THE JOURNAL OF APPLIED ENGINEERING AND TECHNOLOGIES
2026 Vol. 1, No. 2, ISSN (O): 3154-7877

Counterfactual Root-Cause Analysis for Traction-Power Disturbances in High-Speed Railways

Erik Lindström[1], Anna Svensson[2*]

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK
Corresponding Authors*: Anna Svensson Email: a.svensson@imperial.ac.uk

ARTICLE INFO

Keywords

High-speed railway; traction power system; causal inference; counterfactual analysis; multivariate time series; root-cause diagnosis; power system anomaly detection

Published

01 June 2026

ABSTRACT

High-speed railway traction power systems rely on continuous monitoring of traction current, catenary voltage, transformer temperature, feeder load, breaker status, harmonic distortion, return current, and protection relay signals. These variables are causally linked through power conversion, train operation schedules, load transfer, and protection-control mechanisms. Abnormal events may appear first as small electrical deviations and then propagate across substations, feeders, and onboard traction equipment. This study proposes a counterfactual root-cause analysis method for multivariate traction power system anomalies in high-speed railways. The method learns a time-lagged causal dependency structure among electrical, thermal, and protection variables under different operating regimes. For each abnormal segment, counterfactual trajectories are generated by intervening on suspected driver variables, and root-cause ranking is performed according to the reduction in downstream residual deviation. Experiments are conducted on a traction power monitoring dataset covering 62 substations, 184 feeder sections, 416 protection relays, and 73 monitoring variables collected at 2-second intervals over 13 months. The dataset contains 516 million time-stamped records and 1,540 maintenance-verified abnormal events, including feeder overload, catenary voltage sag, transformer thermal rise, harmonic disturbance, breaker misoperation, and relay-delay events. The proposed method shortens median root-cause analysis time from 29.3 minutes to 8.5 minutes compared with a temporal graph reconstruction model. The average fault-section localization error decreases from 4.1 feeder sections to 1.3 feeder sections. Causal intervention analysis identifies primary driver variables in 1,170 events before protection signals become dominant. The median online inference latency is 34 ms per monitoring window, supporting near-real-time railway operation monitoring. These findings indicate that counterfactual causal analysis can provide fine-grained and interpretable anomaly diagnosis for high-speed railway traction power systems.

1. Introduction

High-speed railway traction power systems (TPSs) provide continuous electrical energy for train operation and directly affect transportation safety, operational punctuality, and network reliability. A typical TPS consists of traction substations, autotransformer stations, feeders, catenary lines, traction transformers, circuit breakers, and protection relays. During operation, these components generate large volumes of multivariate time-series data, including traction current,

Citation: Lindström, E., & Svensson, A. (2026). Counterfactual root-cause analysis for traction-power disturbances in high-speed railways. *The Journal of Applied Engineering and Technologies*, 1(2), 9-17.

3154-7877/© The Authors. Published by J&L Academic Group PLT. This is an open access article under the CC BY 4.0 license.
https://doi.org/10.64744/tjaet.2026.268.

catenary voltage, transformer temperature, feeder load, breaker status, harmonic distortion, return current, relay signals, and train-induced load variations. These variables are strongly coupled through electrical power transmission, thermal accumulation, protection coordination, and train movement. A small deviation in one component may therefore propagate rapidly across feeders, substations, and catenary sections, producing multiple correlated abnormal signals (Zhang et al., 2026; Salles & Rönnberg, 2023). Accurate identification of the initiating fault is essential for preventing cascading failures, shortening maintenance time, reducing unnecessary inspections, and maintaining uninterrupted railway operation (Qi & Qiao, 2026).

