CausalPulse: Agentic Copilot for Root Cause Analysis in Smart Manufacturing

Chathurangi Shyalika[1], Utkarshani Jaimini[2], Cory Henson[3], Amit Sheth[1]

[1]Artificial Intelligence Institute, University of South Carolina
[2]Stride Lab, University of Michigan, Dearborn
[3]Bosch Center for Artificial Intelligence, Pittsburgh

Abstract

Modern manufacturing systems demand real-time, trustworthy, and interpretable insights into anomalies and their underlying causes. However, conventional pipelines treat anomaly detection, causal inference, and decision-making as siloed tasks, lacking integration, explainability, and adaptability. We present CausalPulse, an intelligent, multi-agent copilot for automated Root Cause Analysis (RCA) in industrial settings. Built on a modular and extensible architecture, the system leverages standard agentic protocols, including Model Context Protocol (MCP), Agent2Agent (A2A), and LangGraph for dynamic tool and agent discovery and seamless orchestration of tasks. Agents dynamically interact to perform data preprocessing, anomaly detection, causal discovery, and root cause analysis through a neurosymbolic workflow that combines symbolic reasoning with neural methods. Intelligent postprocessing pipelines enable automatic chaining of agent tasks, enhancing contextual awareness and adaptability. CausalPulse is evaluated using both an academic public dataset (i.e., Future Factories) and an industrial proprietary dataset (i.e., Planar Oxygen Sensor Element), and shows that the system outperforms traditional baselines in interpretability, trustworthiness, and operational utility.

Demo Video — https://tinyurl.com/yc4mrd6s

Introduction

Factories today function as highly instrumented cyber-physical systems, producing vast volumes of data that capture critical signals of anomalies and quality deviations. Despite this abundance of data, Root Cause Analysis (RCA), the process of identifying the underlying drivers of faults remains a significant challenge (Steinhauer et al. 2016; Yang, Zhang, and Hoi 2022). Traditional approaches are often heuristic-driven and slow, relying on engineer's prior experience or ad-hoc trial-and-error, which delays timely responses (Rokach and Hutter 2012). These methods are also not generalizable, as handcrafted rules and heuristics rarely transfer effectively across different machines, products, or plants (Roy et al. 2024). Current AI solutions are typically siloed, treating anomaly detection, causal discovery, and RCA as isolated tasks, which limits holistic understanding and decision-making (Shyalika et al. 2025a,b). Existing

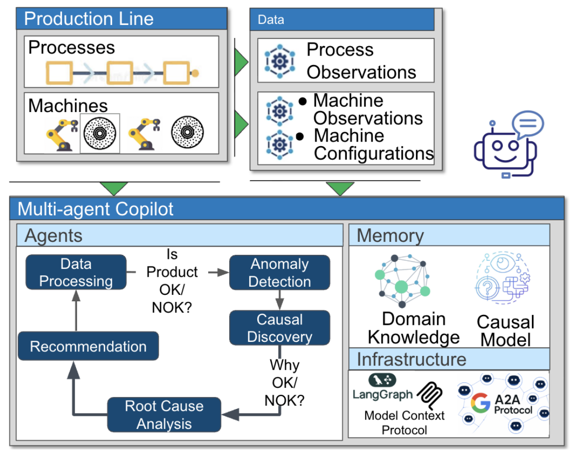

Figure 1: CausalPulse Architecture

black-box models exacerbate this issue by flagging anomalies without offering transparent reasoning or explaining how failures propagate. Real-time constraints in manufacturing demand fast, resource-efficient, and interpretable systems to prevent costly downtime, defects, and safety risks.

To address these gaps, we present **CausalPulse**, an intelligent copilot that unifies anomaly detection to RCA within a modular, agentic framework. It is designed to support human–AI collaboration, enabling factory operators to detect anomalies in real time, incorporate domain knowledge into causal discovery, trace and explain root causes, and receive context-aware recommendations for corrective action. Its overall architecture (Figure 1) integrates production-line data with multi-agent reasoning, memory, and infrastructure components to provide interpretable and real-time analysis.

