



Causal Localization of Process Anomalies in Biopharmaceutical Fermentation

Lukas Schneider¹, Anna Meier^{1*}

¹Department of Biotechnology, ETH Zürich, Zürich 8092, Switzerland

Corresponding Authors*: Anna Meier Email: a.meier@ethz.ch

ARTICLE INFO

Keywords

Biopharmaceutical fermentation; causal inference; fine-grained anomaly localization; multivariate time series; process monitoring; intervention-based residuals; batch anomaly detection

Published

01 June 2026

ABSTRACT

Biopharmaceutical fermentation processes require continuous monitoring of temperature, pH, dissolved oxygen, agitation speed, feed rate, biomass concentration, pressure, conductivity, and metabolite levels. These variables interact through complex biochemical and mechanical mechanisms, making it difficult to identify whether an abnormal signal is a true process fault or a downstream response to another variable. This study proposes a fine-grained causal anomaly localization method for biopharmaceutical fermentation time series. The method combines domain-constrained causal discovery, temporal attention, and intervention-based residual scoring. First, a causal structure is learned under process constraints such as feed-to-growth dependency, oxygen-transfer relationships, and temperature-control feedback. Second, intervention-based residuals are computed by replacing suspected abnormal variables with causally expected values. Finally, abnormal variables are ranked according to their causal contribution to batch-level process deviation. Experiments are conducted on a fermentation monitoring dataset containing 1,120 production batches, 54 process variables, and one-minute observations collected from pilot-scale and production-scale bioreactors over 19 months. The dataset includes 72 million process records and 1,380 expert-labeled abnormal segments, including dissolved-oxygen control failure, feed-pump instability, pH drift, excessive foam formation, abnormal biomass growth, and temperature-control delay. The proposed method reduces median abnormal variable localization time from 16.9 minutes to 5.4 minutes compared with a reconstruction-only Transformer baseline. The false alert rate decreases to 1.8 cases per batch. The top-ranked causal variable matches expert investigation notes in 1,026 abnormal segments, and average event-boundary deviation is reduced by 12.7 minutes. Full batch assessment requires 4.6 minutes on a single GPU. The results demonstrate that causal inference can strengthen fine-grained anomaly localization in non-linear and highly coupled biopharmaceutical time series.

1. Introduction

Biopharmaceutical fermentation is a complex dynamic process governed by closely coupled physical, chemical, and biological variables, including temperature, pH, dissolved oxygen, agitation speed, aeration rate, feed rate, biomass concentration, substrate consumption, and metabolite accumulation. These variables evolve continuously throughout the batch and influence one another through nonlinear feedback mechanisms. A small deviation in feed rate, for example, may alter biomass growth and oxygen demand, subsequently causing changes in dissolved oxygen,

Citation: Schneider, L., & Meier, A. (2026). Causal localization of process anomalies in biopharmaceutical fermentation. *The Journal of Health Science and Medical Technology*, 1(2), 1-8.

3154-7869/© 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/tjhsmt.2026.266>.

agitation speed, pH regulation, and metabolite formation. Consequently, an abnormal signal observed in one variable may represent either the initiating process disturbance or a downstream response propagated from another variable. Distinguishing between these two conditions is essential because inaccurate localization may lead to unnecessary control adjustments, delayed corrective actions, or incorrect assessments of batch quality. Reliable identification of the variables responsible for process abnormalities therefore plays an important role in maintaining fermentation stability, ensuring product consistency, supporting deviation investigation, and improving batch-release decisions (Buttrós et al., 2026; Gui et al., 2025).

