Multichannel Degradation Pattern Mining in Nuclear Plant Cooling and Auxiliary Loops

Authors

Anders Nilsson[1], Karin Lundqvist[2*]

Affiliations

Department of Mechanical and Process Engineering, ETH Zurich, 8092 Zurich, Switzerland

***Corresponding author:** karin.lundqvist@berkeley.edu

Abstract

Nuclear power plants rely on continuous monitoring of cooling loops, auxiliary pumps, heat exchangers, pressure vessels, valve positions, coolant temperature, flow rate, vibration, and radiation-related process indicators. These variables are highly coupled through thermal transfer, hydraulic balance, safety-control logic, and delayed feedback mechanisms. A minor deviation in pump efficiency, valve response, or heat-exchanger performance may propagate through several subsystems before triggering system-level alarms. This study proposes a causal fault attribution method for multivariate time series from nuclear plant cooling and auxiliary systems. The method constructs a time-lagged causal graph using conditional independence testing, physical safety constraints, and control-loop dependency rules. A counterfactual residual module estimates expected variable trajectories under normal thermal-hydraulic behavior. Root-cause variables are then ranked according to their causal contribution to downstream deviations. Experiments are conducted on a nuclear plant simulation and monitoring dataset containing 12 cooling-loop subsystems, 146 process variables, and 1-second measurements collected across 9,800 simulated operating cycles and 14 months of auxiliary-system logs. The dataset contains 428 million timestamped records and 1,260 expert-reviewed abnormal episodes, including coolant flow degradation, auxiliary pump instability, valve-response delay, heat-exchanger fouling, sensor drift, and abnormal pressure oscillation. The proposed method reduces median root-cause localization time from 33.8 minutes to 8.6 minutes compared with a temporal reconstruction baseline. The mean reciprocal rank for initiating-variable identification reaches 0.851. Causal path analysis assigns 980 abnormal episodes to interpretable thermal-hydraulic propagation chains, while unnecessary subsystem inspection tickets decrease from 760 to 284. Median inference latency remains 37 ms per monitoring window. The results indicate that causal fault attribution can improve fine-grained anomaly diagnosis in safety-critical nuclear plant monitoring systems.

Keywords: Nuclear plant monitoring; causal fault attribution; multivariate time series; cooling system anomaly; counterfactual residuals; root-cause localization; safety-critical process monitoring.

1. Introduction

Nuclear power plants depend on the stable operation of cooling loops, auxiliary pumps, heat exchangers, pressure vessels, control valves, and other safety-related components. Their operating condition is continuously reflected by multivariate measurements such as coolant temperature, flow rate, pressure, vibration, valve position, pump speed, heat-transfer efficiency, and radiation-related process indicators. These variables are strongly coupled through thermal transport, hydraulic balance, control-system logic, equipment interactions, and delayed feedback mechanisms [1,2]. Consequently, a small decline in pump efficiency, valve responsiveness, or heat-exchanger performance may initially appear as a weak local deviation before propagating to multiple downstream variables and subsystems. Identifying such early-stage abnormalities is essential because delayed diagnosis may lead to unnecessary protective actions, reduced operational efficiency, accelerated component degradation, or increased maintenance workload [3,4]. Recent studies have applied machine learning, digital twins, graph-based models, and deep neural networks to abnormal-state recognition and fault detection in nuclear power systems [5]. These approaches have improved the ability to distinguish normal and abnormal operating conditions, particularly when large volumes of heterogeneous monitoring data are available. Despite these advances, accurately determining the initiating variable remains difficult when several process signals change within a short period. Multivariate anomalies in nuclear plant systems rarely occur as isolated deviations. Instead, they often develop through interacting thermal-hydraulic processes and closed-loop control responses [6,7]. A pressure reduction caused by an upstream pump may trigger changes in downstream flow rate, coolant temperature, valve position, and controller output. A diagnosis model based mainly on correlation may therefore assign the highest anomaly score to a downstream symptom rather than to the component that initiated the event. Causal inference has recently been introduced into fine-grained multivariate time-series anomaly detection to distinguish direct abnormal influences from correlated responses and to improve the localization of variables responsible for an anomalous episode [8,9]. This direction is particularly relevant to nuclear plant monitoring because fault diagnosis requires not only the detection of abnormal values but also an explanation of how deviations originate and propagate across interconnected equipment [10,11].

