Explainable AI (XAI) Interfaces for Edge–Cloud Forecasting Platforms: Enhancing Operations Manager Trust in Deep Learning Predictive Output

Authors

Asher Noah, Apu Ahsun

Date; September 28, 2026

Abstract

Edge–cloud forecasting platforms increasingly deploy deep learning models for real-time operational prediction, yet operations managers remain reluctant to act on outputs they cannot interpret. This study addresses the gap between predictive accuracy and managerial trust by designing, implementing, and validating an explainable AI (XAI) interface layer for a collaborative hybrid CNN–LSTM edge–cloud forecasting architecture. Using a design-based research methodology combining retrospective retail demand data analysis with a controlled user evaluation (N=48 operations managers), the study tested an XAI dashboard integrating SHAP attributions and counterfactual explanations against a standard prediction-only interface. Results indicate that the XAI-enhanced interface achieved 89.4% managerial trust endorsement compared to 61.2% for the control condition ($p < 0.001$), while preserving the underlying model's 94.7% forecasting accuracy. Feature attribution transparency significantly reduced decision latency and increased appropriate reliance on model outputs. The study concludes that trust-calibrated XAI interfaces constitute a necessary socio-technical layer for edge–cloud forecasting adoption, and provides a replicable framework for designing explanation dashboards that align with operations managers' decision workflows.

Keywords: Explainable AI, edge–cloud computing, operations management, trust calibration, demand forecasting

1. Introduction

1.1 Background

The convergence of edge computing and cloud-based deep learning has enabled a new class of forecasting platforms capable of real-time prediction across distributed operational environments. Aashish et al. (2026) demonstrated that a collaborative hybrid CNN–LSTM architecture deployed on edge–cloud infrastructure achieves 94.7% accuracy in multichannel retail demand forecasting, significantly outperforming standalone LSTM, CNN, and gradient boosting baselines . Similar architectures have shown promise in smart manufacturing, where cross-layered edge–cloud deep learning supports digital twin analytics and predictive maintenance with reduced latency and improved responsiveness . The technical case for edge–cloud forecasting is compelling: local inference at the edge reduces bandwidth dependency and enables sub-second response times, while cloud-layer training provides the computational resources for complex model refinement.

Yet a persistent gap separates technical performance from operational adoption. Operations managers—the individuals responsible for translating forecasts into inventory decisions, production schedules, and resource allocation—frequently exhibit what the human–AI interaction literature terms "algorithmic aversion": a reluctance to rely on model outputs whose reasoning cannot be inspected or challenged . This reluctance is not irrational. In high-stakes operational contexts, managers bear accountability for decisions, and accountability without understanding is untenable. As one systematic review of XAI in sustainability-oriented systems observed, "trust cannot be directly measured as a single scalar variable but is inferred through the joint satisfaction of performance, interpretability, ethical alignment, and robustness" requirements . The edge–cloud forecasting platform, for all its predictive power, presents managers with precisely the kind of opaque output that erodes rather than builds this multidimensional trust.

1.2 Problem Statement

Research on XAI has advanced significantly in recent years, producing a rich toolkit of post-hoc explanation methods including SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-agnostic Explanations), and various attention-based attribution techniques . These methods have been applied to energy forecasting, healthcare diagnosis, and industrial monitoring with encouraging results . Similarly, edge–cloud architectures for predictive analytics have matured rapidly, with demonstrated deployments in manufacturing, retail, and infrastructure management .

However, these two research streams have developed largely in parallel. XAI research has focused predominantly on algorithmic fidelity—whether explanations faithfully represent model behavior—rather than on whether explanations actually improve decision-maker trust and performance in specific operational contexts . Edge–cloud forecasting research has focused on latency, bandwidth, and accuracy metrics, treating the manager primarily as a consumer of predictions rather than an active participant in a human–AI decision loop . The result is a design gap: no validated framework exists for XAI interfaces specifically tailored to operations managers using edge–cloud forecasting platforms. The unsolved issue, therefore, is how to design explanation interfaces that transform technically accurate deep learning predictions into operationally actionable intelligence without overwhelming managers with model internals or, conversely, providing explanations so superficial that they fail to support genuine trust calibration.

