

CausalTrace: A Neurosymbolic Causal Analysis Agent for Smart Manufacturing

Chathurangi Shyalika¹, Aryaman Sharma¹, Fadi El Kalach², Utkarshani Jaimini³,
Cory Henson⁴, Ramy Harik², Amit Sheth¹

¹Artificial Intelligence Institute, University of South Carolina

²Department of Automotive Engineering, Clemson University

³Stride Lab, University of Michigan, Dearborn

⁴Bosch Center for Artificial Intelligence, Pittsburgh

Abstract

Modern manufacturing environments demand not only accurate predictions but also interpretable insights to process anomalies, root causes, and potential interventions. Existing AI systems often function as isolated black boxes, lacking the seamless integration of prediction, explanation, and causal reasoning required for a unified decision-support solution. This fragmentation limits their trustworthiness and practical utility in high-stakes industrial environments. In this work, we present CausalTrace, a neurosymbolic causal analysis module integrated into the SmartPilot industrial CoPilot. CausalTrace performs data-driven causal analysis enriched by industrial ontologies and knowledge graphs, including advanced functions such as causal discovery, counterfactual reasoning, and root cause analysis (RCA). It supports real-time operator interaction and is designed to complement existing agents by offering transparent, explainable decision support. We conducted a comprehensive evaluation of CausalTrace using multiple causal assessment methods and the C3AN framework (i.e. Custom, Compact, Composite AI with Neurosymbolic Integration), which spans principles of robustness, intelligence, and trustworthiness. In an academic rocket assembly testbed, CausalTrace achieved substantial agreement with domain experts (ROUGE-1: 0.91 in ontology QA) and strong RCA performance (MAP@3: 94%, PR@2: 97%, MRR: 0.92, Jaccard: 0.92). It also attained 4.59/5 in the C3AN evaluation, demonstrating precision and reliability for live deployment.

Introduction

Manufacturing is entering an era of hyperautonomous operations, powered by advances in AI-enabled sensing, control, and decision support (Hasan et al. 2026). As production systems grow in complexity, there is an urgent need for solutions that integrate causal reasoning with human-interpretable decision support to ensure adaptive and trustworthy operation in safety-critical environments.

While machine learning models have demonstrated impressive results in demand forecasting and anomaly detection, their lack of interpretability remains a significant barrier to adoption in high-stakes industrial settings. These black-box models often fail to meet the practical needs of shop-floor operators and subject matter experts (SMEs),

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

who require not only accurate predictions but also actionable and understandable insights into system behavior (Singh, Goyat, and Panwar 2024). This has fueled the demand for AI systems that can explain why events occur, identify underlying causes, and offer counterfactual reasoning to support ‘what if’ analysis. These capabilities are essential for real-time decision-making, proactive maintenance, and effective human-machine collaboration in production workflows (Acharya, Kuppan, and Divya 2025; Boskabadi et al. 2025). Yet, existing AI solutions rarely integrate these functionalities into a unified, deployable framework, limiting their utility in live industrial environments.

To address this gap, we present **CausalTrace**, a neurosymbolic causal analysis agent embedded in our prior work, *SmartPilot*^{1,2} industrial CoPilot (Shyalika et al. 2025a,b). SmartPilot builds on the vision of **C3AN**—*Custom, Compact, Composite AI with Neurosymbolic Integration* (Sheth et al. 2025)—an AI paradigm emphasizing robustness, intelligence, and trust of AI systems. *Custom* in C3AN refers to domain-specific data and workflows and *Compact* refers to efficient models deployable across infrastructures, including edge devices. *Composite* refers to the modular orchestration of specialized AI components—neural, symbolic, and decision modules—into a unified system. *Neurosymbolic* denotes the integration of neural learning and symbolic reasoning within each module, enabling transparent, knowledge-aligned, and explainable decision-making. Key contributions of this work are³:

- Design and implement CausalTrace, a neurosymbolic causal analysis agent capable of real-time causal discovery, root cause analysis (RCA), causal effect estimation, and counterfactual reasoning.
- Integrate data-driven causal methods with structured knowledge (ontologies, knowledge graphs), semantic user interfaces, and evaluation pipelines into neurosymbolic agent workflows, enabling human-interpretable reasoning and a deployable decision-support system.
- Develop a comprehensive evaluation methodology grounded in the C3AN framework, incorporating a suite

¹Code: <https://github.com/ChathurangiShyalika/SmartPilot>

²Demo: <https://tinyurl.com/39yy5fx4>

³Extended paper with appendices: <https://arxiv.org/pdf/2510.12033>

of techniques to ensure robustness, intelligence, and trustworthiness. Conduct causal assessments including counterfactual effect estimation, RCA validation, and comparisons with baseline and ablation variants.

- Demonstrate the deployment and practical utility of CausalTrace on an academic rocket assembly testbed, highlighting its readiness for real-world manufacturing environments.

Related Work

Causal Reasoning in Industrial AI Research on causal reasoning in industrial settings spans both symbolic and data-driven approaches. Early systems employ causal models rooted in qualitative physics and control theory to support fault localization and diagnostic reasoning (Bandekar 1989). Other efforts introduce planning-based frameworks that embed causal reasoning into multi-robot coordination and factory automation, enabling optimal plan reuse and failure diagnosis in dynamic production environments (Erdem et al. 2012). Nonparametric entropy methods and information-geometric inference have been used to reconstruct causal networks directly from process data, guiding performance prediction and process optimization in complex industrial workflows (Sun et al. 2024).