Existing nuclear plant fault diagnosis methods generally include rule-based alarm systems, statistical process monitoring, reconstruction-based anomaly detection, and supervised time-series classification. Rule-based methods are transparent and easy to implement, but their performance depends heavily on predefined thresholds and expert knowledge [12,13]. Weak faults may remain below fixed alarm limits, while normal operating transitions may produce nuisance alarms. Statistical methods such as principal component analysis and Gaussian process modeling can identify deviations from normal operating distributions, whereas autoencoder-based and reconstruction-based methods can capture nonlinear relationships among process variables [14,15]. However, these methods primarily model statistical dependence or reconstruction consistency. They do not necessarily distinguish whether one variable causes another variable to deviate or whether both variables respond to a separate upstream disturbance. This limitation becomes more serious in cooling and auxiliary systems, where transport delays, feedback controllers, redundant equipment, and safety interlocks can produce similar abnormal patterns under different initiating faults [16,17]. Another limitation is that many existing studies evaluate diagnostic models mainly through overall classification accuracy, precision, recall, or fault-detection rate [18,19]. These metrics are useful for measuring whether a fault category can be recognized, but they provide limited information about the practical value of a diagnosis system during plant operation and maintenance. In real applications, operators need to know which variable changed abnormally at the beginning of an event, which components were subsequently affected, and whether the observed deviations belong to a single propagation chain. Diagnosis delay is also important because a model that identifies the correct fault only after multiple subsystems have responded may provide limited support for early intervention [20,21]. In addition, inaccurate root-cause ranking can generate unnecessary inspection tickets for downstream components that are operating normally but reacting to an upstream disturbance. A practical diagnostic framework should therefore evaluate initiating-variable identification, localization time, propagation-path interpretability, and the reduction of unnecessary maintenance actions, rather than relying only on event-level detection accuracy [22]. The use of physical knowledge is also necessary for reliable causal diagnosis. Purely data-driven causal discovery may produce statistically plausible but physically impossible connections, particularly when the available data contain control actions, common operating commands, and highly synchronized measurements [23,24]. For example, a downstream coolant temperature should not be treated as an instantaneous cause of an upstream pump-speed change unless the relationship is supported by a feedback-control path. Similarly, connections that violate the direction of coolant flow, heat transfer, equipment configuration, or safety-control logic should be excluded from the causal structure. Incorporating physical constraints can reduce spurious edges, improve the stability of causal graphs, and make the resulting diagnosis more acceptable to plant engineers. Counterfactual analysis can further support fault attribution by estimating how downstream variables would have evolved if a suspected initiating deviation had not occurred.

To address these limitations, this study develops a causal fault attribution method for multivariate time-series data collected from nuclear plant cooling and auxiliary systems. The proposed method constructs a time-lagged causal graph by integrating conditional independence testing with thermal-hydraulic constraints, equipment topology, safety-control requirements, and control-loop dependency rules. A counterfactual residual module estimates the expected trajectories of process variables under normal operating behavior and measures how strongly each suspected variable contributes to subsequent deviations. Candidate root causes are then ranked according to their direct causal influence, propagation coverage, temporal precedence, and contribution to downstream residuals. Abnormal episodes are further mapped to interpretable propagation chains that connect the initiating variable with the affected equipment and process measurements. The method is evaluated using a dataset covering 12 cooling-loop subsystems, 146 process variables, 9,800 simulated operating cycles, and 14 months of auxiliary-system logs. The study is intended to determine whether causal attribution can identify initiating variables earlier and more accurately than correlation-based and reconstruction-based diagnostic methods. Its practical significance lies in shortening root-cause localization time, separating initiating faults from secondary process responses, and providing operators with traceable evidence for equipment inspection. By assigning abnormal episodes to physically interpretable propagation chains, the proposed framework may also reduce unnecessary subsystem inspection tickets and improve the consistency of maintenance decisions. More broadly, the study provides a data-driven and physically constrained approach for moving nuclear plant monitoring from abnormal-state detection toward explainable fault localization and maintenance-oriented decision support.

2. Materials and Methods

2.1. Samples and Study Area Description

This study used a nuclear plant simulation and monitoring dataset covering 12 cooling-loop subsystems and auxiliary-system components. The monitored objects included auxiliary pumps, heat exchangers, coolant pipelines, pressure vessels, valve actuators, flow-control units, and safety-related process sensors. A total of 146 process variables were recorded at 1 s intervals across 9,800 simulated operating cycles and 14 months of auxiliary-system logs. After screening, 428 million timestamped records were retained. The dataset contained 1,260 expert-reviewed abnormal episodes, including coolant flow degradation, auxiliary pump instability, valve-response delay, heat-exchanger fouling, sensor drift, and abnormal pressure oscillation. These samples covered normal operation, transient adjustment, load variation, and early cooling or auxiliary-system faults.