1.3 Objectives of the Study

General objective: To design and validate an XAI interface framework for edge–cloud forecasting platforms that enhances operations manager trust in deep learning predictive outputs.

Specific objectives:

1. To identify the explanation types and presentation formats that operations managers perceive as most useful for evaluating demand forecasts.
2. To design an integrated XAI interface layer that combines SHAP feature attributions with counterfactual explanations within an edge–cloud forecasting architecture.
3. To validate the framework through a controlled user evaluation measuring trust endorsement, decision latency, and appropriate reliance compared to a prediction-only interface.

1.4 Research Questions

1. What combination of XAI explanation techniques and interface design elements most effectively enhances operations manager trust in deep learning forecasting outputs?
2. How does an XAI-enhanced interface compare to a standard prediction-only interface in terms of managerial trust endorsement, decision confidence, and appropriate reliance?
3. What are the technical and organizational barriers to implementing XAI interfaces within existing edge–cloud forecasting workflows?

1.5 Significance of the Study

For practitioners and administrators: This research provides a validated design specification for XAI dashboards that operations managers can use to evaluate, challenge, and ultimately trust forecasting outputs. The 89.4% trust endorsement rate suggests that appropriately designed explanations can substantially reduce the adoption barrier that currently limits edge–cloud forecasting ROI.

For policymakers: The study contributes to the operationalization of transparency requirements in AI governance frameworks, demonstrating that explainability can be engineered into edge–cloud platforms without sacrificing predictive performance or computational efficiency.

For academic literature: The research extends socio-technical trust theory into the specific context of edge–cloud forecasting, bridging XAI algorithmic research and operations management decision science. It provides empirical evidence that explanation quality—not merely explanation presence—drives trust outcomes.

For future researchers: The framework offers a replicable methodology for evaluating XAI interfaces in operational contexts, including validated trust measures and a controlled experimental design that can be adapted to other domains.

1.6 Scope and Limitations

This study focuses on demand forecasting in multichannel retail operations, using publicly available retail transaction data and a simulated edge–cloud deployment environment. The user evaluation involved 48 operations managers from retail and consumer goods organizations, recruited through professional networks. The study excludes: (a) financial forecasting applications, (b) healthcare operational forecasting, and (c) fully autonomous decision systems where no human review occurs. Key limitations include the controlled (non-field) evaluation setting, the use of simulated rather than live edge–cloud infrastructure, and the cross-sectional rather than longitudinal measurement of trust. These limitations are discussed in Section 5.3.

2. Literature Review

2.1 Conceptual Review

Explainable AI (XAI): XAI refers to methods and techniques that make machine learning model outputs understandable to human users. The field distinguishes between intrinsically interpretable models (e.g., decision trees, linear regression) and post-hoc explanation methods applied to black-box models . For deep learning forecasting, post-hoc methods predominate, with SHAP and LIME being the most widely adopted . SHAP, grounded in cooperative game theory, assigns each input feature a contribution value representing its marginal effect on the prediction. LIME constructs local surrogate models to approximate black-box behavior near a specific prediction instance. Extensions such as TimeSHAP and WindowSHAP have been developed specifically for temporal data, grouping time steps or applying pruning strategies to accommodate the sequential structure of forecasting inputs .

Edge–Cloud Forecasting: Edge–cloud architectures distribute computation between local edge devices (for low-latency inference and data preprocessing) and cloud infrastructure (for model training and large-scale analytics). In forecasting applications, this architecture enables real-time prediction from streaming sensor or transaction data while maintaining the computational capacity for complex model refinement . Aashish et al. (2026) demonstrated that a collaborative hybrid CNN–LSTM framework deployed on edge–cloud infrastructure achieved 94.7% accuracy and 0.96 ROC AUC in multichannel retail demand forecasting, with the collaborative learning paradigm reducing both inference latency and bandwidth consumption .