Conventional fermentation monitoring methods primarily rely on statistical process control, soft sensors, multivariate regression, and data-driven predictive models. These approaches have been widely applied to estimate difficult-to-measure quality indicators, monitor critical process parameters, and detect departures from normal operating conditions (Chen et al., 2025; Hema et al., 2022). However, most conventional models are designed to determine whether the overall process is abnormal rather than to identify which variable initiated the abnormal event. Recent research has introduced causal inference into fine-grained multivariate time-series anomaly detection, demonstrating that causal relationships among variables can be used to distinguish influential abnormal variables from correlated responses and thereby improve variable-level anomaly localization (Chen et al., 2025). This finding is particularly relevant to fermentation systems, in which process variables are strongly interconnected and abnormal effects frequently propagate along physiological, material-transfer, and control-feedback pathways (Belda et al., 2025; Liang et al., 2025). Nevertheless, existing causal anomaly detection methods are generally developed for generic multivariate systems and may not adequately account for the operational constraints, stage-dependent behavior, and biological mechanisms specific to biopharmaceutical fermentation. Recent advances in multivariate time-series analysis have further improved anomaly detection by modeling temporal patterns and inter-variable dependencies (Wu & Chen, 2025; Yang et al., 2025). Graph neural networks represent process variables as nodes and their dependencies as edges, allowing abnormal information to be propagated across a learned process graph (Vatter et al., 2023; Yin et al., 2026). Transformer-based models use attention mechanisms to capture long-range temporal relationships and variable interactions, while reconstruction-based architectures identify anomalies according to prediction or reconstruction errors (Yang, 2026; Zhang, 2026). Other approaches compare ongoing batches with historical golden batches or calculate variable-level deviations from normal trajectories. These methods can detect subtle changes that may be missed by fixed-threshold monitoring and are well suited to high-dimensional industrial time-series data (Islam et al., 2026; Zhang et al., 2026). However, high reconstruction error, prediction deviation, or attention weight does not necessarily indicate that a variable is the causal origin of an anomaly. A downstream variable may exhibit a larger numerical deviation than the initiating variable because the disturbance has been amplified through process dynamics (Qi & Qiao, 2026; Su & Dong, 2026). Attention scores may also reflect statistical relevance rather than causal responsibility. As a result, existing models may correctly detect an abnormal batch while producing an inaccurate or unstable ranking of the variables responsible for the event. The difficulty becomes more pronounced in fermentation processes because variable relationships change across growth stages and operating scales (Boulton et al., 2025; Ma, 2026). During rapid biomass growth, dissolved oxygen may decline because of increasing oxygen uptake, whereas a similar decline during another stage may result from insufficient agitation, inadequate aeration, sensor drift, or changes in broth viscosity (Beltrán-Flores et al., 2023; Gao et al., 2026). Feed addition can simultaneously influence substrate concentration, biomass formation, oxygen consumption, heat generation, and metabolite

production. Temperature and pH are regulated through feedback control loops, meaning that controller actions may mask the original disturbance or generate secondary variations in related variables (Bannazadeh, 2024; Gao et al., 2026). Moreover, relationships identified from pilot-scale bioreactors may not transfer directly to production-scale systems because of differences in mixing, oxygen-transfer capacity, heat transfer, and sampling frequency. Models that rely exclusively on statistical associations may therefore produce causal directions that conflict with known process mechanisms (Li & Liu, 2026). Incorporating process knowledge into causal structure learning is necessary to prevent implausible connections, reduce spurious dependencies, and improve the interpretability of anomaly-localization results.

This study aims to develop a fine-grained causal anomaly-localization method for multivariate biopharmaceutical fermentation time series. The proposed method integrates process-constrained causal discovery, temporal attention modeling, and intervention-based residual scoring within a unified diagnostic framework. Process constraints are introduced to guide causal structure estimation according to known fermentation mechanisms, control relationships, and feasible directions of influence. Temporal attention is used to capture stage-dependent dependencies and delayed interactions among process variables, while intervention-based residual scoring evaluates how strongly each candidate variable contributes to the observed abnormal state. Instead of ranking variables solely according to reconstruction error or statistical deviation, the proposed method estimates their causal contribution and separates likely initiating variables from downstream responses. The method is evaluated using 1,120 production batches collected over 19 months, containing 54 process variables from both pilot-scale and production-scale bioreactors. The evaluation focuses on variable-localization accuracy, false-alert reduction, diagnostic consistency, and localization efficiency under different process stages and operating scales. By providing interpretable rankings of abnormal variables and revealing the pathways through which disturbances propagate, the study seeks to shorten root-cause investigation time, support more targeted corrective actions, reduce unnecessary interventions, and strengthen the reliability of fermentation monitoring. The proposed framework also provides a practical connection between causal time-series analysis and bioprocess knowledge, offering a more transparent and operationally useful approach to anomaly diagnosis, process control, product-quality assurance, and batch-release decision-making in complex biopharmaceutical manufacturing.