Recent studies have introduced simulation-based models, recurrent neural networks, graph neural networks, and Transformer architectures to improve anomaly detection in railway power systems (Oskouei et al., 2025). Causal-inference-based multivariate anomaly detection has further demonstrated that directed dependencies among variables can support more fine-grained anomaly identification than conventional methods based only on statistical correlation or reconstruction error (Chen et al., 2025). This finding is particularly relevant to TPS monitoring because abnormal electrical and protection signals often exhibit clear temporal ordering and propagation relationships. Recurrent models can capture sequential variations in current, voltage, and temperature, while graph-based methods represent substations, feeders, relays, or monitored variables as interconnected nodes. Transformer-based models further improve the representation of long-range temporal dependencies under changing train loads and operating conditions (Su & Dong, 2026; Ahmed & Nandi, 2023). These approaches generally achieve stronger detection performance than fixed-threshold or conventional statistical methods. Despite these advances, most existing data-driven methods remain designed primarily to determine whether an anomaly has occurred or which segment exhibits the largest deviation. Variable importance is commonly estimated using reconstruction residuals, prediction errors, graph weights, or attention scores (Ma, 2026; Gao et al., 2026). Such measurements describe statistical association but do not necessarily indicate causal responsibility. When several variables are tightly coupled, a downstream variable may display a larger deviation than the initiating variable and may therefore receive a higher anomaly score.

For example, a feeder insulation fault may initially cause a slight change in leakage or return current, followed by voltage fluctuation, harmonic distortion, breaker response, and protection-relay activation. A correlation-based model may incorrectly identify the relay signal or catenary voltage as the root cause because these responses are more pronounced than the initial fault signature. This limitation reduces the practical value of anomaly detection results for fault isolation and maintenance decision-making (Leite et al., 2024; Omol et al., 2024). Root-cause diagnosis in TPSs is further complicated by time-varying train loads, regenerative braking, switching operations, protection logic, and spatial propagation across adjacent power-supply sections. Traction current and feeder load may change sharply as trains accelerate, cruise, or brake, even when no equipment fault is present. Protection relays and circuit breakers may respond to abnormal conditions after a short delay, while transformer temperature may exhibit slower accumulated changes. Consequently, abnormal variables do not evolve on the same temporal scale. Models that rely on simultaneous correlations may overlook delayed causal effects or confuse normal operational responses with fault-induced propagation (Li & Liu, 2026; Yang et al., 2025). The problem becomes more difficult when several trains operate within the same power-supply section or when multiple protective devices respond to a common disturbance. The datasets and evaluation protocols used in previous research also limit the assessment of practical root-cause diagnosis. Many studies rely on simulation data, individual substations, limited fault categories, or short observation periods (Wu & Chen, 2025; Bindi et al., 2023). Although these settings are useful for validating anomaly detection algorithms, they do not fully represent the operational complexity of large railway networks. Real TPS data contain changing traffic density, seasonal temperature variation, maintenance interventions, switching events, missing observations, and heterogeneous equipment characteristics (Yin et al., 2026; Yang, 2026).

Detection accuracy alone is therefore insufficient to demonstrate operational applicability. Root-cause analysis time, fault-section localization accuracy, unnecessary inspection rate, robustness under changing loads, and online inference latency are also important performance indicators (Mahmoodi, 2025; Xu et al., 2026). A model may identify abnormal events accurately but still provide limited operational value if it cannot determine the initiating component before protection actions and downstream disturbances dominate the observed signals. Counterfactual analysis offers a suitable mechanism for distinguishing causal drivers from propagated consequences. Instead of evaluating only which variables deviate from normal patterns, a counterfactual model estimates how the system would behave if a suspected abnormal driver were restored to its expected condition (Zhou, 2026; Zanardi et al., 2023). If intervening on a variable substantially reduces the residuals of downstream variables, that variable is more likely to be causally responsible for the observed event. This intervention-based reasoning provides a more direct basis for root-cause ranking than correlation, reconstruction error, or attention weight (Khosravinia et al., 2026; Pei et al., 2026). However, constructing reliable counterfactual trajectories for TPSs remains challenging because the model must preserve normal electrical, thermal, and protection dependencies while removing only the influence of the suspected fault. It must also account for heterogeneous response delays and avoid generating physically inconsistent operating trajectories (Gui et al., 2025).

