



## Infrastructure-Level Intelligence: Embedding AI into Data Movement and Validation Layers

Bhargavaram Potharaju\*

University of Bridgeport, USA

\* Corresponding Author Email: [bhargavarampotharaju@gmail.com](mailto:bhargavarampotharaju@gmail.com) - ORCID: 0000-0002-0047-9550

### Article Info:

DOI: 10.22399/ijcesen.4944  
Received : 10 February 2026  
Revised : 18 February 2026  
Accepted : 21 February 2026

### Keywords

Infrastructure Intelligence,  
AI-Driven Data Validation,  
Intelligent Data Pipelines,  
Platform Engineering,  
Compliance Automation

### Abstract:

Data volumes and complexity push enterprise and governance systems beyond their design assumptions. Rule-based data validation techniques can no longer appropriately model modern data ecosystems. Infrastructure-Level Intelligence (ILI) refers to applying AI to data movement and validation infrastructure, rather than application-level data validation. ILI aims to help quickly spot problems, adjust data validation rules on the fly, fix data movement automatically, improve data quality as needed, and provide other features by processing data close to where The implementations of such compliance-driven use cases have led to huge improvements in these areas with dramatic reductions of data quality issues, faster validation times, more accurate processes, and lower numbers of manual reconciliations. Customary static architectures have failed to accommodate changing data behaviors, regulations, and increasing interdependencies between systems. Adopting real-time intelligent infrastructure allows error detection, dynamic logic, and resilience, which are particularly useful in areas with compliance needs. Infrastructure artificial intelligence will be an essential part of future national regulatory systems, helping to solve problems related to data accuracy and compliance that traditional methods couldn't handle.

## 1. Introduction

### 1.1 Background Context

Data movement and validation layers are at the core of modern digital platforms, ingesting, transforming and validating large heterogeneous data volumes from diverse sources. Enterprise and regulatory requirements demand that these systems be very accurate and process data in real-time, but the designs of these systems are often quite fixed. They depend on both deterministic rules and threshold parameters. Such manual exception handling is common in many industries.

Modern systems process more complex data structures. Tabular data is a core element of regulatory systems, financial systems, and healthcare. Explainable AI (XAI) methods for tabular data have improved significantly, making it easier for regulators to understand and verify decision-making. Adding explainability to the validation steps helps solve the ongoing problem of AI being unclear in regulatory compliance situations.

### 1.2 Contemporary Challenges

However, the quantity of records that existing validation methods can process constrains them, leading to larger-scale operational issues. Static rules do not identify subtle anomalies and cannot adapt to changes in patterns in the data. Upstream data drift causes cascading failures that occur after processing is complete.

Machine learning can improve data quality management. Rule-based approaches do not work well with heterogeneous and fast-changing data sources. Machine learning models can learn leading features not easily defined through deterministic rules and adapt to changes in underlying distributions without reprogramming. Classification algorithms allow us to identify data quality issues across multiple dimensions, while clustering algorithms are used to group similar anomalies. Neural networks are able to consider complex interdependencies of data attributes [2]. This feature addresses critical shortcomings of customary validation architectures.

### 1.3 Proposed Solution

Infrastructure-Level Intelligence aims to solve these problems by integrating AI into the parts that manage and check data, putting smart technology closer to the infrastructure and allowing the platform to detect issues before they happen. They enable dynamic alteration of validation logic, greatly improving accuracy and resilience in compliance-critical applications.

In contrast to application layer intelligence, which processes data after the fact, framework intelligence focuses on raw data both in motion and at rest. This enables the detection of anomalies at an earlier stage of the process. It prevents the propagation of errors into downstream systems and helps improve recovery before data quality issues result in operational or compliance problems. This paper introduces Infrastructure-Level Intelligence (ILI) as a distinct architectural paradigm that embeds adaptive artificial intelligence directly into data movement and validation layers, rather than treating intelligence as an application-level or post-processing concern, thereby enabling early anomaly detection, deterministic remediation, and compliance-ready auditability by design. This infrastructure-first treatment of artificial intelligence aligns with emerging regulatory guidance that emphasizes embedding AI risk management, transparency, and governance directly into system architecture rather than relying on downstream application controls [NIST].