Operations Manager Trust: Trust in automation and AI has been theorized as a multidimensional construct encompassing performance trust (confidence in accuracy), process trust (confidence in how outputs are generated), and purpose trust (alignment with user goals) . In operational contexts, trust is behaviorally manifested as appropriate reliance: neither blind acceptance nor reflexive rejection of AI outputs. The XAI literature has established that explanation presence alone does not guarantee trust improvement; explanations must be relevant, comprehensible, and aligned with the user's decision task . A user study of XAI in no-code ML platforms found that novices rated explanations as useful and trustworthy when presented with Partial Dependence Plots, Permutation Feature Importance, and KernelSHAP, while experts were more critical of explanation sufficiency .

2.2 Theoretical Framework

Prospect Theory and Managerial Decision-Making Under Uncertainty: Prospect theory, developed by Kahneman and Tversky, posits that decision-makers evaluate outcomes relative to a reference point and exhibit loss aversion—losses loom larger than equivalent gains. In operational forecasting contexts, managers evaluate predictions against a reference expectation (e.g., historical demand patterns). An unexplained forecast that deviates from expectation triggers loss-aversion responses: managers may discount the prediction to avoid the "loss" of being wrong. XAI interfaces mitigate this by reframing prediction deviations as explainable feature effects, shifting evaluation from outcome comparison to process assessment. When a manager sees that a demand spike prediction is attributable to a specific promotional calendar or weather pattern, the deviation becomes intelligible rather than threatening.

Socio-Technical Trust Theory: This framework conceptualizes trust in AI systems as emerging from the joint satisfaction of technical reliability, interpretability, ethical alignment, and robustness requirements . Trust is not a property of the model alone but of the entire socio-technical system comprising the model, the interface, the organizational context, and the user's cognitive and motivational state. This theory directly informs the present study's design: the XAI interface is positioned not as an add-on but as the trust-generating mechanism that connects model performance to managerial action.

2.3 Empirical Review

Aashish et al. (2026) developed and evaluated a collaborative hybrid CNN–LSTM framework for real-time multichannel retail demand forecasting using edge–cloud computing. Their architecture achieved 94.7% model accuracy and 0.96 ROC AUC, outperforming LSTM, CNN, RNN, and LightGBM baselines. However, the study did not incorporate any explanation mechanism; the interface between model outputs and managerial decision-making was left unexamined .

A systematic review of XAI in energy forecasting examined 50 studies and identified four dimensions of explanation quality: Global Transparency, Local Fidelity and Robustness, User Relevance, and Operational Viability. The review found that while SHAP and LIME were widely applied, few studies explicitly evaluated explanations with actual decision-makers or tailored explanation content to specific user roles .

A user study on XAI in no-code ML platforms (N=20) integrated Partial Dependence Plots, Permutation Feature Importance, and KernelSHAP into the DashAI workflow. Results showed high task success rates (≥80%) and that novices rated explanations as useful, accurate, and trustworthy on the Explanation Satisfaction Scale, while experts were more critical. The study concluded that explanation design must accommodate varying expertise levels .

Research on operationalizing selective transparency in clinical AI diagnosis systems found that "more" transparency is not always better; complex or irrelevant explanations can backfire and decrease trust. The study emphasized that explanations must be tailored to the user's decision context and expertise .

A three-layer architecture for digital twin analytics in smart manufacturing incorporated an AI conversational assistant with XAI capabilities, immersive VR/AR environments, and interactive dashboards. The human-centric interface layer was designed to translate complex technical outputs into "simple and useful insights" for managers, with the AI assistant answering "why" questions about model predictions .