2. Materials and Methods

2.1 Samples and Study Area

This study used a biopharmaceutical fermentation dataset collected from pilot-scale and production-scale bioreactors over 19 months. The dataset contained 1,120 production batches and 54 process variables, including temperature, pH, dissolved oxygen, agitation speed, feed rate, biomass concentration, pressure, conductivity, and key metabolites. All variables were recorded at one-minute intervals, producing 72 million process records. A total of 1,380 abnormal segments were labeled by experienced process engineers. These segments covered dissolved-oxygen control failure, feed-pump instability, pH drift, excessive foam formation, abnormal biomass growth, and temperature-control delay. All batches followed standard operating procedures, and culture conditions were kept consistent to reduce process variation.

2.2 Experimental Design and Control Setup

Two groups were used to assess the proposed method. The experimental group used process-constrained causal discovery and intervention-based residual scoring. The control group used a

reconstruction-only Transformer model without causal information. Both groups were tested on the same fermentation batches to ensure a fair comparison. The control model represented common deep learning monitoring methods that detect abnormal patterns from reconstruction errors. This setting allowed direct comparison of localization time, false alert rate, ranking accuracy, and interpretability.

2.3 Measurement Methods and Quality Control

Process variables were measured using online sensors connected to the bioreactors and supervisory control systems. Sensors were calibrated regularly according to manufacturer instructions and internal quality standards. Raw records were checked for drift, sudden spikes, and missing values before analysis. Short missing intervals were filled by linear interpolation, while clear sensor faults were removed from the dataset. Manual sampling records and production logs were used to verify online measurements. Redundant sensors were also compared to confirm measurement consistency. These procedures helped ensure reliable data for causal analysis and model validation.

2.4 Data Processing and Model Formulation

All time-series variables were normalized with z-scores to reduce scale differences. A causal graph $G=(V,E)$ was built using conditional independence tests and process knowledge, including feed-growth dependence, oxygen-transfer links, and temperature-control feedback. For each suspected abnormal variable, the expected value under normal causal conditions was estimated as:

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

where $X_{pa(i)}$ denotes the parent variables of X_i , τ is the time lag, and f_i is a regression function trained on normal batches. The intervention residual was then calculated as:

$$R_i(t)=|X_i(t)-\hat{X}_i(t)|$$

where a larger $R_i(t)$ indicates a stronger deviation from the expected causal state. Residuals were summarized along causal paths to rank abnormal variables by their contribution to batch-level deviation.

2.5 Statistical Analysis and Validation

Model performance was assessed using mean reciprocal rank, false alert rate per batch, median localization time, and event-boundary deviation. The proposed method and the control model were compared using paired t-tests and Wilcoxon signed-rank tests at a significance level of 0.05. Sensitivity tests were conducted by changing the time-lag value and anomaly threshold. All analyses were performed in Python 3.10 using NumPy, SciPy, and networkx. Expert-labeled abnormal segments were used as reference records to validate the ranking results and confirm the practical value of the proposed method.

3. Results and Discussion

3.1 Detection Performance across Fermentation Abnormalities

The proposed causal anomaly localization method showed stable detection performance for several fermentation faults, including dissolved-oxygen control failure, feed-pump instability, pH drift, foam formation, abnormal biomass growth, and temperature-control delay. Compared with the reconstruction-only Transformer baseline, the method reduced missed detections in weak and delayed abnormal segments (Asad et al., 2026; Subhan et al., 2026). As shown in Fig. 1, fermentation variables often change together over time, which makes reconstruction error less reliable when several signals deviate at the same time. Previous fermentation monitoring studies

have shown that deep learning and synthetic time-series data can improve soft-sensor prediction. However, these methods mainly focus on signal prediction and provide limited support for causal localization. In contrast, the proposed method uses causal links and intervention residuals to identify the variable with the strongest contribution to batch-level deviation.