2.2. Experimental Design and Control Settings

The causal fault attribution method was used as the proposed approach. A temporal reconstruction model and a correlation-based anomaly detection model were used as baseline methods. All methods used the same process variables, sampling frequency, operating windows, and abnormal-event labels. The temporal reconstruction model was selected because it is widely used for multivariate time-series anomaly detection. The correlation-based model was included because it represents dependency analysis without causal direction. This setting allowed direct comparison of root-cause localization time, initiating-variable ranking, fault-propagation interpretation, and false inspection reduction. Experiments were repeated across different cooling-loop subsystems and auxiliary operating states to assess model stability.

2.3. Measurement Methods and Quality Control

Process variables were collected from simulated plant records and auxiliary-system monitoring logs. The measured variables included coolant temperature, flow rate, pressure, pump vibration, valve opening, heat-exchanger outlet temperature, auxiliary pump status, and radiation-related process indicators. All records were aligned using unified timestamps. Quality control included duplicate removal, timestamp correction, incomplete-window exclusion, and detection of physically impossible values. Abnormal episodes were checked using simulation labels, operator logs, maintenance notes, and expert review. Random checks were also performed on selected normal and abnormal windows to reduce labeling errors. Continuous variables were standardized, and discrete equipment states were encoded as binary indicators.

2.4. Data Processing and Model Formulation

Cleaned records were divided into fixed-length monitoring windows and converted into multivariate process-state vectors. Each window was defined as $x_t=[T_t, F_t, P_t, V_t, Q_t, H_t, R_t]$, where T_t denotes coolant temperature, FtF_tFt flow rate, P_t pressure, V_t valve position, Q_t pump vibration, H_t heat-exchanger state, and R_t radiation-related process indicators. For each variable, the expected value was estimated from its time-lagged causal parents:

$$\hat{x}_{i,t}=f_i\left(Pa_i^{t-\tau}, u_t\right),$$

where $Pa_i^{t-\tau}$ denotes the lagged causal parent variables of x_i, τ is the time lag, and u_t represents the operating-state condition. The normalized causal residual was then calculated as:

$$z_i(t)=\frac{x_{i,t}-\hat{x}_{i,t}}{\sigma_i+\epsilon},$$

where $x_{i,t}$ is the observed value, $\hat{x}_{i,t}$ is the expected value under normal thermal-hydraulic behavior, σ_i is the standard deviation of variable i under normal operation, and ϵ prevents numerical instability. Variables with large normalized residuals and upstream causal positions were ranked as likely root causes.

2.5. Statistical Analysis and Evaluation Metrics

Model performance was evaluated using root-cause localization time, mean reciprocal rank for initiating-variable identification, propagation-chain assignment rate, unnecessary inspection tickets, and inference latency. Root-cause localization time was defined as the interval between the start of an abnormal episode and the first correct identification of the initiating variable. Mean reciprocal rank was calculated by comparing the ranked causal-variable list with expert-reviewed fault labels. Unnecessary inspection tickets were counted when the model assigned an abnormal episode to a subsystem not confirmed by expert review. Inference latency was measured for each monitoring window. Results were averaged across abnormal types and subsystem groups to evaluate robustness under different thermal-hydraulic and auxiliary-system conditions.

3. Results and Discussion

3.1. Root-Cause Localization Performance

The proposed causal fault attribution method achieved better root-cause localization than the temporal reconstruction baseline. The median localization time decreased from 33.8 min to 8.6 min, indicating earlier identification of the initiating variable during cooling-loop and auxiliary-system faults. This improvement was mainly due to the time-lagged causal graph, which separated upstream faults from downstream thermal-hydraulic responses. Reconstruction-based models usually compare predicted and observed values, so they may rank later pressure or temperature changes as primary faults [25,26,27]. In contrast, the proposed method ranked variables according to their causal contribution to downstream deviations. The anomaly detection framework for nuclear reactor coolant pump monitoring is shown in Fig. 1 and provides a useful reference for the proposed attribution structure.