2.4 Research Gap

The empirical literature establishes that XAI methods can be technically implemented and that explanation presence can influence trust in controlled settings. However, no validated XAI interface framework exists specifically for operations managers using edge–cloud forecasting platforms. Existing XAI research has either focused on technical users (data scientists, model developers) or on domains such as healthcare where clinician–AI interaction patterns differ substantially from operations management. The edge–cloud forecasting literature, meanwhile, has not addressed the explanation interface as a design concern. This study fills that gap by designing, implementing, and empirically validating an XAI interface layer that connects deep learning predictive outputs to operations manager decision-making within an edge–cloud architecture.

3. Methodology

3.1 Research Design

This study employed a design-based research methodology combining retrospective data analysis with a controlled user evaluation. This design was chosen for three reasons: (a) it enables iterative refinement of the XAI interface based on empirical trust outcomes, (b) it permits controlled comparison of explanation conditions while maintaining realistic forecasting task demands, and (c) it produces both a designed artifact (the XAI framework) and theoretical knowledge (trust mechanisms in operational forecasting) .

3.2 Study Area / Population

The study population comprised operations managers in retail and consumer goods organizations (N=48) with responsibility for demand forecasting, inventory planning, or supply chain coordination. Participants were recruited through professional operations management networks and LinkedIn, screened for minimum two years of experience in forecasting-related roles. The forecasting task context was multichannel retail demand, using a publicly available transaction dataset.

3.3 Sample Size and Sampling Technique

A purposive sampling strategy was employed, targeting managers with direct forecasting decision authority. The final sample (N=48) was stratified by organizational role: 18 inventory planners, 16 supply chain managers, and 14 operations directors. This sample size was determined through a priori power analysis for a two-group comparison with medium effect size ($d=0.5$), $\alpha=0.05$, power=0.80, yielding a minimum of 34 participants; we oversampled to accommodate stratification and potential attrition.

3.4 Data Collection Methods

Forecasting data source: A publicly available multichannel retail dataset comprising two years of daily sales transactions across 45 product categories and three sales channels (in-store, online, mobile app). The dataset contains approximately 1.2 million transaction records with timestamps, product identifiers, channel indicators, promotional flags, and pricing information.

User evaluation data: Each participant completed a structured forecasting evaluation task consisting of 12 demand scenarios. For each scenario, participants viewed a forecast (either with XAI explanations or without, depending on condition assignment) and rated: (a) trust in the forecast on a 7-point Likert scale, (b) decision confidence, (c) intended action (increase order, decrease order, maintain), and (d) perceived explanation usefulness (XAI condition only). Response latency was recorded for each scenario.

3.5 Research Instruments

Forecasting model: The collaborative hybrid CNN–LSTM architecture proposed by Aashish et al. (2026) was adapted as the base forecasting model . The model processes multivariate inputs (lagged sales, promotional indicators, day-of-week, channel identifiers) through parallel CNN and LSTM branches, with a fusion layer producing demand predictions. This architecture was selected because its edge–cloud deployment paradigm aligns with the study's operational context and because its published accuracy (94.7%) provides a validated performance baseline.

XAI implementation: SHAP values were computed for each prediction instance using the DeepExplainer implementation for hybrid architectures. Temporal extension followed the WindowSHAP approach, aggregating SHAP values across meaningful time windows (e.g., week-level attribution) to reduce cognitive load . Counterfactual explanations were generated by identifying the minimal feature perturbation (e.g., "if promotional intensity were 20% lower") that would flip the predicted demand category from "high" to "moderate."

XAI interface: The dashboard was implemented as a React-based web application with three explanation components: (1) a feature attribution bar chart showing the top five contributors to the prediction, (2) a temporal heatmap displaying SHAP values aggregated by week, and (3) a counterfactual panel presenting two alternative scenarios with their predicted outcomes. The interface was designed for clarity rather than technical completeness, following the selective transparency principle that explanations must match user decision needs .

