processing for Mobile Augmented Reality in Real-Time Object Monitoring. In 2026 3rd International Conference on Integrated Intelligence and Communication Systems (ICIICS) (pp. 1-7). IEEE.
32. Meda, R. (2026). Agentic AI-Driven Smart Supply Chain Orchestration for Paint Manufacturing and Retail Ecosystems. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 916-930.
33. Jiang, Y., Yin, S., & Kaynak, O. (2020). Data-driven monitoring and safety control of industrial cyber-physical systems: Basics and beyond. IEEE Transactions on Industrial Electronics, 67(6), 4734–4746.
34. Bae, C., Choi, E., & Lee, S. (2025). Technologies, applications, and challenges of digital twin across industries: A systematic review of the state-of-the-art literature. IEEE Access.
35. 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.
36. Madhavi, K. R., Gottimukkala, V. R. R., Pandiri, L., Sriram, H. K., Malempati, M., & Adusupalli, B. (2025, November). Hybrid Transformer–Federated Learning Model for Secure Release Engineering in Global Payment Networks. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
37. Suura, S. R. (2026). Deep Learning-Enabled Cell-Free DNA Analytics for Precision Reproductive and Preventive Healthcare. International Journal of Science, Research and Technology, 9(3), 801-815.
38. Dunzer, S., Zelt, S., Matzner, M., & Rinderle-Ma, S. (2019). Conformance checking: A state-of-the-art literature review. Business & Information Systems Engineering, 61(6), 607–628.
39. Cimino, C., Negri, E., & Fumagalli, L. (2021). Review of digital twin applications in manufacturing. Computers in Industry, 123, 103297.
40. Davuluri, P. N., Segireddy, A. R., & Sheelam, G. K. (2026). Comment on "Interpretable machine learning model for predicting refeeding syndrome after colorectal cancer surgery". Clinical Nutrition ESPEN, 75.
41. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on "Predicting self-image satisfaction after adult spinal deformity surgery: a machine learning approach using patient phenotypes". Spine Deformity, 1-3.
42. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 33–44.
43. Singireddy, J., & Sheelam, G. K. (2026). Generative AI Models for Process Optimization in Semiconductor Wafer Design and Yield. Artificial Intelligence: Theory and Applications: Proceedings of AITA 2025, Volume 3, 3, 220.
44. Seenu, A., Aitha, A. R., Gottimukkala, V. R. R., Singireddy, J., Meda, R., & Garapati, R. S. (2025, November). Hybrid Multi-Agent Reinforcement Learning and Blockchain Framework for Real-Time Transaction Integrity in Cloud-Driven Financial Systems. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
45. 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.
46. Rizzi, W., Comuzzi, M., Di Francescomarino, C., Ghidini, C., Lee, S., Maggi, F. M., & Nolte, A. (2024). Explainable predictive process monitoring: A user evaluation. Process Science, 1, Article 3.
47. 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.
48. Maguluri, K. K. (2026). Cloud-Integrated Machine Learning System for Ebola Virus Disease Prediction and Epidemic Intelligence in Smart Healthcare Systems. Journal of Advances in Management, Engineering and Science (JAMES), 1(03), 1-9.
49. Mahadevan, S. (2024). Clouding the Future: Innovating Towards Net-Zero Emissions. International Journal of Computing and Engineering, 6(2), 17-23.
50. Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., & Müller, K.-R. (Eds.). (2019). Explainable AI: Interpreting, explaining and visualizing deep learning. Springer.