Fig. 1. Time-series changes in fermentation process variables during abnormal events.

3.2 Root-Cause Localization and Variable Ranking

The proposed method improved abnormal variable localization by ranking variables according to causal contribution rather than reconstruction loss. The median localization time decreased from 16.9 min to 5.4 min, and the top-ranked causal variable matched expert investigation notes in 1,026 abnormal segments. This result indicates that process-constrained causal discovery can distinguish primary faults from secondary process responses (Xu et al., 2026; Zhou, 2026). As illustrated in Fig. 2, causal graph-based anomaly detection models capture directed relations among variables, which is important in strongly coupled process systems. Existing graph-based and Transformer-based models can detect abnormal windows, but their variable rankings often depend on attention scores or prediction errors. These scores may highlight downstream variables when oxygen transfer, feeding, biomass growth, and temperature control interact. The intervention-based residual score used in this study provides a clearer basis for root-cause localization.




Fig. 2. Causal graph framework for anomaly detection and root-cause localization in multivariate time series.

3.3 False Alert Reduction and Batch-Level Reliability

The false alert rate decreased to 1.8 cases per batch, which was lower than the reconstruction-only baseline. This result is important for biopharmaceutical fermentation because frequent false alerts increase operator workload and may delay corrective action (Mitra & Murthy, 2022; Pei et al., 2026). The reduction mainly came from residual scoring along causal paths. A variable with a large deviation was not directly regarded as the root cause unless its change matched the learned causal structure and process constraints. This rule reduced the effect of random sensor noise, short-

term control fluctuation, and harmless process variation. Compared with conventional multivariate monitoring methods, the proposed method showed better stability across pilot-scale and production-scale bioreactors.

3.4 Practical Value and Remaining Limitations

The method was validated on 1,120 production batches, 54 process variables, and 1,380 expert-labeled abnormal segments collected over 19 months. Full batch assessment required 4.6 min on a single GPU, suggesting potential for near real-time monitoring. The results show that causal inference can improve detection accuracy and diagnostic clarity in nonlinear fermentation systems. The method can help process engineers shorten investigation time, reduce false alerts, and obtain interpretable abnormal variable rankings. However, several limitations remain. Localization performance depends on the learned causal graph, time-lag settings, and completeness of monitored variables. Future studies should examine adaptive causal graph updating, faster inference, and validation across more cell lines, reactor types, and production scales.

4. Conclusion

This study introduced a causal anomaly localization method for detailed detection in biopharmaceutical fermentation time series. By combining process-based causal discovery, temporal attention, and intervention-based residual scoring, the method identified abnormal variables and separated primary causes from downstream effects. Experiments on 1,120 production batches with 54 variables showed that the method reduced median localization time from 16.9 to 5.4 minutes, lowered the false alert rate to 1.8 cases per batch, and produced interpretable variable rankings that matched expert assessments in most abnormal segments. These results demonstrate that causal reasoning can improve both detection accuracy and diagnostic clarity in complex, nonlinear fermentation processes. The approach can support real-time process monitoring, process control, and batch quality management. Its performance depends on the quality of the causal graph, the choice of time-lag parameters, and the completeness of monitored variables. Future studies should explore adaptive causal graph updates, faster inference, and validation across different fermentation systems and reactor scales.

Acknowledgments

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

References

1. Asad, M., Ullah, I., Sistu, G., & Madden, M. G. (2026). Contextual graph modeling for open-set object detection in multi-view autonomous driving. *IEEE Open Journal of Vehicular Technology*.
2. Bannazadeh, F. (2024). Model predictive control for dissolved oxygen and temperature to study adeno-associated virus (AAV) production in bioreactor [Doctoral dissertation, Université d'Ottawa/University of Ottawa].
3. Belda, I., Izquierdo-Gea, S., Benitez-Dominguez, B., Ruiz, J., & Vila, J. C. (2025). Wine fermentation as a model system for microbial ecology and evolution. *Environmental Microbiology*, 27(4), e70092.