### 1.4 Theoretical Foundation

The framework draws on cloud infrastructure architecture, artificial intelligence-based data quality management, platform engineering, and compliance governance and seeks to address fundamental gaps in building large-scale data platforms. This conceptual advance makes it possible to treat AI as an infrastructure primitive rather than a business analytics tool.

Platform engineering is based on the principles of reliability, scalability, and operability, which are also the foundations for Infrastructure-Level Intelligence. Unlike static infrastructure, intelligent infrastructure learns during operation. This approach foresees issues before they arise, transforming a platform's performance under pressure.

## 2. Traditional Architecture Limitations

### 2.1 Static Rule Dependency and Reactive Detection

Rule-based validation assumes consistent data behavior, while reality is different: data distributions constantly change, and product modifications break existing patterns. Regulatory changes result in shifts in reporting requirements, and upstream systems change, creating new data patterns. Static rules produce false positives when the pattern changes and miss real errors that are not within their bounds.

Threshold-based validation also entangles architecture; for instance, a rule checking for a few percent increment in the volume fails on real business growth. The main failure mode is false positives, where non-anomalous data is flagged as anomalous. Also, it is difficult to determine gradual degradation below the bounds. Finally, the mechanism does not adapt to legitimate changes.

Furthermore, when looking for anomalies in a real-time data stream, all data points should be processed as soon as they are generated. Batch-mode validation suffers from unacceptable latency. Stream processing can only model an algorithm on unbounded streams. AI-based anomaly detection supports human interpretation with alerts. Machine learning models can analyze streams of data using a low-latency approach. They can also detect patterns over time, and ensemble models of multiple detection methods can be used together. These methods are particularly relevant for data-intensive platforms [3].

### 2.2 Manual Exception Management and Adaptation Constraints

In the majority of architectures, faults cannot be detected until a job has been fully processed. When faults are detected after job processing, several issues arise. Rework is required when the error is detected too late, and processing tends to be delayed in more than one dependent system. Failing to do so results in operational errors for downstream consumers.

Additionally, post-processing validation introduces latency between ingesting data and detecting errors in it. Root cause analysis can also become complicated, as not all errors can be captured early in processing, and engineers must dig through a data pipeline to find the error.

This slows down the regulatory cycles, introduces variability in compliance regimes, and can lead to different analysts assessing the same patterns differently.

On the other hand, adaptive data rate algorithms use smart methods from deep learning to change system settings based on what they see happening. A common example of adaptive data rate algorithms is LoRaWAN networks that derive suitable

configuration parameters based on their network conditions over time. Resource utilization improves without manual intervention. Energy utilization is reduced, without sacrificing service quality [4]. These are applied to data validation pipelines, where validation is adaptive to the data. Processing resources can be dynamically allocated to high-risk parts of the data.

### 2.3 Scalability and Economic Constraints

Manually reviewing data is not scalable and creates more exceptions as the data scale grows. The trade-off between comprehensiveness and review time is due to compliance regulations. Because speed is so important, the quality of the research is frequently compromised.

However, at scale, the costs become prohibitively high because a human reviewer is required. Exceptions interrupt analysts' workflows, introducing context switching and overhead. Organizations struggle to find enough qualified candidates as domain expertise becomes increasingly important.

Infrastructure-level intelligence reduces these economic constraints as systems scale linearly with the data processed, with the marginal cost per unit of processing decreasing. Model construction is costly, and operational costs do not depend on data size, so they are better suited to high-volume settings. Table 1 presents the fundamental limitations of traditional data validation architectures and contrasts them with the capabilities enabled by infrastructure-level intelligence approaches across key operational dimensions.

## 3. Infrastructure-Level Intelligence Framework

### 3.1 Intelligent Data Ingestion

Machine learning models continuously monitor the ingestion streams to detect abnormal volume patterns. Data anomalies in the time series, such as unusual spikes, can indicate upstream failures or data duplication, while drops can indicate outages or missing sources. Systems catch problems before they reach the pipeline.