Neurosymbolic and Knowledge-Infused Systems In safety-critical domains like manufacturing and cybersecurity, neurosymbolic systems have shown promise in improving explainability and behavior under uncertainty by combining neural networks with explicit knowledge graphs (Piplai et al. 2023). In scientific domains such as geochemistry, symbolic rules extracted from expert literature via LLMs have been integrated with learning models to guide predictions and improve domain alignment (Chen et al. 2025).

Agentic AI and Multiagent Architectures in Industry Agentic AI represents a paradigm shift toward autonomous systems capable of goal-driven, adaptive decision-making. Surveys of Agentic AI emphasize its core attributes: autonomy, proactivity, learning, and reactivity as enablers for complex real-world applications (Acharya, Kuppan, and Divya 2025). Agentic systems have been integrated with digital twins to support real-time monitoring and control of nonlinear systems (Kusiak 2025; Boskabadi et al. 2025). Recent frameworks such as Intelligent Design 4.0 (Jiang et al. 2025), intent-based industrial automation (Romero and Suyama 2025), and multi-agent coordination architectures (Panigrahy 2025) propose layered systems where LLM-powered agents decompose high-level goals and orchestrate downstream execution across various functions.

Summary and Motivation While prior work in causal reasoning, neurosymbolic AI, and agentic systems has made meaningful strides, current solutions often fall short in integration, interactivity, and operational readiness. Symbolic methods struggle to scale; neural models lack transparency; and agentic systems lack semantic grounding and are not designed for operator-in-the-loop decision support. These gaps hinder adoption in complex industrial environments where explainability, adaptability, and human collaboration are critical. SmartPilot is built to address these limitations

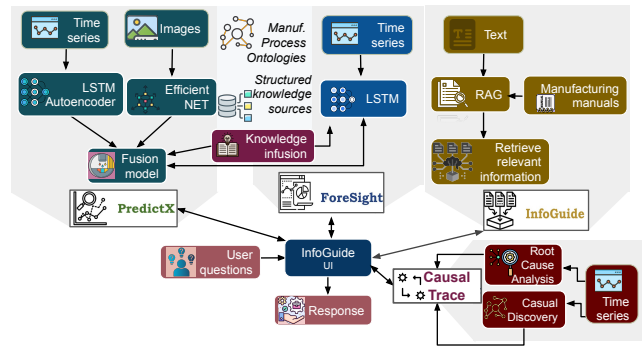


Figure 1: Multi-agent architecture of SmartPilot

holistically. By fusing causality, neurosymbolic methods, and agentic AI, it lays a practical and extensible foundation for intelligent manufacturing.

System Design

SmartPilot Architecture

The earlier version of the SmartPilot system adopts a multi-agent architecture, comprising three specialized agents that collectively support predictive analytics and operator assistance. The algorithmic details of these agents are available at (Shyalika et al. 2025c,a,b). PredictX agent specializes in anomaly prediction using multimodal sensor data (time series and images). It leverages decision-level fusion enhanced by transfer learning and knowledge-infused learning (via process ontologies) in predicting anomalies. ForeSight agent employs a Knowledge-Infused LSTM (KIL-LSTM) model to predict throughput across production cycles. It integrates time series data with domain priors to improve long-term forecasting. InfoGuide is a question-answering agent that provides real-time responses to operational, safety, maintenance, and troubleshooting queries. It uses retrieval-augmented generation over curated manufacturing manuals and integrates with other agents to address questions related to anomaly detection, production forecasting, and causal reasoning. SmartPilot’s natural language understanding and explanation are powered by LLaMA3-70B-8192 model.

Integrating the CausalTrace Agent

The CausalTrace agent extends SmartPilot’s capabilities beyond prediction to include explanation, attribution, and simulation needed for efficient RCA. This agent is tightly coupled into SmartPilot (Figure 1) and supports real-time and historical data analysis through an interactive interface. The causal analysis pipeline (Figure 2) includes the following components and features.

Data Loader and Feature Selector. SmartPilot’s Data Loader supports two modes of data ingestion: direct connection to Programmable Logic Controllers (PLCs) for real-time streaming, or batch upload of historical sensor data for offline analysis. Once data is ingested, the Feature Selector enables the selection of data-driven feature selection methods or manual selection based on domain expertise, allowing operators to prioritize variables most relevant to the analysis.