51. Gadi, A. L., Anitha, S., Basha, N., & Kapilas, D. (2026, July). Enhancing Smart Grid Privacy with Adaptive Fog-Based Differential Privacy Models. In Signal Processing, Telecommunication & Embedded Systems: Automation and Sustainability Applications: Proceedings of Tenth International Conference on Microelectronics Electromagnetics and Telecommunications (ICMEET 2025), Volume 4 (Vol. 4, p. 344). Springer Nature.
52. Soori, M., Arezoo, B., & Dastres, R. (2023). Digital twin for smart manufacturing: A review. Sustainable Manufacturing and Service Economics, 2, 100017.
53. Pamisetty, A. (2026). The Impact of Climate Change on Coastal Ecosystems. SCIENTIFIC CULTURE, 12(1), 1-11.
54. Recharla, M., Garapati, R. S., Bandi, V. D. V. K., Yandamuri, U. S., & Mangalampalli, B. M. (2026, March). Generative AI-Enhanced Data Engineering Pipelines for Predictive Biomarker Discovery in Alzheimer's and Kidney Disease. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
55. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. Int. J. Innov. Sci. Res, 9, 593-595.
56. Madhavi, K. R., Rongali, S. K., Polineni, T. N. S., Kummari, D. N., Challa, K., & Challa, S. R. (2026, February). Explainable AI (XAI)-Driven Predictive Analytics Framework for Ethical and Scalable Automation in Cloud-Native Architectures with Enterprise and Healthcare Interoperability. In 2026 International Conference on Electronics and Renewable Systems (ICEARS) (pp. 31-36). IEEE.
57. PSL, N. D., Segireddy, A. R., & Sheelam, G. K. (2026). Comment on" Effectiveness of AI-assisted ESI triage on accuracy and selected outcomes in emergency nursing: A systematic review". International emergency nursing, 87, 101846-101846.
58. Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. Harvard Journal of Law & Technology, 31(2), 841–887
59. Recharla, M., Arockiaraj, M. C., Kamalakkannan, D., & Gamini, P. (2026, April). An IoT–WSN–Machine Learning and Cloud Integrated Framework for Real-Time Smart Irrigation and Crop Health Monitoring. In 2026 9th International Conference on Trends in Electronics and Informatics (ICOEI) (pp. 1521-1526). IEEE.
60. Komaragiri, V. B. (2026). Harnessing AI Neural Networks and Generative AI for the Evolution of Digital Inclusion: Transformative Approaches to Bridging the Global Connectivity Divide. Transcriptome, 1(1), 13-21.
61. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on "From laboratory promise to decision-grade practice: strengthening reproducibility and casework relevance in AI-assisted bloodstain ageing". Forensic Science, Medicine and Pathology, 1-2.
62. Rao, S., Annapareddy, V. N., Sriram, H. K., Kannan, S., & Komaragiri, V. B. (2026, January). The Role of Cloud Computing in Scalable Solar Power Infrastructure: Ensuring Reliability Through AI and ML-Based Grid Management. In Smart Computing Paradigms: Human-Centric Systems for Sustainable Development: Proceedings of Seventh International Conference on Smart Computing and Informatics (SCI 2025), Volume 4 (Vol. 4, p. 295). Springer Nature.
63. Pamisetty, A. (2026). Cloud-Native Big Data Architecture for Real-Time Tax Analytics, Fraud Detection, and Fiscal Governance. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 902-915.
64. Corallo, A., Lazoi, M., & Lezzi, M. (2021). Shop floor digital twin in smart manufacturing: A systematic literature review. Sustainability, 13(23), 12987.
65. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115.
66. Rajamanickam, V., Singireddy, S., Davuluri, P. N., Sheelam, G. K., Aitha, A. R., & Vakkalagadda, T. (2026). AI-Enabled Visual Evidence Intelligence for Detecting Manipulated Digital Media. International Journal of Special Education, 41(14s), 115-124.
67. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. Int. J. Innov. Sci. Res, 9, 593-595.
68. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.
69. Zhang, R., Wang, F., Cai, J., Wang, Y., Guo, H., & Zheng, J. (2023). Digital twin and its applications: A survey. The International Journal of Advanced Manufacturing Technology, 127, 1–22.
70. Rao, S., Annapareddy, V. N., Sriram, H. K., Kannan, S., & Komaragiri, V. B. (2026, January). The Role of Cloud Computing in Scalable Solar Power Infrastructure: Ensuring Reliability Through AI and ML-Based Grid Management. In Smart Computing Paradigms: Human-Centric Systems for Sustainable Development: Proceedings of Seventh International Conference on Smart Computing and Informatics (SCI 2025), Volume 4 (Vol. 4, p. 295). Springer Nature.
71. Kumar, M. V. K., Kolla, S. H., Pamisetty, V., Pandiri, L., Yandamuri, U. S., & Valiki, D. (2026, May). Enterprise-Scale Generative AI Agents for Secure and Governed Automation in Insurance and Public Financial Management. In

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104003

2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE) (pp. 1-6). IEEE.

72. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., & Gebru, T. (2019). Model cards for model reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency, 220–229.

73. Meijer, A., Lorenz, L., & Wieringa, R. (2025). Explainable and trustworthy artificial intelligence for industrial decision support: Emerging challenges and research directions. Artificial Intelligence Review.

74. Singireddy, J. (2026). Generative AI for Autonomous Accounting and Tax Filing: A Cloud-Driven Framework for Smart Financial Services. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 889-901.

75. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv.

76. Zhu, Y., Cheng, J., Liu, Z., Cheng, Q., Zou, X., Xu, H., Wang, Y., & Tao, F. (2023). Production logistics digital twins: Research profiling, application, challenges and opportunities. Robotics and Computer-Integrated Manufacturing, 83, 102592.

77. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).

78. Annapareddy, V. N. (2026). AI-Driven Cloud-Native Solar Energy Intelligence Platform for Smart Educational IT Ecosystems. International Journal of Research and Applied Innovations, 9(3), 597-609.