System Overview

CausalPulse is designed as a four-layer agentic framework (Figure 2:a).

User-Interface (UI) Layer: Offers a browser-based interface for copilot interaction. The LangGraph-based Workflow Engine, positioned between the UI and Agent layers, comprises *Planner*, *Executor*, and *Replanner* nodes: the Planner interprets queries and determines which agents to invoke, the

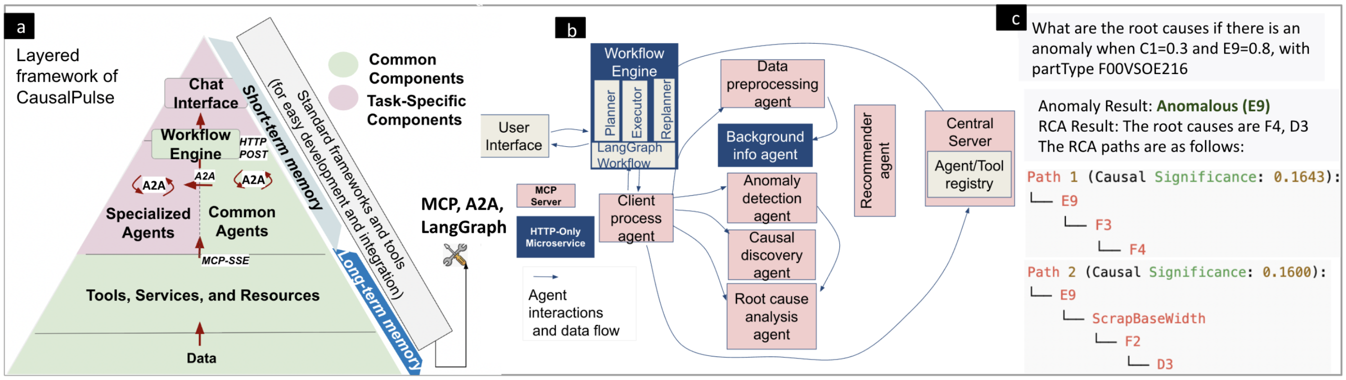

Figure 2: (a) Layered framework of CausalPulse, (b) Agent interactions and data flow, (c) Interactive interface responding to user queries on Root cause analysis showing potential failure paths, root causes and causal discovery with causal graph visualization

Executor runs them sequentially, and the Replanner adaptively revises the workflow (Figure 2:b).

Agent Layer: Comprises two categories: *common agents*, which are reusable across domains, and *specialized agents*, which are tailored for root cause analysis. Common agents include (i) *Client Process Agent*, which routes user queries to appropriate tools; (ii) *Data Preprocessing Agent*, which cleans, imputes, and encodes sensor logs; (iii) *Background Info Agent*, which enriches analysis with semantic knowledge from sensor specifications; and (iv) *Recommender Agent*, which suggests context-aware next actions through post-process chaining. Specialized agents include: (i) *Anomaly Detection Agent*, which supports detection through an LLM-based sensor–text fusion classifier (Shyalika et al. 2025c) and domain-informed thresholding; (ii) *Causal Discovery Agent*, which constructs causal graphs using the Peter-Clark (PC) (Spirtes, Glymour, and Scheines 2000) and Greedy Equivalence Search (GES) (Chickering 2002) algorithms refined with domain knowledge; and (iii) *RCA Agent*, which applies the PRORCA methodology (Dawoud and Talupula 2025) to build Structural Causal Models, and rank candidate root causes. Agents are implemented as FastAPI microservices, described by agent cards, and coordinated through A2A calls with registry-based orchestration.

Utility Layer: Provides the tools, services, and resources. *Tools* are custom MCP-based functions to support copilot tasks. To enrich reasoning and enhance interpretability, *Services* integrates domain-informed services such as process ontologies, LLMs, knowledge graphs, curated technical manuals, and rule bases. *Resources* are artifacts, including agent cards, preprocessed documents, stored causal graphs, RCA results, and precomputed causal knowledge.