Causal Discovery Engine. The engine supports Independent Component Analysis-based Linear Non-Gaussian

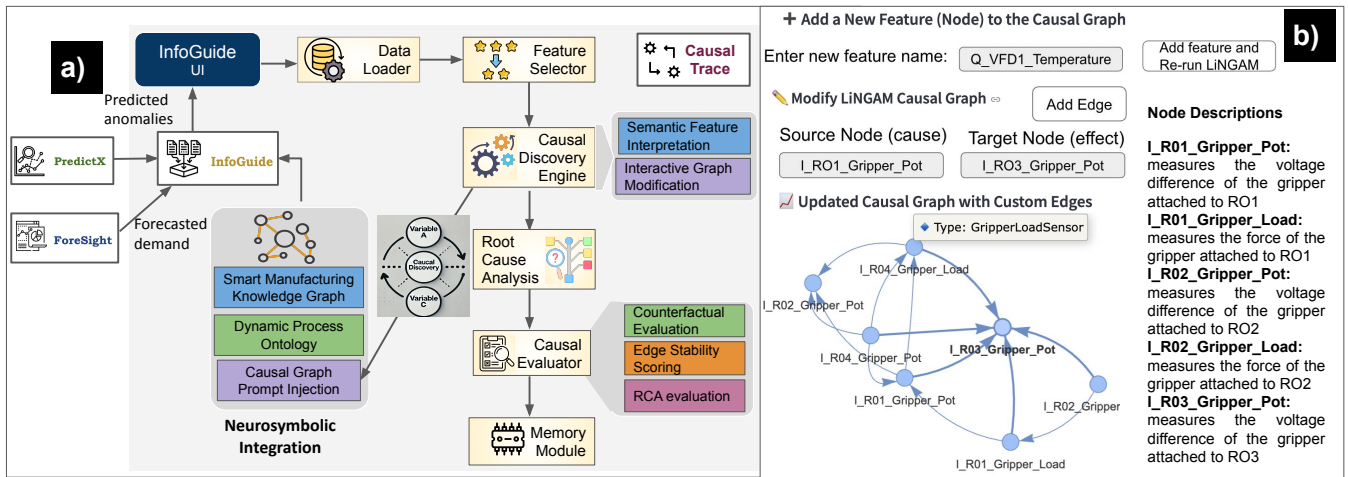


Figure 2: (a) *Architecture of CausalTrace agent*: PredictX provides predicted anomalies, ForeSight offers forecasted demand, both feeding into InfoGuide. Data flows through Data Loader and Feature Selector before entering the Causal Discovery Engine. Discovered causal graphs support Root Cause Analysis. Causal graphs and RCA results are validated in the Causal Evaluator. InfoGuide responses are enhanced via neurosymbolic integration, combining a manufacturing knowledge graph, process ontology, and causal graph prompts. Final results are stored in the Memory Module. (b) *Interactive User Interface*: Causal graph visualization and allows causal graph editing. Enhances graphs with ontology-derived descriptions, types, and units. Tooltip metadata is enriched using a smart manufacturing knowledge graph.

Acyclic Model (ICA-based LiNGAM) (Shimizu et al. 2006) and Differentiable Causal Discovery with Attention (Dif-fAN) (Sanchez et al. 2022) to construct Directed Acyclic Graphs (DAGs) from multivariate sensor data. To enhance reliability, we integrate bootstrap-based edge stability analysis (Algorithm 1) directly into the discovery workflow. Causal graphs are repeatedly estimated on bootstrap-resampled datasets, and edge stability scores computed from their occurrence frequency are used to down-weight or remove low-confidence edges (Algorithm 1). This yields more trustworthy graphs while preserving meaningful causal structure. For each retained edge, the total causal effect is computed (Algorithm 2) to quantify both direct and indirect influences between variables, and results are visualized in the user interface for interactive exploration.

Root Cause Analysis. For each anomaly, the RCA module combines expert-defined, cycle-aware sensor tolerance ranges with the learned causal topology to generate a ranked list of candidate root causes, based on causal path effect strengths and sensor deviation analysis.

Neurosymbolic Integration. SmartPilot adopts a multi-layered neurosymbolic integration framework that tightly couples structured knowledge with data-driven reasoning. The system leverages a *smart manufacturing knowledge graph* encoded in RDF to represent entities such as sensors, machines, parts, and anomalies. This knowledge is injected into the reasoning layer using `rdflib`, enriching InfoGuide’s responses with semantic context. A *dynamic process ontology*⁴, implemented in Neo4j, encodes process semantics, variable relationships, and operational constraints. Real-time `Cypher` queries enable on-demand retrieval of explanations, tolerance ranges, and sensor-function map-

⁴<https://github.com/revathyramanan/Dynamic-Process-Ontology>

Algorithm 1: Bootstrap-Based Edge Stability Analysis

Input: Dataset \mathcal{D} , number of bootstrap samples N , causal discovery method M

Output: Stability score s for each edge $A \rightarrow B$

- 1: Initialize list $\mathcal{W}_{A \rightarrow B} = []$ for each $A \rightarrow B$
- 2: **for** $i = 1$ to N **do**
- 3: Resample dataset $\mathcal{D}^{(i)}$ with replacement from \mathcal{D}
- 4: Run M on $\mathcal{D}^{(i)}$ to estimate causal graph $\mathcal{G}^{(i)}$
- 5: **for** each directed edge $A \rightarrow B$ in $\mathcal{G}^{(i)}$ **do**
- 6: Record causal strength w_i of $A \rightarrow B$
- 7: Append w_i to $\mathcal{W}_{A \rightarrow B}$
- 8: **end for**
- 9: **end for**
- 10: **for** each edge $A \rightarrow B$ with samples $\mathcal{W}_{A \rightarrow B} = w_1, \dots, w_N$ **do**
- 11: Compute mean strength: $\bar{w} = \frac{1}{N} \sum_{i=1}^N w_i$
- 12: Compute standard deviation: $\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (w_i - \bar{w})^2}$
- 13: Compute stability score: $s = \frac{1}{1+\sigma}$
- 14: **end for**
- 15: **Edge Interpretation:**
 - $s \geq 0.9$: Very strong and stable
 - $0.8 \leq s < 0.9$: Reliable
 - $0.6 \leq s < 0.8$: Moderately stable (use with caution)
 - $s < 0.6$: Unstable (typically excluded)
- 16: Retain edges with $s \geq 0.6$ for the final causal graph