Schema drift detection is automatic; models learn the expected schema from the historical ingestion. Deviations are signaled immediately to catch changes upstream in the production cycle. Teams can address schema mismatches before processing begins.

Analyzing latency patterns can provide early indicators of infrastructure degradation. Ingestion time establishes an operational baseline.

Anomalous delays may indicate network or source system problems or a shortage of resources that can be corrected.

Self-healing data pipelines, unlike customary pipelines that crash during an error, automatically detect errors in the pipeline and consequently repair the components. It uses diagnostic algorithms to determine the root cause. Repair strategies are executed automatically. Temporary errors are also retried, using exponential backoff. Data format discrepancies trigger transformation modules, and missing dependencies start fallback processes [5]. This reduces downtime and maintains consistent processing.

### 3.2 AI-Driven Validation and Reconciliation

Static thresholds are replaced with learned data behavior patterns. AI models learn from historical behavior across dimensions. They detect outliers at the record, aggregate, and distribution levels. This multi-granularity approach offers more perception than single-point rules, capturing otherwise missed errors.

Cross-system reconciliation becomes intelligent, as models learn the normal data flow patterns across systems. Timing mismatches emerge, and volume discrepancies induce investigation, while structural inconsistencies also appear, leading to lower reconciliation failure rates.

Adaptive validation logic can evolve with information. Models are updated to reflect normal behavior as business needs change, and with fewer false alarms and better detection rates, they perform exceptionally well on regulatory datasets. Continuous learning has resulted in fewer false positive alerts.

Neural networks can handle complex validation problems while deep learning architectures process multi-dimensional feature spaces. While convolutional neural networks are good at recognizing patterns in grid-like data and recurrent neural networks excel at understanding sequences, attention mechanisms help focus on the most challenging parts of the data, doing better than older statistical methods for complicated data quality issues.

### 3.3 Self-Adaptive Data Movement Pipelines

Intelligence in the infrastructure allows for dynamic optimizations, e.g., pipeline parallelism adjustment based on data size and complexity. It can automatically adjust resource allocations to processing needs and increase throughput without manual tuning. By switching processing from degraded pipelines to healthy ones, intelligent

routing also enhances resilience. Instead of re-running a pipeline in its entirety, partial reprocessing will happen. This allows faster recovery with less disruption.

Processing efficiency becomes very favorable. End-to-end validation cycles become brief. Infrastructure resource consumption is improved. Manual intervention during peak processing windows is reduced. Systems can handle increasing loads without performance degradation.

In terms of actions, self-healing pipelines can react depending on the error type. Syntax errors can trigger data transformation modules, while semantic issues trigger data enrichment procedures. Imputation algorithms are executed in case of missing data. Deduplication logic is executed for duplicate input. It learns to perform corrections over time based on whether the previous methods were successful, through reinforcement learning.

### 3.4 Compliance-Aware Observability and Feedback Loops

Infrastructure-Level Intelligence is designed to be a system that continuously learns. Audit findings can feed back into model learning, regulatory feedback can improve validation logic, and models can improve without a meaningful reprogramming effort.

Infrastructure metrics have correlations to data quality results, as systems learn to anticipate expected conditions. Proactive monitoring is observability to prevent data accuracy issues before they happen, ensuring infrastructure and data quality.

Explainable AI outputs are compliant, providing a justification for every anomaly detected. Audit trails provide linkages between AI decisions and particular data features, thus enabling regulatory reviewers to understand and trust AI's reasoning.

Explainable AI for regulatory compliance also addresses industry-specific needs. For example, financial services need transparency in automated decision-making, while healthcare regulations also require interpretable AI. Model-agnostic explanation methods can be used with most algorithms. LIME provides local explanations for individual predictions. SHAP values measure feature contributions to an individual prediction, while rule extraction (writing a ruleset, tree, etc.) can approximate a complex model with one that meets regulatory demands without sacrificing performance [6]. Decisions are subject to review by compliance officers, and the reasoning behind an automated decision is logged. Table 2 delineates the four primary architectural layers of the Infrastructure-Level Intelligence framework,

detailing the embedded AI capabilities and operational mechanisms that transform passive data pipelines into adaptive intelligent systems.