Data Layer: Forms the foundation for reasoning and analysis by aggregating data from cyber-physical systems. It encompasses sensor and process logs, event and cycle metadata, sensors, process stages, and anomalies, all of which serve as inputs for higher-level agentic operations.

Demonstration

User Interfaces and Interactivity: The user interface, implemented in `Streamlit`, connects with backend implemented in `Python`, allowing operators to upload sensor data or connect directly to Programmable Logic Controller (PLC) servers for real-time inputs. Interaction with the copilot is supported through natural language queries (e.g., `preprocess the dataset`, `run causal discovery`, `find anomalies and root causes`), to which the system responds in real time (Figure 2:c). PyVis-powered visualizations present sensor variables, process states, and causal links as interactive graphs, with nodes and edges annotated by semantic metadata such as sensor type, function, and tolerance limits. Natural language understanding and explanation are powered by the `llama-3.1-8b-instant` model.

Workflow Execution and Observability: The system generates dynamic workflows that adapt to user queries, responses, and contextual state, with agents discovered and invoked as needed. Workflow planning, replanning, and tool execution are exposed through interactive terminal logs, providing transparency, traceability, and observability.

Adaptability and Extensibility: CausalPulse is scalable and demo-ready, supporting the seamless integration of new agents and services. Extending the system requires only lightweight modifications: defining an agent card, creating the corresponding MCP server and tool, implementing service logic in the stages directory, specifying input–output schemes within the *Client Process Agent*, and adding visualization routines in `main.py`. Additional flexibility is provided through the integration of expert-defined rules and optional A2A chaining for agent coordination. This modular design facilitates live demonstration of extensibility and also ensures adaptability to diverse use cases and evolving industrial requirements.

References

Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov): 507–554.

Dawoud, A.; and Talupula, S. 2025. ProRCA: A Causal Python Package for Actionable Root Cause Analysis in Real-world Business Scenarios. *arXiv preprint arXiv:2503.01475*.

Rokach, L.; and Hutter, D. 2012. Automatic discovery of the root causes for quality drift in high dimensionality manufacturing processes. *Journal of Intelligent Manufacturing*, 23(5): 1915–1930.

Roy, D.; Zhang, X.; Bhave, R.; Bansal, C.; Las-Casas, P.; Fonseca, R.; and Rajmohan, S. 2024. Exploring LLM-Based Agents for Root Cause Analysis. In *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering*, FSE 2024, 208–219. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706585.

Shyalika, C.; Prasad, R.; Al Ghazo, A.; Eswaramoorthi, D.; Kaur, H.; Muthuselvam, S. S.; and Sheth, A. 2025a. SmartPilot: A Multiagent CoPilot for Adaptive and Intelligent Manufacturing. In *2025 IEEE Conference on Artificial Intelligence (CAI)*, 1–8. IEEE.

Shyalika, C.; Prasad, R.; Al Ghazo, A.; Eswaramoorthi, D. L.; Shree Muthuselvam, S.; and Sheth, A. 2025b. SmartPilot: Agent-Based CoPilot for Intelligent Manufacturing. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, 3053–3055.

Shyalika, C.; Prasad, R.; El Kalach, F.; Venkataramanan, R.; Zand, R.; Harik, R.; and Sheth, A. 2025c. NSF-MAP: Neurosymbolic Multimodal Fusion for Robust and Interpretable Anomaly Prediction in Assembly Pipelines. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI) 2025*. Special Track on AI4Tech: AI Enabling Critical Technologies; currently available as an arXiv preprint arXttarXiv:2505.06333.

Spirtes, P.; Glymour, C. N.; and Scheines, R. 2000. *Causation, prediction, and search*. MIT press.

Steinhauer, H. J.; Karlsson, A.; Mathiason, G.; and Helldin, T. 2016. Root-cause localization using restricted boltzmann machines. In *2016 19th international conference on information fusion (FUSION)*, 248–255. IEEE.

Yang, W.; Zhang, K.; and Hoi, S. C. 2022. A causal approach to detecting multivariate time-series anomalies and root causes. *arXiv preprint arXiv:2206.15033*.