3.6 Validity and Reliability

Content validity: The explanation components and interface design were reviewed by three domain experts (two operations managers, one XAI researcher) prior to deployment. Revisions addressed terminology (replacing "feature attribution" with "key drivers") and visual density (reducing the number of displayed features from eight to five).

Predictive validity: The forecasting model was validated on a held-out test set comprising the final six months of transaction data. Model accuracy (94.7%) and ROC AUC (0.96) matched the performance reported by Aashish et al. (2026), confirming that the adapted architecture maintained predictive integrity .

Inter-rater reliability: Trust ratings from a subset of participants ($n=12$) who completed the evaluation twice, two weeks apart, yielded a test-retest correlation of $r=0.82$, indicating acceptable stability of trust measurements.

3.7 Data Analysis Techniques

Model performance was evaluated using accuracy, precision, recall, F1-score, and ROC AUC. Trust endorsement was operationalized as the proportion of scenarios in which participants rated trust ≥ 5 on the 7-point scale. Between-group comparisons used independent-samples t-tests for continuous variables and chi-square tests for categorical variables. Effect sizes were calculated using Cohen's d. Feature importance rankings were derived from mean absolute SHAP values across the test set. Cross-validation for model selection used a rolling-window scheme appropriate for time-series data, with a forecast horizon of seven days and a look-back window of 28 days .

3.8 Ethical Considerations

All data used in this study were de-identified and publicly available. No personally identifiable information was accessed. User evaluation participants provided informed consent and were compensated for their time. The study protocol was reviewed and determined to be exempt from IRB oversight under the institutional exemption for research involving public datasets and non-sensitive decision-making tasks.

4. Results

4.1 Data Presentation

Table 1. Descriptive Statistics of Forecasting Task Performance

Indicator	XAI Condition (n=24)	Control Condition (n=24)
Mean trust rating (1–7)	5.68 (SD=0.94)	4.31 (SD=1.21)
Trust endorsement rate (≥5)	89.4%	61.2%
Mean decision confidence (1–7)	5.42 (SD=1.03)	4.58 (SD=1.15)
Mean response latency (seconds)	18.4 (SD=6.2)	24.7 (SD=8.1)
Appropriate reliance rate	76.3%	58.9%
Explanation usefulness rating	5.91 (SD=0.87)	—

Note. Appropriate reliance defined as agreement with model prediction when confidence was high and deviation when confidence was low.

Table 1 presents descriptive statistics for the primary outcome measures. The XAI condition demonstrated higher mean trust (5.68 vs. 4.31), higher trust endorsement rate (89.4% vs. 61.2%), and shorter decision latency (18.4s vs. 24.7s) compared to the control condition.

4.2 Analysis of Results

Model performance: The CNN–LSTM forecasting model achieved 94.7% classification accuracy and 0.96 ROC AUC on the held-out test set, consistent with published benchmarks for this architecture . No significant performance degradation was observed when SHAP explanation generation was enabled, indicating that the XAI layer operates as a parallel process rather than modifying the underlying prediction.

Trust outcomes: Trust endorsement was significantly higher in the XAI condition (89.4%) than in the control condition (61.2%), $\chi^2(1, N=48) = 12.84$, $p < 0.001$. Mean trust ratings were also significantly higher, $t(46) = 4.31$, $p < 0.001$, $d = 1.24$. Decision confidence followed a similar pattern, $t(46) = 2.67$, $p = 0.011$, $d = 0.77$.

Decision latency: Participants in the XAI condition reached decisions faster (M=18.4s) than control participants (M=24.7s), $t(46) = -3.02$, $p = 0.004$, $d = -0.87$. This suggests that explanations reduced cognitive effort by providing a structured framework for evaluating predictions.

Appropriate reliance: The XAI condition demonstrated higher appropriate reliance (76.3% vs. 58.9%), $\chi^2(1, N=48) = 6.72$, $p = 0.010$, indicating that explanations supported calibrated trust—neither blind acceptance nor reflexive rejection.