## 4. Explainability and Compliance Governance

### 4.1 Transparency Requirements and Explainable Techniques

Black-box models can introduce regulatory compliance risks in applications requiring high standards of transparency and explainability from automated decision-making systems. Infrastructure-Level Intelligence considers explainability as part of the design, not simply an afterthought.

To justify these detections, explainable AI methods, such as feature importance analysis, identify the features that triggered the warning, while counterfactual explanations describe the changes that would render flagged data acceptable. Decision trees can offer interpretable reasoning paths, making the AI reasoning accessible to non-technical auditors.

The latest databases are embedding AI and machine learning capabilities to transform infrastructure intelligence, instead of simply reading and writing data. Hence, modern databases also embed AI capabilities in their storage layer. Vector databases enable semantic search on unstructured data. In-database machine learning allows models to run without moving data. Some query optimizers are designed using artificial intelligence, and automated indexing takes query patterns into account [7]. These embedded capabilities reduce latency and improve efficiency, as the AI processing occurs in secure database environments where data never leaves. This architecture improves privacy and compliance.

### 4.2 Model Governance and Audit Traceability

Model governance includes processes for regulatory compliance. Versioning tracks all changes made to a model. The training data lineage and performance are periodically reviewed. Processes are in place to ensure that models meet accuracy requirements before deployment, and model drift is monitored.

Model updates are subject to change management processes, including impact assessment. Testing can be performed on representative data, rollback mechanisms can undo the release, gradual rollout minimizes the risk of deploying, and A/B testing allows comparison between the new and existing models.

Immutable audit logs link the AI's decision back to the data, including the reasoning behind every

validation decision. Inputs, model versioning, confidence scores and prediction outcomes are all logged for regulatory traceability. It helps in analyzing system behavior when conducting an audit.

Real-time measurements of data allow for continuous observability to be achieved in terms of environmental impact assessments, and similar principles apply to data quality monitoring, ultimately providing real-time tracking. Historical data helps to identify trends and deploy predictive maintenance. It can be displayed on dashboards. The ability to notify operators of certain conditions [8] allows operations to remain transparent for regulatory bodies.

### 4.3 Regulatory Alignment and Compliance Assurance

Regulatory frameworks with AI governance requirements exist. There are banking regulators for financial services. Healthcare systems comply with patient privacy standards. Government systems comply with public accountability standards. Infrastructure-Level Intelligence embraces these requirements by design.

Besides calibration, information about the model's hyperparameters and possible points of failure can help. For example, performance benchmarks establish baseline expectations, while limitation disclosures identify scenarios where models may fail. The documentation supports regulatory examinations.

Fairness and bias considerations may influence the validation logic and bias characteristics based on demographic imbalances across data. Regular bias assessments find bias patterns, which can be reduced by adjusting the model's behavior. Fairness metrics complement accuracy metrics in model assessment.

When required by regulation, compliance reporting includes AI-generated audit trails and records of validation activities. By documenting and highlighting data quality over time via incident reports, resolution reports, and performance dashboards, regulators are convinced that these systems can be automated. Table 3 outlines the explainability techniques and governance mechanisms integrated into Infrastructure-Level Intelligence systems to ensure regulatory compliance, transparency, and auditability in automated decision-making processes.

## 5. Impact and Applications

### 5.1 National-Scale Deployment and Cross-Institution Collaboration

Transformation is possible when data validation is near real-time, many records are processed at once with high accuracy, and reporting deadlines for compliance are shortened. Financial and regulatory platforms are especially well positioned to leverage this potential. Organizations have increased confidence in meeting regulatory deadlines.

Systemic risk is reduced if anomalies can be detected early before any cascading effects pass through the interconnected systems. Financial stability is also improved by better quality of data collected by regulators. The early warning capabilities improve oversight.