4. Beltrán-Flores, E., Tayar, S., Blaquez, P., & Sarra, M. (2023). Effect of dissolved oxygen on the degradation activity and consumption capacity of white-rot fungi. *Journal of Water Process Engineering*, 55, 104105.
5. Boulton, R., Nelson, J., & Knoesen, A. (2025). A review of sensors for the monitoring, modeling, and control of commercial wine fermentations. *Fermentation*, 11(6), 329.
6. Buttrós, V. H., Dias, D. R., & Schwan, R. F. (2026). Smarter fermenters, better foods: Automation and AI in the plant-based revolution. In *Next-gen plant based fermented foods: Innovations for a sustainable future* (pp. 263–325). Springer Nature Switzerland.
7. Chen, F., Liang, H., Li, S., Yue, L., & Xu, P. (2025). Design of domestic chip scheduling architecture for smart grid based on edge collaboration.
8. 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.
9. 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.
10. Gao, G., Gao, R., Lu, C., Gao, R., & Kuang, Y. (2026, March). Security governance methods and quantitative evaluation for enterprise SMS and verification code systems. In *2026 International Conference on Generative Artificial Intelligence and Information Security (GAIS)* (pp. 455–458). IEEE.
11. Gui, H., Wang, B., Lu, Y., & Fu, Y. (2025). Computational modeling-based estimation of residual stress and fatigue life of medical welded structures.
12. Hema, P., Sathish, E., Maheswari, M., Khosla, A., Pant, B., & Ambethkar, M. R. (2022). Robust soft sensor systems for industry: Evaluated through real-time case study. *Measurement: Sensors*, 24, 100542.
13. Islam, M. S., Hossain, R., & Chowdhury, I. (2026). Change-point detection and early warning systems. *Scientific Reports*.
14. 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.
15. Liang, R., Ye, Z., Liang, Y., & Li, S. (2025). Deep learning-based player behavior modeling and game interaction system optimization research.
16. 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.
17. Mitra, S., & Murthy, G. S. (2022). Bioreactor control systems in the biopharmaceutical industry: A critical perspective. *Systems Microbiology and Biomanufacturing*, 2(1), 91–112.
18. Pei, Z., Huang, Q., & Wang, S. (2026). When LLMs develop languages: Symbolic communication for efficient multi-agent reasoning. *arXiv preprint arXiv:2606.29354*.

19. 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.
20. Subhan, R. M., Lee, Y. D., & Koo, I. (2026). Autoencoder-based self-supervised anomaly detection in wireless sensor networks: A taxonomy-driven meta-synthesis. *Applied Sciences*, 16(3), 1448.
21. Su, D., & Dong, Y. (2026). Classroom-based assessment with Bayesian learning analytics for instructional decision-making in ASD inclusive education.
22. Vatter, J., Mayer, R., & Jacobsen, H. A. (2023). The evolution of distributed systems for graph neural networks and their origin in graph processing and deep learning: A survey. *ACM Computing Surveys*, 56(1), 1–37.
23. Wu, C., & Chen, H. (2025). Research on system service convergence architecture for AR/VR system.
24. Xu, T., Zhang, J., & Zhu, W. (2026). Reproducible modeling pipelines and cross-window stability in subprime auto loan credit risk assessment. SSRN. SSRN 6893861.
25. Yang, J. (2026). Stage-coupled computational framework for stratified accessibility and equity analysis in community-based elderly care services.
26. Yang, M., Wu, J., Tong, L., & Shi, J. (2025). Design of advertisement creative optimization and performance enhancement system based on multimodal deep learning.
27. Yin, J., Huang, Y., & Rao, H. (2026). A study on user behavior signal-driven predictive models for digital platform consumption trends. SSRN. SSRN 6607298.
28. Zhang, Z. (2026). A study on return optimization in e-commerce for complex consumer goods driven by installation information quality. SSRN. SSRN 6734720.
29. Zhang, Z., Gao, Y., & Tong, Y. (2026). Semantic-temporal graph learning for demand forecasting and resource allocation in public information systems. SSRN. SSRN 6792279.
30. Zhou, Y. (2026). A study on identifying the causal effects of traffic signal reconfiguration on urban congestion and cross-city transfer learning.