Feature importance: The top five predictors of demand, ranked by mean absolute SHAP value, were: (1) lagged sales (7-day), (2) promotional intensity, (3) day-of-week, (4) channel mix, and (5) seasonal index. This ranking was consistent across product categories, suggesting that the explanation content provides stable guidance for managerial evaluation.

5. Discussion

5.1 Interpretation

The finding that XAI explanations increased trust endorsement from 61.2% to 89.4% while preserving 94.7% predictive accuracy supports the socio-technical trust framework's proposition that trust emerges from the joint satisfaction of performance and interpretability requirements . This is not merely a "halo effect" of explanation presence; the reduced decision latency and improved appropriate reliance indicate that explanations functioned as cognitive scaffolds, enabling managers to evaluate predictions against their domain knowledge rather than accepting or rejecting them holistically.

The feature importance results align with the demand forecasting literature's emphasis on lagged demand, promotional effects, and temporal patterns as key drivers . The explanation interface made these drivers visible at the individual prediction level, transforming a black-box output into a reviewable claim about demand determinants. This finding extends the selective transparency principle from healthcare AI to operations management: explanations that match the decision-maker's evaluation criteria improve trust without requiring full model disclosure.

The counterfactual explanations appeared particularly influential in the qualitative feedback, with participants reporting that seeing "what would change the prediction" helped them assess the boundary conditions of model reliability. This aligns with prospect theory's prediction that decision-makers evaluate uncertainty relative to reference points; counterfactuals provide alternative reference points against which the primary prediction can be assessed.

5.2 Implications

Academic implications: This study extends socio-technical trust theory by demonstrating that trust-relevant explanation dimensions (global transparency, local fidelity, user relevance) are not merely theoretical constructs but measurable design variables with quantifiable trust effects. It also bridges the XAI algorithmic literature and operations management decision research, introducing a framework for evaluating explanation quality in terms of decision outcomes rather than algorithmic fidelity alone.

Practical implications: Operations managers evaluating edge–cloud forecasting platforms should prioritize XAI interface capabilities alongside accuracy metrics. The framework presented here—SHAP attribution bars, temporal heatmaps, and counterfactual panels—provides a concrete design specification. Organizations deploying forecasting platforms should monitor trust endorsement rates as a key performance indicator; a rate below 80% suggests that explanation design requires attention, even if model accuracy is high.

For system designers: The edge–cloud architecture must accommodate explanation generation with minimal latency overhead. SHAP computation for hybrid CNN–LSTM models can be optimized through caching strategies and background precomputation, ensuring that explanations are available at inference time without degrading the sub-second response that motivates edge deployment.

5.3 Limitations

1. **Controlled evaluation setting:** The user evaluation was conducted in a structured task environment rather than live operational conditions. Trust measured in controlled settings may not fully predict trust in high-stakes, time-pressured decisions.
2. **Simulated edge–cloud deployment:** While the forecasting model and XAI components were functionally deployed, the edge–cloud infrastructure was simulated rather than implemented on physical edge devices. Latency and bandwidth characteristics may differ in production environments.
3. **Single domain:** The study focused on retail demand forecasting. Generalizability to manufacturing, logistics, or energy forecasting requires replication.

4. **Cross-sectional trust measurement:** Trust was measured at a single time point. The stability of trust over repeated interactions and its relationship to actual adoption behavior remain unexamined.

5.4 Future Research Directions

1. **Longitudinal field deployment:** Conduct a six-month field study tracking trust endorsement, decision quality, and adoption behavior among managers using the XAI interface in live operations.
2. **Cross-domain validation:** Replicate the framework in manufacturing predictive maintenance and energy load forecasting contexts to assess generalizability.
3. **Explanation personalization:** Investigate whether adapting explanation type (attribution vs. counterfactual vs. example-based) to individual manager characteristics (analytical orientation, risk tolerance, experience level) improves trust outcomes further.
4. **Explanation drift monitoring:** Develop methods for detecting when explanation patterns change (indicating model drift) and for alerting managers to shifts in the reliability of familiar explanations.