Federated learning enables multiple institutions to collaborate without sharing patient data across systems, maintaining data privacy while creating a model. This is achieved by training local models on institutional data. Model updates from institutions are shared, not data. Aggregation is performed centrally after learning has occurred in each institution, without transferring sensitive data via privacy-preserving mechanisms. Differential privacy adds noise to protect against reconstruction attacks. Secure aggregation protocols protect model updates during transit [9]. These techniques allow for cross-institution validation while maintaining data sovereignty and have been leveraged by regulatory platforms when legal frameworks restrict data sharing.

### 5.2 Operational Efficiency and Economic Benefits

Provides important operational cost savings; replaces manual reconciliation with automation. Staff focus shifts from exception handling to calculated activity. Infrastructure efficiencies are better leveraged to reduce costs. The economic case becomes more compelling as data grows.

The total cost of ownership decreases over time but requires substantial upfront investments in infrastructure and expertise. However, models have a low development cost, and training prepares workers for new forms of work. Operating costs do not increase with data; for example, labor is not scalable. Infrastructure costs grow linearly, not exponentially.

Infrastructure-Level Intelligence is most cost-effective on high-volume platforms, typically breaking even within a few months. Cost savings result from avoiding the expense of errors and regulatory fines and from considerably improving accuracy. Competitive advantages with faster processing fuel revenue growth.

### 5.3 Regulatory Adaptation and Compliance Agility

Customary platforms necessitate extensive re-engineering to adapt to regulatory changes, but updates propagate the rules. Infrastructure-Level Intelligence systems undergo longer testing cycles and adapt better. Under the new law, models learn patterns based on initial data sets, and validation logic autonomously adapts. Deployment timelines compress dramatically.

Dynamic regulation, where market dynamics and policy objectives prompt regulators to modify the rules, presents a favorable scenario for Infrastructure-Level Intelligence, enabling the organization to more easily comply with guidance. Competitive advantages arise from superior agility. Infrastructure intelligence also benefits regulatory technology, or regtech, making compliance easier. AI can assist in identifying violations of regulation in regulatory reports via validation errors, identify violations in real-time, and document the audit trail. These capabilities both simplify and streamline compliance processes.

**5.4 Cross-Domain Applicability and Industry Adoption**

The framework can also be generalized to other domains, such as health care, with the corresponding data verification problem and supply chain platforms requiring real-time anomaly detection. Government reporting systems require very high accuracy and responsiveness under strict deadlines. The principles described in Infrastructure-Level Intelligence can apply just as

well to any of these domains. Smart pipelines are also useful in healthcare data validation. Electronic health records contain both structured and unstructured data. To ensure the safety of the patient under evaluation, clinical documentation must be precise. Public health reporting needs to be timely for outbreak detection. Infrastructure-Level Intelligence formalizes these diverse requirements into a common architecture.

Supply chain optimization relies on real-time detection of anomalies and sharing of inventory information by suppliers, manufacturers, and distributors. Demand forecasting is done with historical data. Disruption detection enables proactive response. Quality control is improved through automated inspection data and infrastructure intelligence provides supply chain data reliability. Automated remediation extends infrastructure intelligence into the operational sphere and is used within cloud environments to correct drift. Security systems are capable of automated threat responses, and DevOps systems can perform auto-healing. Remediation playbooks are standard operating procedures that are triggered automatically when their conditions are met. Remediation success is measured [10]. These ideas can be generalized to self-healing data pipelines. In turn, such approaches yield intelligent infrastructure networks that largely self-manage with little human involvement. Table 4 shows how Infrastructure-Level Intelligence can be used in different industries, pointing out specific ways it can be applied and the advantages gained from using AI in data validation systems.