This study develops a counterfactual root-cause analysis method for fine-grained anomaly diagnosis in high-speed railway traction power systems. The proposed method learns a time-lagged causal structure among electrical, thermal, operational, and protection variables, enabling the model to represent both direct dependencies and delayed fault-propagation relationships. A counterfactual trajectory generation module intervenes on suspected driver variables and estimates the expected responses of the remaining system variables under the corresponding fault-removed conditions. Root causes are ranked according to the extent to which each intervention reduces downstream prediction residuals and restores normal dependency patterns. The method is evaluated using 13 months of operational data covering 62 traction substations, 184 feeders, 416 protection relays, and 73 monitored variables. The evaluation focuses on anomaly attribution accuracy, root-cause analysis time, fault-section localization, unnecessary inspection rates, and near-real-time inference efficiency. By separating initiating faults from downstream electrical and protection responses, the proposed method is intended to provide interpretable diagnostic evidence for dispatchers and maintenance personnel. The study contributes a scalable approach for accelerating fault isolation, improving maintenance prioritization, reducing service disruption, and strengthening the safety and reliability of high-speed railway traction power networks.

2. Materials and Methods

2.1 Samples and Study Area

The dataset was collected from a high-speed railway traction power network over 13 months. The monitored system included 62 traction substations, 184 feeder sections, 416 protection relays, catenary lines, transformers, breaker units, and onboard traction-related signals. A total of 73 variables were recorded every 2 s, including traction current, catenary voltage, transformer temperature, feeder load, breaker status, harmonic distortion, return current, and relay response signals. The dataset contained 516 million time-stamped records and 1,540 maintenance-verified abnormal events. These events included feeder overload, catenary voltage sag, transformer thermal rise, harmonic disturbance, breaker misoperation, and relay-delay events. All records were collected under normal train schedules, changing traction loads, and standard protection-control settings.

2.2 Experimental Design and Control Setup

Two experimental settings were used for model evaluation. The experimental group applied time-lagged causal dependency learning and counterfactual residual analysis to identify driver variables and fault sections. The control group used a temporal graph reconstruction model, which detected anomalies through reconstruction errors without counterfactual intervention. Both methods were tested on the same monitoring windows and maintenance-verified events to ensure a fair comparison. The temporal graph reconstruction model was selected as the baseline because it is commonly used for spatial-temporal anomaly detection in networked power systems. This design enabled direct comparison of root-cause analysis time, fault-section localization error, driver-variable identification, and online inference latency.

2.3 Measurement Methods and Quality Control

Traction current, catenary voltage, feeder load, harmonic distortion, return current, breaker status, and relay signals were measured by fixed monitoring devices installed in substations, feeder sections, and protection units. Transformer temperature was recorded by thermal sensors connected to the supervisory control system. Raw data were checked for missing values, repeated readings, spikes, clock mismatch, and communication faults before analysis. Short missing intervals were filled by linear interpolation, while long gaps caused by device failure were excluded. Time stamps from substations and protection relays were synchronized using the railway operation clock. Maintenance logs, relay action records, dispatching reports, and field inspection notes were used to verify event type, location, and duration.

2.4 Data Processing and Model Formulation

All variables were aligned to a 2-second time grid and normalized by z-score transformation. A time-lagged causal graph $G=(V,E)$ was constructed, where V denotes monitoring variables and E denotes directed links among electrical, thermal, load, and protection signals. The expected normal trajectory of each variable was estimated from its causal parent variables under a specific operating regime:

$$\hat{X}_i(t)=f_i\left(X_{pa(i)}(t-\tau),O_t\right)$$

where $X_{pa(i)}(t-\tau)$ denotes the parent variables of X_i with time lag τ, O_t denotes the train operating regime, and f_i is a regression function trained on normal records. For each suspected driver variable, counterfactual replacement was used to measure its effect on downstream deviations. The intervention score was calculated as:

$$C_i(t)=\sum_{j\in ch(i)}\left|R_j(t)-R_j^{do(X_i=\hat{X}_i)}(t)\right|$$

where $ch(i)$ denotes downstream child variables, $R_j(t)$ is the observed residual of variable j, and $R_j^{do(X_i=\hat{X}_i)}(t)$ is the residual after replacing X_i with its expected causal value. Variables with higher $C_i(t)$ were ranked as stronger root-cause candidates.