ings. The *causal graph prompt injection* mechanism serializes the total causal effect matrix from the causal discovery step and passes it into LLM prompts. This allows the system to generate explanations and reasoning that are grounded in

Algorithm 2: Total Causal Effect Computation

Input: Causal effect matrix \mathbf{B} , identity matrix \mathbf{I} **Output:** Total causal effect matrix $\mathbf{T} = (\mathbf{I} - \mathbf{B})^{-1}$

- 1: Start with a structural equation model (SEM): $\mathbf{X} = \mathbf{B}\mathbf{X} + \varepsilon$, where \mathbf{X} is a vector of observed variables and ε is noise.
 - 2: Rearranged to: $(\mathbf{I} - \mathbf{B})\mathbf{X} = \varepsilon$
 - 3: Solve for \mathbf{X} using: $\mathbf{X} = (\mathbf{I} - \mathbf{B})^{-1}\varepsilon$.
 - 4: Therefore, the total effect matrix is: $\mathbf{T} = (\mathbf{I} - \mathbf{B})^{-1}$.
 - 5: Each entry T_{ij} captures total effect of variable j on variable i
 - 6: This can be decomposed as $\mathbf{T} = \mathbf{I} + \mathbf{B} + \mathbf{B}^2 + \mathbf{B}^3 + \dots$.
 - 7: Interpretation:
 - \mathbf{I} : Self influence (identity effect),
 - \mathbf{B} : Direct effects (1-hop causal influence),
 - \mathbf{B}^k : Indirect effects through k -hop paths.
 - 8: Series expansion valid when graph is acyclic or spectral radius of $\mathbf{B} < 1$
 - 9: \Rightarrow Total causal effect aggregates all direct and mediated influences
-

the structure of the causal graph.

Interactive User Interface. The interface visualizes causal graphs with rich semantic metadata derived from the process ontology and the knowledge graph. Each node is annotated with descriptions, types, units, and anomaly relevance, and tool tips display this semantic information for interactive exploration. CausalTrace is integrated with InfoGuide’s conversational query interface, enabling operators to ask natural language questions about graph structure, causal reasoning, and root cause analysis. Example queries are shown in Figure 3, with a complete set of competency questions provided in Appendix. Users can also interactively modify the discovered causal graph by adding or removing nodes and edges; all edits are validated against the ontology to ensure semantic consistency, allowing domain experts to refine and adapt the graph in real time.

Memory Module. This enables persistent, context-aware reasoning by storing and retrieving information across sessions. It consists of three components: *episodic memory*, which logs time-stamped interactions such as causal discovery and RCA runs for longitudinal tracking; *semantic memory*, which stores structured annotations of sensors and entities to support enriched, context-aware explanations; and *procedural memory*, which retains user preferences (e.g., chosen algorithms or display settings) to enable personalized interactions. These memory traces are stored in JSON format and are injected into InfoGuide’s natural language responses to ensure continuity and context-aware dialogue.

Evaluation and Results

Dataset Description. We use the publicly available manufacturing dataset generated by the Future Factories (FF) Lab at the McNair Aerospace Research Center, University of South Carolina (Harik et al. 2024). This dataset captures synchronized sensor and image data from a prototype

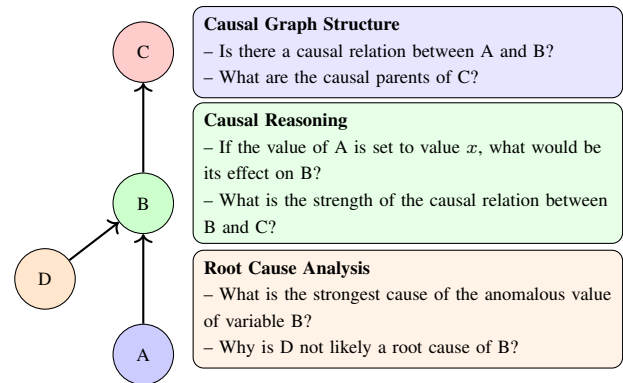


Figure 3: Illustration of a structured causal graph with categorized query types: structure, reasoning, and RCA.

rocket assembly pipeline designed to emulate industrial-grade manufacturing processes. It contains 166K records sampled at 1.95 Hz over 30 hours, covering 285 complete rocket assembly–disassembly cycles. Each cycle is segmented into 21 distinct operational states. The dataset includes time-series measurements from potentiometers, load cells, drive temperatures, and robot kinematics. The annotated version used in this work (Shyalika et al. 2025c)⁵ includes cycle state labels along with ground-truth anomaly types. These anomalies indicate missing rocket components and include six types, such as NoNose, NoBody2, and combined cases like NoBody2, NoBody1. In the evaluation, causal discovery was performed using ICA-based LiNGAM and DiffAN on eight key variables identified through XGBoost feature selection and validated by domain experts. LiNGAM produced 20 directed causal edges, while DiffAN yielded 15. The subsections below detail the evaluation methods conducted on these causal graphs.