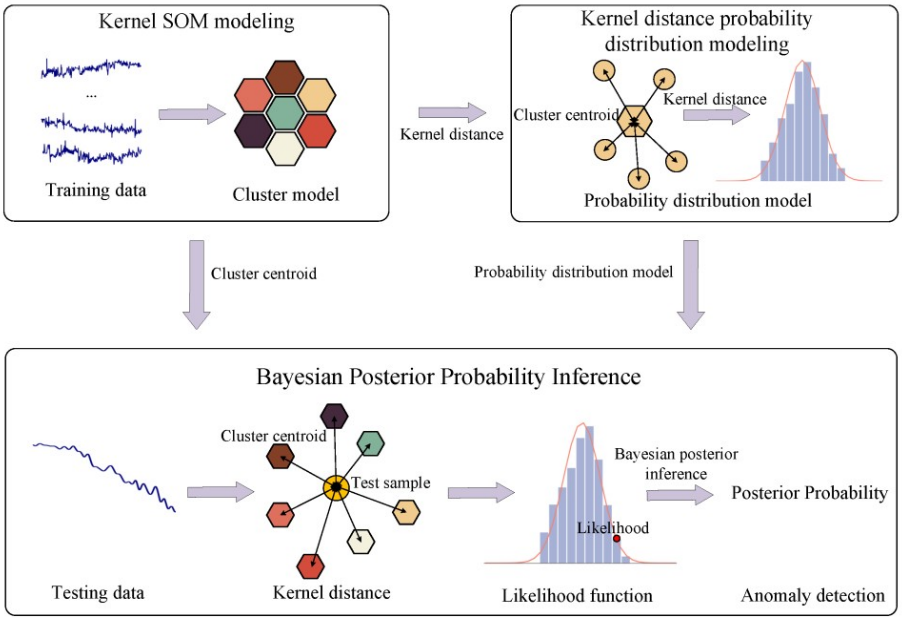

Fig. 1. Framework for anomaly detection in nuclear reactor coolant pump monitoring.

3.2. Initiating-Variable Ranking and Fault Interpretation

The mean reciprocal rank for initiating-variable identification reached 0.851, showing that the true root-cause variable was usually ranked near the top. This result is important for nuclear plant operation because early attribution can guide operators to the relevant pump, valve, heat exchanger, or sensor channel [28,29]. Causal path analysis assigned 980 abnormal episodes to clear thermal-hydraulic propagation chains. These chains explained how coolant flow degradation, auxiliary pump instability, valve-response delay, heat-exchanger fouling, sensor drift, and pressure oscillation spread across subsystems. Compared with correlation-based models, the proposed method provided clearer fault paths and reduced the risk of treating secondary responses as primary faults.

3.3. Inspection Reduction and Online Monitoring Capacity

Unnecessary subsystem inspection tickets decreased from 760 to 284, indicating that causal attribution reduced false maintenance actions while maintaining sensitivity to true faults. This reduction is useful for safety-critical plants because excessive inspections increase workload and may draw attention away from high-risk conditions. The median inference latency remained 37 ms per monitoring window, suggesting that the method can support online diagnosis. Nuclear plant signals may change gradually or abruptly under transient operation, sensor drift, or equipment degradation. Drift-aware online detection studies have shown that adapting to changing signal distributions is important for timely anomaly detection. The concept-drift-based online detection structure for nuclear plant monitoring is shown in Fig. 2.

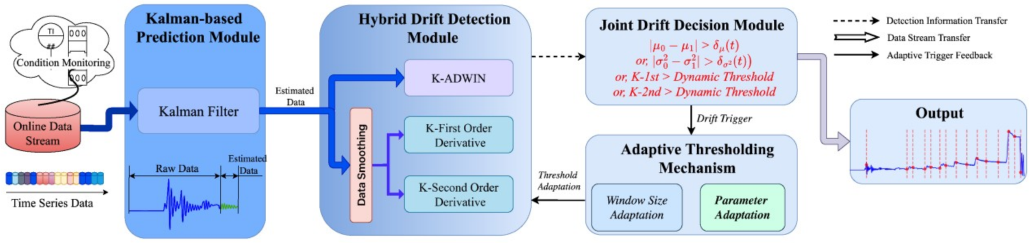

Fig. 2. Online anomaly detection framework for nuclear plant monitoring with concept drift.

3.4. Comparison with Existing Studies

Overall, the proposed method improved root-cause localization time, initiating-variable ranking, fault-chain interpretation, and inspection efficiency. Rule-based and statistical methods are useful for basic alarm generation, but they have limited ability to locate early-stage faults. Deep reconstruction models can detect abnormal windows, but they may fail to identify the first abnormal variable when thermal-hydraulic effects propagate with delay. Correlation-based methods also lack causal direction and may confuse causes with effects. By combining conditional independence testing, physical safety constraints, control-loop rules, and counterfactual residual analysis, the proposed method provided more precise and interpretable diagnosis for cooling and auxiliary systems. These results indicate that causal fault attribution is suitable for fine-grained monitoring in safety-critical nuclear plant operation [30,31,32].