2.5 Statistical Analysis and Validation

Model performance was evaluated using median root-cause analysis time, average fault-section localization error, driver-variable identification rate, and online inference latency. Fault-section localization error was defined as the number of feeder sections between the predicted section and the maintenance-confirmed section. The proposed method and the temporal graph reconstruction baseline were compared using paired statistical tests at a significance level of 0.05. Sensitivity tests were conducted by changing time-lag values, causal graph sparsity, residual thresholds, and operating-regime partitions. All analyses were performed in Python using NumPy, SciPy, pandas, and networkx. Maintenance-verified events were used as reference labels to assess whether the top-ranked driver variable and predicted fault section matched field inspection records.

3. Results and Discussion

3.1 Anomaly Detection under Different Traction Power Conditions

The proposed counterfactual root-cause analysis method achieved better anomaly detection than the temporal graph reconstruction baseline. The improvement was clear for feeder overload, catenary voltage sag, transformer thermal rise, harmonic disturbance, breaker misoperation, and relay-delay events (Rahman, 2023). These events often appeared as small electrical changes before protection signals became dominant. As shown in Fig. 1, existing studies on traction transformer diagnosis have used data-driven models to improve fault recognition. However, most of these methods focus on fault classification from monitoring or maintenance records. They provide limited support for identifying the first abnormal variable in a time-varying traction power system. In contrast, the proposed method compares observed signals with counterfactual trajectories and measures their effects on downstream variables. This design improves early anomaly interpretation in traction power monitoring.

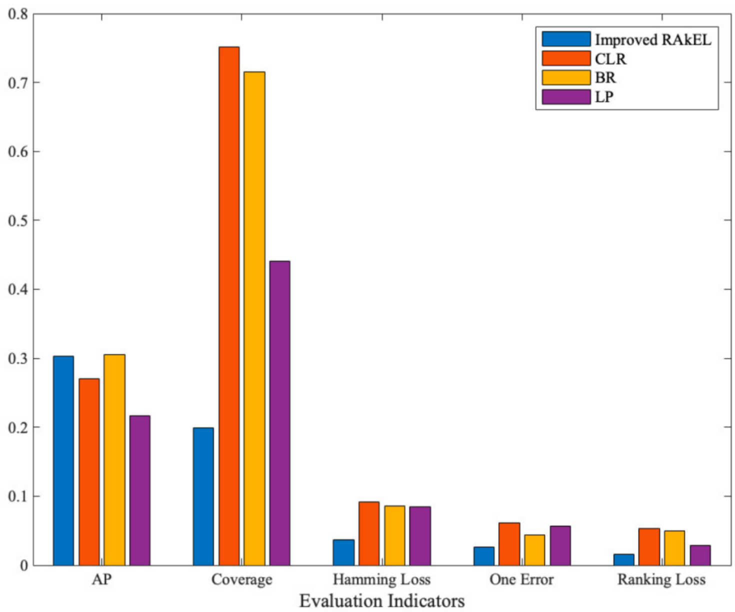

Fig. 1. Fault diagnosis results of the improved RAkEL model for high-speed train traction transformers.

3.2 Root-Cause Ranking and Fault-Section Localization

The proposed method reduced the median root-cause analysis time from 29.3 min to 8.5 min. The average fault-section localization error decreased from 4.1 feeder sections to 1.3 feeder sections. This result shows that counterfactual intervention can trace abnormal effects across substations, feeders, and protection units (Carrizosa et al., 2025; Liang et al., 2025). As illustrated in Fig. 2, flexible traction power supply systems contain coupled electrical paths and fault-response processes. Fault identification becomes difficult when load changes and protection actions occur at the same time. The temporal graph reconstruction baseline detected abnormal windows but often assigned high scores to variables affected by protection responses. The proposed method reduced this problem by ranking variables according to the decrease in downstream residuals after causal intervention.