Counterfactual Effect Computation. CausalTrace provides an interactive module for validating causal links via counterfactual effect computation. From the user interface, users can select an edge $A \rightarrow B$, apply an intervention on A , and compare the predicted change in B (from the learned total effect; Algorithm 3) with observed changes in held-out data. This enables direct plausibility checks and graph refinement based on domain knowledge.

Root Cause Analysis Evaluation. We evaluated the RCA module on anomaly events from the FF dataset with known fault cases. For each anomaly, CausalTrace produced a ranked list of candidate root causes using causal path significance and sensor deviation analysis. Here, K denotes the cut-off rank that is, the number of top predictions considered when computing the metrics. Results showed close agreement with expert-defined causes, achieving Mean Average Precision@K (MAP@3) of 94%, Precision@K (PR@2) of 97%, Mean Reciprocal Rank (MRR) of 0.92 and a Jaccard index of 0.92.

Comparison with Baseline and Ablation Variants To assess CausalTrace, we compared it against: (1) a **correlation-based root cause ranking**, which identifies candidate causes based on Pearson correlation

⁵<https://github.com/ChathurangiShyalika/NSF-MAP>

Algorithm 3: Counterfactual Validation of Causal Effects

Input: Dataset \mathcal{D} , total effect matrix τ , tolerance ϵ , effect threshold δ

Output: Validation table with $(A, B, \tau_{A \rightarrow B}, \Delta B_{\text{pred}}, \Delta B_{\text{obs}}, \text{Error})$ for each valid causal pair $A \rightarrow B$

- 1: **for** each variable pair (A, B) where $A \neq B$ and $|\tau_{A \rightarrow B}| > \delta$ **do**
 - 2: Estimate baseline (a_1) and intervention (a_2) values. If unspecified, compute:
 $a_1 = Q_1(A)$ (1st quartile), $a_2 = Q_3(A)$ (3rd quartile)
 - 3: Compute predicted change using total effect:
$$\Delta B_{\text{pred}} = (a_2 - a_1) \cdot \tau_{A \rightarrow B}$$
 - 4: Extract subsets of \mathcal{D} :
 - Baseline group: rows where $A \approx a_1$ within ϵ
 - Counterfactual group: rows where $A \approx a_2$ within ϵ
 - 5: Compute observed change in B :
$$\Delta B_{\text{obs}} = \mathbb{E}[B \mid A \approx a_2] - \mathbb{E}[B \mid A \approx a_1]$$
 - 6: Compute absolute error: $\text{Error} = |\Delta B_{\text{pred}} - \Delta B_{\text{obs}}|$
 - 7: Record result: $(A, B, \tau_{A \rightarrow B}, \Delta B_{\text{pred}}, \Delta B_{\text{obs}}, \text{Error})$
 - 8: **end for**
 - 9: Compile results into a table for global causal graph assessment
 - 10: **Interpretation:**
 - Small error \Rightarrow Causal effect from A to B is well-supported by data
 - Large error \Rightarrow Possible misspecification or unobserved confounding
-

Method	ROUGE-1	Jaccard	MAP@3	PR@2	MRR
RCA baseline	–	0.33	44%	51%	0.5
CausalTrace (no KG ontology)	0.56	–	–	–	–
CausalTrace (full)	0.91	0.92	94%	97%	0.92

Table 1: Comparison of CausalTrace and baselines

with the anomalies, and (2) a **CausalTrace variant without KG and ontology grounding**, which omits semantic tool tips, concept filtering, and knowledge-informed edge validation. Results show that the correlation baseline yields high false positives due to spurious associations and fails to capture causal directionality (Jaccard:0.33,MAP@3:44%,PR@2:51%,MRR:0.50). The ontology-ablated variant degrades interpretability and occasionally yields invalid causal paths. In contrast, full CausalTrace achieves higher agreement with expert annotations and generates explanations rated as more trustworthy by domain evaluators, with a ROUGE-1 score of 0.91 compared to 0.56 for the variant without ontology (Table 1).

C3AN Principles Evaluation. We assessed CausalTrace using the C3AN framework, which defines 14 principles for robust, intelligent, and trustworthy AI systems. A subset of 10 principles was selected based on operational relevance, assessability, and alignment with CausalTrace’s ca-

pabilities, as in Table 2. To evaluate the system against the C3AN principles, we manually crafted 10 targeted questions per principle using expert guidelines and industry best practices. The ground truth answers were obtained from two sources: (1) domain experts familiar with the manufacturing processes, and (2) official manufacturing manuals and documentation. The ground truth answer was also modeled keeping the C3AN principle in focus⁶.