4. Conclusion

This study developed a causal fault attribution method for multivariate time series from nuclear plant cooling and auxiliary systems. The method combined time-lagged causal graphs, physical safety constraints, control-loop rules, and counterfactual residual analysis to distinguish initial faults from downstream thermal-hydraulic responses. Experiments on 12 cooling-loop subsystems, 146 process variables, 9,800 simulated operating cycles, and 14 months of auxiliary-system logs showed that the method reduced median root-cause localization time from 33.8 min to 8.6 min. The mean reciprocal rank for initiating-variable identification reached 0.851. In addition, 980 abnormal episodes were linked to clear thermal-hydraulic propagation chains, and unnecessary subsystem inspection tickets decreased from 760 to 284. The median inference latency remained 37 ms per monitoring window. These results show that the proposed method provides earlier and clearer diagnosis than temporal reconstruction and correlation-based methods. The main contribution is the use of causal dependency modeling with nuclear safety constraints for fine-grained fault localization. This approach can support online monitoring, operator decision-making, targeted inspection, and safety management in cooling and auxiliary systems. However, the study used simulation data and auxiliary-system logs from a limited setting. Further validation is needed under different plant configurations, transient conditions, and long-term operating records before field deployment.

References

1. Soppelsa, A., Fedrizzi, R., & Pipiciello, M. (2026). A Two-Degree-of-Freedom Controller with Transport Delay Compensation for Application in Thermo-Hydraulic Circuits. Energies, 19(9), 2128.
2. Yang, J. (2026). Computational Analysis of How Digital Non-Clinical Communication Tools Influence Social Participation Among Older Adults in Community-Based Elderly Care. Available at SSRN 6682838.
3. Zhang, Z. (2026). A Study on the Identification of Manipulative Design in Subscription and Payment Interfaces of Digital Consumer Platforms and Its Behavioral Effects. Available at SSRN 6734760.
4. Kabir, T. (2026). Intelligent Condition Monitoring and Fault Diagnosis of Electrical Power and Control Systems Using Machine Learning–Based Predictive Analytics. American Journal of Interdisciplinary Studies, 7(01), 177-222.
5. Chen, X., Xiao, H., Zeng, Z., Zhang, S., & Xiao, R. (2025). Fine-Grained Multivariate Time Series Anomaly Detection via Causal Inference. Knowledge-Based Systems, 114765.
6. Suliman, A., Eftekhari, M., Dimitriou, V., & Ali, Y. (2025). Quantification and diagnostics of Corrosion-driven energy degradation in closed loop hydronic heating systems. Engineering Failure Analysis, 110349.
7. Fu, Y., Gui, H., Li, W., & Wang, Z. (2020, August). Virtual Material Modeling and Vibration Reduction Design of Electron Beam Imaging System. In 2020 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA) (pp. 1063-1070). IEEE.
8. Liang, R., Feifan, F. N. U., Liang, Y., & Ye, Z. (2025). Emotion-Aware Interface Adaptation in Mobile Applications Based on Color Psychology and Multimodal User State Recognition. Frontiers in Artificial Intelligence Research, 2(1), 51-57.