*Table 1: Comparison of Traditional and Intelligent Validation Architectures [3, 4]*

<b>Limitation Category</b>	<b>Traditional Architecture Characteristics</b>	<b>Infrastructure-Level Intelligence Approach</b>
Rule Adaptation	Static thresholds with fixed validation parameters that require manual updates for pattern changes	Dynamic learning models that automatically adapt to evolving data distributions and business changes
Error Detection Timing	Post-processing identification requiring complete pipeline execution before anomaly discovery	Real-time anomaly detection during data ingestion and movement preventing downstream propagation
Exception Handling	Manual review processes with subjective analyst judgment introducing inconsistency and delays	Automated intelligent validation with explainable reasoning and self-correcting pipeline mechanisms
Scalability Model	Linear cost growth with data volume requiring proportional increase in human resources	Automated processing with stable operational costs independent of data volume growth
Adaptation Speed	Prolonged re-engineering cycles for regulatory changes with extensive testing requirements	Rapid adaptation through model learning from initial data under updated regulatory frameworks

*Table 2: Infrastructure-Level Intelligence Framework Architecture [5, 6]*

<b>Framework Layer</b>	<b>Embedded AI Capabilities</b>	<b>Operational Mechanisms</b>
Intelligent Data Ingestion	Machine learning models for stream monitoring and schema drift	Abnormal volume pattern identification, structural inconsistency flagging, latency

	detection	anomaly analysis
AI-Driven Validation	Multi-dimensional outlier detection and cross-system reconciliation learning	Record-level analysis, aggregate pattern recognition, distribution-level anomaly identification
Self-Adaptive Pipelines	Dynamic parallelism adjustment and intelligent routing algorithms	Automated resource allocation, degraded component bypass, partial reprocessing triggers
Compliance Observability	Continuous feedback loops and explainable AI output generation	Audit result integration, infrastructure-quality correlation, regulatory feedback incorporation
Autonomous Remediation	Error classification and correction strategy selection	Syntax transformation, semantic enrichment, imputation execution, deduplication logic

**Table 3: Explainability and Governance Framework for Regulatory Compliance [7, 8]**

Governance Component	Implementation Techniques	Regulatory Alignment Features
Explainable AI Methods	Feature importance analysis, counterfactual explanations, decision tree interpretation	LIME for local predictions, SHAP values for feature contribution, rule extraction for model approximation
Model Governance	Version control tracking, training data lineage documentation, performance metric review	Impact assessment protocols, representative data testing, gradual rollout procedures
Audit Traceability	Immutable logging of AI decisions with complete reasoning chains	Input data capture, model version recording, confidence score documentation
Transparency Mechanisms	Model-agnostic explanation generation across diverse algorithms	Human-readable logic paths, non-technical auditor accessibility, supervisory confidence building
Fairness Assurance	Regular bias assessment and demographic disparity analysis	Fairness metric integration, equitable treatment validation, mitigation strategy implementation

**Table 4: Cross-Domain Applications and Implementation Benefits [9, 10]**

Application Domain	Infrastructure Intelligence Use Cases	Operational Benefits
Financial Services	National-scale regulatory reporting, systemic risk detection, compliance deadline management	Early anomaly detection preventing cascading failures, compressed reporting timelines, enhanced supervisory visibility
Healthcare Systems	Electronic health record validation, clinical documentation accuracy, public health outbreak detection	Patient safety through accurate documentation, fraud prevention in insurance claims, timely epidemic response
Supply Chain Networks	Real-time inventory data validation, demand forecasting accuracy, disruption detection	Supplier-manufacturer-distributor data reliability, proactive response capabilities, automated inspection validation
Government Platforms	Cross-institution collaboration, regulatory change adaptation, compliance reporting automation	Data sovereignty preservation through federated learning, rapid policy adjustment, transparent audit trail generation
RegTech Solutions	Automated compliance monitoring, real-time violation detection, policy enforcement embedding	Reduced compliance burden, automated documentation support, competitive advantage through agility