To evaluate the system based on the questionnaire, we implemented an **LLM-as-a-Judge pipeline**. The pipeline consists of state-of-the-art LLMs (GPT-4o-mini and LLaMA3-70B-8192), each provided with three inputs: the question, ground truth, and the CausalTrace-generated answer. Using these inputs as the context, the LLM is prompted to evaluate the generated answer based on the ground truth and provide an integer score from 1 to 5, where 1 signifies complete misalignment with the ground truth and 5 denotes maximum alignment. The LLM is also prompted to explain its score. This enables the model to use more tokens, improving judgment accuracy and reducing bias. The LLM-as-a-Judge framework enabled efficient, low-bias evaluation of system responses. However, since it lacks special domain knowledge, CausalTrace was also assessed by six human evaluators—three experts in the manufacturing sector and three computer scientists. These evaluators scored responses on a scale of 1–5 based on alignment with the ground truth, following the same rubric. (Table 2).

The results in Table 3, indicate that GPT-4o-mini was slightly more conservative in its scoring (*Weighted Cohen’s Kappa* (κ) mean = 4.19) than LLaMA3-70B-8192 (mean = 4.32). In contrast, human evaluators rated more favorably, with an average score of 4.59. The κ values further validated the consistency and reliability of these evaluations. Agreement was generally higher within evaluator groups than across them. Specifically, GPT-4o-mini and LLaMA3-70B-8192 demonstrated substantial inter-model agreement ($\kappa = 0.56$), whereas their agreement with human evaluators was notably lower ($\kappa = 0.52$ for GPT-4o-mini and $\kappa = 0.58$ for LLaMA3-70B-8192). Human evaluators exhibited moderate internal agreement ($\kappa = 0.58$). These findings suggest that both LLMs and humans tend to align more closely with members of their own group when assessing responses against the ground truth.

Pathway To Deployment

CausalTrace is integrated as a core agent in SmartPilot, which is currently under deployment in the FF Lab. The full deployment pipeline proceeds through two primary phases: virtual and real-world deployment (Figure 4). In the initial phase, SmartPilot is virtually deployed using the Testbed as a Service (TaaS) framework (McCormick and Wuest 2025), which emulates the robotic assembly line operations. TaaS functions as a data replay mechanism, publishing data from a pre-existing dataset one timestamp at a time, mimicking the behavior of live data streams. In this setting, SmartPilot subscribes to specific MQTT (Message Queuing Telemetry Transport) topics defined by TaaS and processes the incoming data streams as if operating in a live production envi-

⁶Questions and results available at: <https://shorturl.at/pKW2>

C3AN Principle, Category and Evaluation Description	LLM1 Avg. Sc.	LLM2 Avg. Sc.	Human Avg. Sc.
Reliability [R] : Tested for hallucinations in responses.	4	4.5	4.2
Consistency [R] : Assessed stability using paraphrased queries and comparing output similarity.	3.5	4.5	4.4
Abstraction [I] : Evaluated how low-level data are mapped to high-level concepts in knowledge graph and ontology	4	3.3	4.6
Causality [I] : Tested via causal discovery, total effects, root cause analysis, and counterfactuals.	4.9	4.4	4.7
Reasoning [T] : Measured the system’s ability to logically infer conclusions.	3.8	4.5	4.6
Planning [I] : Evaluated multi-agent workflows from data input to final output.	4	5	4.6
Grounding [T] : Checked if answers used context from knowledge graph and ontology.	4.4	4.6	4.6
Attribution [T] : Verified traceability of decisions using agent logs and contextual sources.	4.5	4.5	4.2
Interpretability [T] : Users could inspect model settings, selected features, and causal graphs.	4	3.3	4.5
Explainability [T] : Judged clarity of explanations for and causal outputs.	5	5	5

Table 2: Average performance on selected C3AN principles, evaluated by two LLMs (GPT-4o-mini and LLaMA3-70B-8192) and six human evaluators. Categories: **R**= Robustness, **I**= Intelligence, **T**= Trustworthiness.

	LLM1	LLM2	E1	E2	E3	E4	E5	E6
LLM1	1							
LLM2	0.61	1						
E1	0.68	0.68	1					
E2	0.48	0.56	0.56	1				
E3	0.50	0.45	0.61	0.53	1			
E4	0.61	0.59	0.64	0.86	0.58	1		
E5	0.33	0.66	0.52	0.57	0.54	0.52	1	
E6	0.52	0.46	0.51	0.57	0.86	0.49	0.37	1

Table 3: Inter-annotator agreements (κ) (range: -1 to 1). Green: perfect agreement, yellow: moderate, red: low. LLM1: GPT-4o-mini, LLM2: LLaMA3-70B-8192, E1-E6: Human evaluators.

ronment. This virtual prototyping setup offers a controlled and repeatable environment for evaluating system behavior under a range of operational conditions.