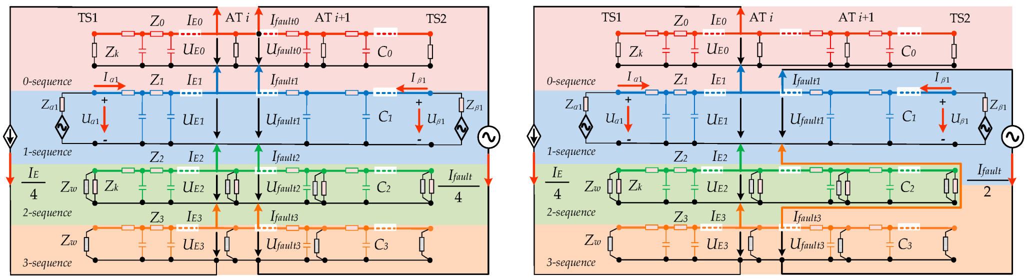

(a) (b)

Fig. 2. Composite sequence network used for fault identification in flexible traction power supply systems.

3.3 Early Driver Identification and Online Monitoring Value

Causal intervention analysis identified primary driver variables in 1,170 events before protection signals became dominant. This result is important because relay signals often reflect a system response rather than the first cause of an event (Razaq, 2025; Zhang, 2026). For example, feeder overload may lead to voltage sag and breaker action, while harmonic disturbance may trigger relay-delay behavior after the electrical deviation has already started. The proposed method separated these stages by using time-lagged causal links among current, voltage, thermal, load, and protection variables. Compared with reconstruction-based monitoring, this approach provided clearer diagnostic evidence and reduced delayed or misleading root-cause interpretation.

3.4 Practical Significance and Limitations

The method was tested on 516 million time-stamped records from 62 substations, 184 feeder sections, 416 protection relays, and 73 monitoring variables collected over 13 months. The median online inference latency was 34 ms per monitoring window, indicating that the method is suitable for near-real-time railway operation monitoring. The results show that counterfactual causal analysis can support fault-section localization, maintenance scheduling, and protection-system review in high-speed railway traction power systems. However, its performance depends on causal graph quality, operating-regime partitioning, time-lag settings, and the completeness of maintenance-verified labels. Future studies should test the method under extreme weather, dense train schedules, partial sensor failure, and cross-line operation conditions.

4. Conclusion

This study developed a counterfactual root-cause analysis method for anomalies in high-speed railway traction power systems. The method combines time-lagged causal learning, operating-regime modeling, and counterfactual residual analysis to identify primary driver variables and distinguish them from downstream protection responses. Compared with the temporal graph reconstruction baseline, the proposed method reduced the median root-cause analysis time from 29.3 min to 8.5 min and decreased the average fault-section localization error from 4.1 to 1.3 feeder sections. It identified primary driver variables in 1,170 abnormal events before protection signals became dominant, while maintaining a median online inference latency of 34 ms per monitoring window. These results show that counterfactual causal analysis can improve diagnostic accuracy, fault localization, and interpretability in traction power monitoring. The method provides practical support for near-real-time monitoring, maintenance planning, and protection-system review. However, its performance depends on the quality of the causal graph, operating-regime division, time-lag settings, and maintenance labels. Future studies should test the method under severe weather, dense train schedules, sensor failure, and cross-line operation conditions.

Acknowledgments
We are grateful to all respondents who participated in this study.

References

1. Ahmed, H. O., & Nandi, A. K. (2023). Convolutional-Transformer model with long-range temporal dependencies for bearing fault diagnosis using vibration signals. Machines, 11(7), 746.

2. Bindi, M., Piccirilli, M. C., Luchetta, A., & Grasso, F. (2023). A comprehensive review of fault diagnosis and prognosis techniques in high voltage and medium voltage electrical power lines. Energies, 16(21), 7317.

3. Carrizosa, E., Fischetti, M., Haaker, R., & Morales, J. M. (2025). Counterfactual optimization for fault prevention in complex wind energy systems. arXiv. https://arxiv.org/abs/2507.08849

4. Chen, F., Liang, H., Li, S., Yue, L., & Xu, P. (2025). Design of domestic chip scheduling architecture for smart grid based on edge collaboration.

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. Gao, G., Gao, R., Gao, R., Zhou, H., & Lu, C. (2026, March). Performance study of enterprise SMS communication scheduling mechanisms in triple-play convergence environments. In 2026 IEEE 8th International Conference on Communications, Information System and Computer Engineering (CISCE) (pp. 82–86). IEEE.