## 6. Conclusions

Infrastructure-Level Intelligence is a model shift in data platform architecture from customary approaches, embedding machine learning and artificial intelligence into the application layer (for analytics and prediction). Strategically, Infrastructure-Level Intelligence can provide major operational advantages using intelligence in the infrastructure (for validation and in the data

pipeline. Real-time anomaly detection and adaptive validation logic replace post-processing error detection and static rule sets. Self-correcting pipelines minimize manual intervention and support compliance with an explainable audit trail using explainable AI. Together, these show that the improvements have real business value in many deployment cases. So this type of intelligent, learning-based generalization at the infrastructure level will have to be a first principle for future

national scale systems. Compliance-critical environments can gain from built-in AI features because this method aims to solve problems with data platform reliability and meeting regulations that were once thought impossible, as organizations deal with growing complexity and regulatory demands that traditional systems can't handle. Infrastructure-Level Intelligence systems are platforms that can achieve accuracy, resiliency and compliance at a scale that was never possible with traditional architectural patterns. Future research may focus on federated learning for regulatory compliance across institutions, while maintaining privacy. AI regulatory policy enforcement may occur at infrastructure controls. Autonomous remediation may also be researched to eliminate breaches of critical compliance controls without human intervention. Observability may become more tightly integrated with predictive analytics to prevent issues. Explainable AI methods will continue to evolve along with regulatory pressures for transparency. This article area will expand from the Infrastructure-Level Intelligence use case into new use cases while retaining the architecture principles of embedded intelligence as well as autonomous operation on top of a uniform architecture.

### Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

### References

1. Maria Sahakyan, et al., "Explainable Artificial Intelligence for Tabular Data: A Survey," IEEE Xplore, 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9551946>
2. Shreyas B, "Machine Learning Techniques for Data Quality Management," Dview, 2024. [Online]. Available: <https://dview.io/blog/machine-learning-for-data-quality-management>
3. Shakir Poolakkal Mukkath, "AI-driven anomaly detection in real-time streaming: enhancing human decision-making," World Journal of Advanced Engineering Technology and Sciences, 2025. [Online]. Available: <https://journalwjaets.com/sites/default/files/fulltext/pdf/WJAETS-2025-0583.pdf>
4. Luca Leonardi, et al., "A deep learning-based Adaptive Data Rate algorithm for LoRaWAN networks," Internet of Things, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2542660525003002>
5. Yashwanth Boddu, "Self-Healing Data Pipelines with Autonomous Error Correction," TIJER, 2025. [Online]. Available: <https://tijer.org/tijer/papers/TIJER2506049.pdf>
6. Nikhil Gupta, "Explainable AI for Regulatory Compliance in Financial and Healthcare Sectors: A comprehensive review," International Journal of Advances in Engineering and Management (IJAEM), 2025. [Online]. Available: [https://ijaem.net/issue\\_dcp/Explainable%20AI%20for%20Regulatory%20Compliance%20in%20Financial%20and%20Healthcare%20Sectors%20%20A%20comprehensive%20review.pdf](https://ijaem.net/issue_dcp/Explainable%20AI%20for%20Regulatory%20Compliance%20in%20Financial%20and%20Healthcare%20Sectors%20%20A%20comprehensive%20review.pdf)
7. Robert Gravelle, "The Rise of Embedded AI/ML Capabilities in Modern Databases," Navicat, 2025. [Online]. Available: <https://www.navicat.com/en/company/aboutus/blog/3350-the-rise-of-embedded-ai-ml-capabilities-in-modern-databases.html>
8. Tamírís Pacheco da Costa, et al., "A systematic review of real-time data monitoring and its potential application to support dynamic life cycle inventories," Environmental Impact Assessment Review, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0195925524000039>
9. Micah J. Sheller, et al., "Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data," Sci Rep, 2020. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC7387485/>
10. Joseph Barringhaus, "Automated Remediation: Key Benefits, Best Practices & Industry Use Cases," Tamnoon, 2024. [Online]. Available: <https://tamnoon.io/blog/automated-cloud-remediation-guide/>
11. National Institute of Standards and Technology (NIST), "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," U.S. Department of Commerce, 2023. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

1. Maria Sahakyan, et al., "Explainable Artificial Intelligence for Tabular Data: A Survey," IEEE