Following successful virtual validation, SmartPilot progresses to real-world deployment. The process begins with configuring the OPC UA Server at the assembly line level. Subsequently, analog sensor data is streamed via the OPC UA Server, while image data is acquired directly from cameras positioned in the lab. This phase encompasses internal data processing tasks, including structuring raw inputs and developing a robust data ingestion pipeline. Next, the knowledge graph, dynamic process ontology, and previously generated in-memory causal graphs are instantiated via Neo4j, RDF, and HTTP SSE, respectively, and integrated with the data infrastructure. Finally, leveraging the ingested timeseries signals, visual data, and semantic insights from

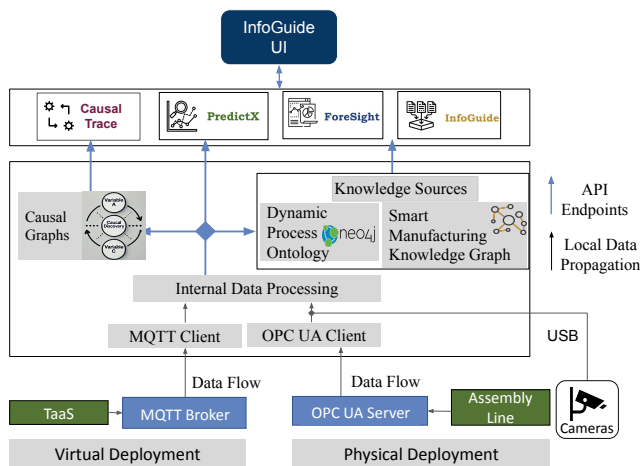


Figure 4: End-to-end deployment pipeline of SmartPilot with virtual and real-world integration

knowledge sources, SmartPilot’s four specialized agents operate in real time to perform anomaly detection, production forecasting, question answering, and causal analysis. The current deployment architecture is optimized for minimal latency in data acquisition and processing, as demonstrated in (El Kalach et al. 2025). However, its centralized nature presents scalability limitations as system complexity increases. To address this, future iterations of the system will transition to a decentralized processing model. Furthermore, the direct connection to cameras imposes spatial constraints, as the system responsible for image ingestion must remain within the physical range of the camera cabling.

In the second iteration of the deployment, the OPC UA Server and direct camera connections are to be replaced by a public MQTT broker, enabling SmartPilot to be hosted in a publicly accessible environment. This revised architecture introduces an additional layer beneath the data processing stage: a data publishing module that streams data from the local servers and cameras to the MQTT broker. SmartPilot is also publicly hosted⁷ to enable community access, leveraging both in-house and community-driven abstractions for interfacing with its back-end services. This setup enables global real-time access, improving scalability and accessibility, but it introduces latency, which we plan to address with edge computing and asynchronous communication.

Conclusion and Further Work

We introduced CausalTrace, a neurosymbolic causal analysis agent for intelligent manufacturing. Integrated with knowledge sources, it supports interpretable, trustworthy reasoning over complex industrial data. Evaluated via the C3AN framework, CausalTrace demonstrates strong alignment with robustness, intelligence, and trustworthiness. Its phased deployment from controlled virtual prototyping to live integration shows practical viability and scalability. Future work will extend capabilities such as intervention planning, continual causal graph learning, and safety-aware instruction-following features to further optimize performance and adaptability in diverse environments.

⁷<https://aiisc.ai/smartpilot/>

Acknowledgments

This work was supported in part by NSF grant #2119654, “RII Track 2 FEC: Enabling Factory to Factory (F2F) Networking for Future Manufacturing”.