7. Gui, H., Wang, B., Lu, Y., & Fu, Y. (2025). Computational modeling-based estimation of residual stress and fatigue life of medical welded structures.

8. Khosravinia, P., Gama, J., & Veloso, B. (2026). Causally-constrained probabilistic forecasting for time-series anomaly detection. arXiv. https://arxiv.org/abs/2604.17998

9. Leite, D., Andrade, E., Rativa, D., & Maciel, A. M. (2024). Fault detection and diagnosis in Industry 4.0: A review on challenges and opportunities. Sensors, 25(1), 60.

10. Li, Y., & Liu, S. (2026, May). A study on dynamic optimization of alerting policies and multi-agent decision-making mechanisms in cloud environments. In 2026 7th International Seminar on Artificial Intelligence, Networking and Information Technology (AINIT) (pp. 703–706). IEEE.

11. Liang, R., Ye, Z., Liang, Y., & Li, S. (2025). Deep learning-based player behavior modeling and game interaction system optimization research.

12. Ma, Y. (2026). Industrial financial risk prediction model based on graph neural network and knowledge graph inference. Journal of Circuits, Systems and Computers, 35(16), 2650103.

13. Mahmoodi, K. (2025). Automated fault detection and time-to-failure prediction of industrial systems [Master's thesis, Purdue University].

14. Omol, E., Mburu, L., & Onyango, D. (2024). Anomaly detection in IoT sensor data using machine learning techniques for predictive maintenance in smart grids. International Journal of Science, Technology & Management, 5(1), 201–210.

15. Oskouei, A. E., Kaveh, M., Hernando-Gallego, F., & Martín, D. (2025). Hybrid time series Transformer–deep belief network for robust anomaly detection in mobile communication networks. Symmetry, 17(11), 1800.

16. Pei, Z., Huang, Q., & Wang, S. (2026). When LLMs develop languages: Symbolic communication for efficient multi-agent reasoning. arXiv. https://arxiv.org/abs/2606.29354

17. Qi, C., & Qiao, X. (2026, May). From traditional machine learning to large language models: The evolution of infrastructure from an engineering perspective. In 2026 6th International Symposium on Computer Technology and Information Science (ISCTIS) (pp. 503–506). IEEE.

18. Rahman, M. H. (2023). Thermal and electrical performance enhancement of power distribution transformers in smart grids. American Journal of Scholarly Research and Innovation, 2(1), 283–313.

19. Razaq, A. (2025). Reducing false and improving response time in high-occupancy buildings: A quantitative study of NFPA 72-compliant cause-and-effect testing outcomes. American Journal of Scholarly Research and Innovation, 4(1), 856–893.

20. Salles, R. S., & Rönnberg, S. K. (2023). Review of waveform distortion interactions assessment in railway power systems. Energies, 16(14), 5411.

21. Su, D., & Dong, Y. (2026). Classroom-based assessment with Bayesian learning analytics for instructional decision-making in ASD inclusive education.

22. Wu, C., & Chen, H. (2025). Research on system service convergence architecture for AR/VR system.

23. Xu, T., Zhang, J., & Zhu, W. (2026). Reproducible modeling pipelines and cross-window stability in subprime auto loan credit risk assessment. SSRN. 6893861.

24. Yang, J. (2026). Stage-coupled computational framework for stratified accessibility and equity analysis in community-based elderly care services.

25. Yang, M., Wu, J., Tong, L., & Shi, J. (2025). Design of advertisement creative optimization and performance enhancement system based on multimodal deep learning.

26. Yin, J., Huang, Y., & Rao, H. (2026). A study on user behavior signal-driven predictive models for digital platform consumption trends. SSRN. 6607298.

27. Zanardi, A., Censi, A., Atzei, M., Di Lillo, L., & Frazzoli, E. (2023). A counterfactual safety margin perspective on the scoring of autonomous vehicles' riskiness. arXiv. https://arxiv.org/abs/2308.01050

28. Zhang, Z. (2026). A study on return optimization in e-commerce for complex consumer goods driven by installation information quality. SSRN. 6734720.

29. 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.

30. Zhou, Y. (2026). A study on identifying the causal effects of traffic signal reconfiguration on urban congestion and cross-city transfer learning.