6. Conclusion

This study demonstrates that an XAI interface layer designed specifically for operations managers can achieve 89.4% trust endorsement of edge–cloud forecasting outputs while preserving the underlying model's 94.7% predictive accuracy. The framework—integrating SHAP feature attributions, temporal heatmaps, and counterfactual explanations—transforms deep learning predictions from opaque outputs into reviewable decision inputs. For administrators, the practical takeaway is clear: investing in explanation interface design is not a compliance expense but a trust-generating capability that reduces decision latency and improves appropriate reliance. As edge–cloud forecasting platforms proliferate across operational domains, the question is no longer whether deep learning can predict, but whether managers will act on those predictions. The answer, this study suggests, depends on whether the platform explains.

References

1. Aashish, K. C., Ahmed, K. R., Goyal, S. K., Chowdhury, M. A. R., Salam, T., & Alvi, M. T. A. (2026, April). A collaborative hybrid CNN–LSTM framework for real-time multichannel retail demand forecasting using edge–cloud computing. In *2026 IEEE 2nd International Conference on Quantum Photonics, Artificial Intelligence & Networking (QPAIN)* (pp. 1–6). IEEE.
2. Abarca, N., Carvallo, A., Moncada, C., & Bravo-Marquez, F. (2025). Explaining AI without code: A user study on explainable AI. In *Advances in Neural Information Processing Systems 38*. NeurIPS.
3. Altukhi, Z., & Pradhan, S. (2026). C-XplainEd: A conceptual framework for trustworthy XAI educational systems. In *Human–Computer Interaction and AI in Education* (pp. 23–45). Taylor & Francis.
4. Bento, J., Saleiro, P., Cruz, A. F., Figueiredo, M. A. T., & Bizarro, P. (2021). TimeSHAP: Explaining recurrent models through sequence perturbations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (pp. 2565–2573). ACM.
5. Guillemé, M., et al. (2026). Explainable semi-supervised methods for multivariate time-series analysis: A review. *Applied Mathematics and Computer Science*, *36*(1), 145–168.
6. HumAlne Project. (2026). *D3.1: XAI methods and benchmark suite for human–AI systems* (Version 1). European Commission.
7. Kumar, A., et al. (2026). Next-generation digital twin analytics in smart manufacturing using cross-layered edge–cloud deep learning. In *2026 IEEE International Conference on Industrial Technology*. IEEE.
8. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765–4774).
9. Mashhido, T., et al. (2025). Intelligent predictive maintenance: Multivariate ML model optimization in an Edge–Fog–Cloud environment. *Computing*, *107*(11), 1–28.
10. Meske, C., Bunde, E., Schneider, J., & Gersch, M. (2022). Explainable artificial intelligence: Objectives, stakeholders, and future research opportunities. *Information Systems Management*, *39*(1), 53–63.
11. Rosenbacke, R., et al. (2024). Operationalizing selective transparency using progressive disclosure in AI clinical diagnosis systems. *HAL Open Science*.
12. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. In *Proceedings of the 4th Machine Learning for Healthcare Conference* (pp. 359–380). PMLR.

13. Tran, N. K., et al. (2025). The costs of maritime supply chain disruptions: The case of the Suez Canal blockage. *International Journal of Production Economics*, *279*, 109464.
14. Wang, Y., et al. (2026). COGEniX: Interpretable multi-horizon forecasting and carbon-aware scheduling for industrial edge deployment. *Scientific Reports*, *16*, Article 29397.
15. Zaid, A., et al. (2026). Explainable AI for strategic decision-making in management information systems: A critical review. *Journal of Information and Data Management*, *8*(2), 112–138.
16. Zhang, H., et al. (2025). Explainable multi-step heating load forecasting using SHAP values and temporal attention mechanisms. *Energy and AI*, *18*, 100412.