References

- Acharya, D. B.; Kuppan, K.; and Divya, B. 2025. Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. *IEEE Access*, 18912 – 18936.
- Bandekar, V. R. 1989. Causal models for diagnostic reasoning. *Artificial Intelligence in Engineering*, 4(2): 79–91.
- Boskabadi, M. R.; Cao, Y.; Khadem, B.; Clements, W.; Gerek, Z. N.; Reuthe, E.; Sivaram, A.; Savoie, C. J.; and Mansouri, S. S. 2025. Industrial Agentic AI and generative modeling in complex systems. *Current Opinion in Chemical Engineering*, 48: 101150.
- Chen, W.; Zhang, J.; Li, W.; Que, X.; Li, C.; and Ma, X. 2025. Integrating Neuro-Symbolic AI and Knowledge Graph for Enhanced Geochemical Prediction in Copper Deposits. *Applied Computing and Geosciences*, 100259.
- El Kalach, F.; Farahani, M.; Wuest, T.; and Harik, R. 2025. Real-time defect detection and classification in robotic assembly lines: a machine learning framework. *Robotics and Computer-Integrated Manufacturing*, 95: 103011.
- Erdem, E.; Haspalamutgil, K.; Patoglu, V.; and Uras, T. 2012. Causality-based planning and diagnostic reasoning for cognitive factories. In *Proceedings of 2012 IEEE 17th International Conference on Emerging Technologies & Factory Automation (ETFA 2012)*, 1–8. IEEE.
- Harik, R.; Kalach, F. E.; Samaha, J.; Clark, D.; Sander, D.; Samaha, P.; Burns, L.; Yousif, I.; Gadow, V.; Tarekegne, T.; et al. 2024. Analog and Multi-modal Manufacturing Datasets Acquired on the Future Factories Platform. *arXiv preprint arXiv:2401.15544*.
- Hasan, M. K.; Latif, Z.; Hlali, A.; Xunping, L.; and Aka, S. A. B. 2026. Hyperautomation: The Next Frontier in Supply Chain (SC) and Manufacturing. In *Emerging Trends in Smart Logistics Technologies*, 1–44. IGI Global Scientific Publishing.
- Jiang, S.; Xie, M.; Chen, F. Y.; Ma, J.; and Luo, J. 2025. Intelligent Design 4.0: Paradigm Evolution Toward the Agentic AI Era. *arXiv preprint arXiv:2506.09755*.
- Kusiak, A. 2025. Agentic Manufacturing System = Digital Twin + Manufacturing. *Journal of Intelligent Manufacturing*, 36: 2221–2222.
- McCormick, M.; and Wuest, T. 2025. Testbed as a Service (TaaS): A Scalable Ecosystem for Smart Manufacturing and Industry 4.0 Collaboration.
- Panigrahy, S. 2025. Multi-Agentic AI Systems: A Comprehensive Framework for Enterprise Digital Transformation. *Journal of Computer Science and Technology Studies*, 7(6): 86–96.
- Piplai, A.; Kotal, A.; Mohseni, S.; Gaur, M.; Mittal, S.; and Joshi, A. 2023. Knowledge-enhanced Neuro-Symbolic AI for Cybersecurity and Privacy. *arXiv preprint arXiv:2308.02031*.
- Romero, M. L.; and Suyama, R. 2025. Agentic AI for Intent-Based Industrial Automation. *arXiv preprint arXiv:2506.04980*.
- Sanchez, P.; Liu, X.; O’Neil, A. Q.; and Tsafaris, S. A. 2022. Diffusion models for causal discovery via topological ordering. *arXiv preprint arXiv:2210.06201*.
- Sheth, A. P.; Roy, K.; Venkataramanan, R.; Nadimuthu, V.; and Shyalika, C. 2025. Composite AI With Custom, Compact, Neurosymbolic Models: The Emergent Enterprise Artificial Intelligence Paradigm. *IEEE Internet Computing*, 29(2): 37–49.
- Shimizu, S.; Hoyer, P. O.; Hyvärinen, A.; Kerminen, A.; and Jordan, M. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10).
- Shyalika, C.; Prasad, R.; Al Ghazo, A.; Eswaramoorthi, D. L.; Shree Muthuselvam, S.; and Sheth, A. 2025a. SmartPilot: Agent-Based CoPilot for Intelligent Manufacturing. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems*, 3053–3055.
- Shyalika, C.; Prasad, R.; Ghazo, A. A.; Eswaramoorthi, D.; Kaur, H.; Muthuselvam, S. S.; and Sheth, A. 2025b. SmartPilot: A Multiagent CoPilot for Adaptive and Intelligent Manufacturing. In *Proc. of the IEEE Conference on Artificial Intelligence*.
- Shyalika, C.; Prasad, R.; Kalach, F. E.; Venkataramanan, R.; Zand, R.; Harik, R.; and Sheth, A. 2025c. NSF-MAP: Neurosymbolic Multimodal Fusion for Robust and Interpretable Anomaly Prediction in Assembly Pipelines. *arXiv preprint arXiv:2505.06333*.
- Singh, M.; Goyat, R.; and Panwar, R. 2024. Fundamental pillars for industry 4.0 development: implementation framework and challenges in manufacturing environment. *The TQM Journal*, 36(1): 288–309.
- Sun, Y.-N.; Pan, Y.-J.; Liu, L.-L.; Gao, Z.-G.; and Qin, W. 2024. Reconstructing causal networks from data for the analysis, prediction, and optimization of complex industrial processes. *Engineering Applications of Artificial Intelligence*, 138: 109494.

Appendix

Appendix A: CausalTrace Competency Questions

This section presents a structured set of competency questions designed to evaluate the capabilities of the CausalTrace across key reasoning dimensions. These questions are grounded in the example causal graph $G = \{A \rightarrow B, B \rightarrow C, D \rightarrow B\}$, which is assumed to be generated by the CausalTrace and span four core categories: (1) Causal Graph Structure, (2) Causal Reasoning, (3) Root Cause Analysis, and (4) Causal Discovery. The goal is to assess the agent's ability to interpret, reason over, and explain causal relationships in response to natural language queries.

Questions about Causal Graph Structure

1. Does A cause B?
2. Does B cause A?
3. Does A cause C?
4. Is there a causal relation between A and B (or between B and A)?
5. Is there a causal relation between A and D?
6. What variables have a direct causal effect on B (i.e., causal parents of B)?
7. Is there a causal relation between A and Z?

Questions about Causal Reasoning

1. What is the strength of the causal relation between A and B?
2. Which variable has the strongest causal effect on B?
3. If the value of A were set to x , what would be the effect on B?

Questions about Root Cause Analysis

1. What is the most likely root cause of the anomalous value of variable B?
2. What are the three most likely root causes of the anomalous value of variable B, ranked in descending order?
3. Is A a likely root cause of the anomalous value of variable B?
4. Which is more likely to be the root cause: A or D?
5. Why is D not considered a likely root cause of the anomalous value of variable B?

Questions about Causal Discovery

1. What algorithm was used to learn the causal graph?
2. What is the stability score of the edge $A \rightarrow B$?
3. Which edges in the graph are considered the least reliable?
4. How many bootstrap iterations were used during causal discovery?