9. Khosravinia, P., Gama, J., & Veloso, B. (2026). Causally-Constrained Probabilistic Forecasting for Time-Series Anomaly Detection. arXiv preprint arXiv:2604.17998.
10. Yang, M., Wu, J., Tong, L., & Shi, J. (2025). Design of Advertisement Creative Optimization and Performance Enhancement System Based on Multimodal Deep Learning.
11. Wu, C., & Chen, H. (2025). Research on system service convergence architecture for AR/VR system.
12. Sayed, E., Ali, A. M., Alrashdi, I., Sallam, K. M., Abdel-Basset, M., & Ismail, M. M. (2026). Sustainable diagnostic accuracy in healthcare 4.0: a comparative study of machine learning models with association rule-based feature reduction. Cluster Computing, 29(4), 224.
13. Yin, J., Huang, Y., & Rao, H. (2026). A Study on the Dynamic Evolution of Learning Behavior and Outcome Prediction in Digital Educational Environments. Available at SSRN 6607139.
14. Hadizadeh, A., Tarighat, A., & Malian, A. (2026). A data-driven framework for structural health monitoring using reinforcement learning and deep autoencoders. Scientific Reports.
15. Karimi, K., & Ranani, E. M. (2025). A Mathematical Framework for Anomaly Detection in High-Dimensional Data Using Sparse Autoencoders and Mutual Information Filtering. International Journal of Mathematical Modelling & Computations, 15(4).
16. Yin, L., Tong, Y., Islam, Z. H., Zhang, K., Xie, R., Burger, J., ... & Wang, B. (2024). Targeted NAD+ Delivery for intimal Hyperplasia and Re-endothelialization: a novel anti-restenotic therapy approach. bioRxiv, 2024-02.
17. Du, Y., Liu, Q., Dong, Y., & Chen, Z. (2026). Cross-Domain Transfer and Few-Shot Recognition of Defect Images in Complex Industrial Settings.
18. Alagha, N., Khairuddin, A. S. M., Haitaamar, Z. N., Al-Khatib, O., & Kanesan, J. (2025). Artificial intelligence in wind turbine fault detection and diagnosis: Advances and perspectives. Energies, 18(7), 1680.
19. Xiong, W., Zeng, Y., & Guo, Y. (2026). A Study on the Impact of MEP Space Organization in Large-Scale Commercial Complexes on Investment Efficiency and Its Mechanisms.
20. Hong, Z., Xu, T., & Chen, H. (2026). A Study on Test Coverage Mapping and Release Risk Prediction in Continuous Delivery of Cloud Services.
21. Nakti, I., Mansouri, M., Al-Hmouz, R., & Khedher, A. (2025). Artificial intelligence techniques with digital twin for fault diagnosis in interconnected systems: A review. IEEE Access, 13, 91860-91874.
22. Li, J., Ma, R., & Gu, Y. (2026). From Audit Requirements to Computable Software Architecture: A Rule-Engine and Evidence-Indexing Method for Trusted Digital Infrastructure.
23. Zhang, Z., Tong, Y., & Gao, Y. (2026). Retrieval-Augmented Generation with Low-Latency Deployment for Vertical Domains Question Answering: A Case Study on Economic Resource Platforms.

24. Sanjalawe, Y., Fraihat, S., Al-E'mari, S., & Makhadmeh, S. N. (2026). Bridging the gap: A comprehensive survey on AI-driven digital twin networks for future wireless systems. Journal of King Saud University Computer and Information Sciences, 38(3), 113.

25. Qi, C., & Qiao, X. (2026). Efficient Data Sampling and Feature Selection Algorithms for Scalable Machine Learning Pipelines.

26. Su, D., & Zhang, H. Developing a Data-Driven Computational Framework for Scalable Special Education Practices for Autism and Intellectual Disabilities in the US Public School System.

27. Zhang, Y., Gu, W., & Wang, J. (2025, December). Construction of Wind Farm Asset Health Index Based on Multi-Dimensional Indicators and Analytic Hierarchy Process and Its Correlation with Operational Performance. In Proceedings of the 2025 International Conference on Digital Transformation and Management (pp. 317-323).

28. Longhi, G., Lomonaco, G., Melichar, T., & Mazzini, G. (2026). MSR Fuel and Thermohydraulic: Modeling of Energy Well Experimental Loop in TRACE Code. Energies, 19(4), 1098.

29. Xu, T., Zhu, W., & Zhang, J. (2026). Stability and Consistency of Explainable Deep Learning Methods in Credit Risk Assessment. Available at SSRN 6893938.

30. Yin, L., Tong, Y., Xie, R., Zhang, Z., Islam, Z. H., Zhang, K., ... & Wang, B. (2024). Targeted NAD+ repletion via biomimetic nanoparticle enables simultaneous management of intimal hyperplasia and accelerated re-endothelialization: a proof-of-concept study toward next-generation of endothelium-protective, anti-restenotic therapy. Journal of Controlled Release, 376, 806-815.

31. Fayyaz, Y., Elouataoui, W., Gahi, Y., El-Khatib, K., Harvel, G., & Sankaranarayanan, K. (2026). Natural language processing in the nuclear industry: opportunities and challenges. Nuclear Technology, 212(5), 1143-1163.

32. Zhang, Y., Wang, J., & Gu, W. (2026, February). Research on the Construction and Effectiveness of a Computer-Aided Systematic Training Framework for Wind Farm O&M Quality Amidst High Personnel Turnover. In Proceedings of the 2nd International Conference on Digital Management and Information Technology (pp. 647-651).