SERC Doctoral Students Research Forum · 2026

From Correlation to Intervention: Causal Root-Cause Diagnosis in Nonstationary Industrial Time Series

Cansu Yalim
cyali001@odu.edu

Engineering Management & Systems Engineering
Old Dominion University

Accepted in

Wiley International Journal of Intelligent Systems (2026)

IEEE SysCon 2026
ODU Knowledge & Creativity Expo

Industrial Fault Diagnosis: A Harder Problem to Solve

The Predictive Maintenance Chain

DETECT	Is something wrong?
DIAGNOSE	What is the root cause? GAP
PROGNOSE	How long until failure?
DECIDE	What action to take?

Why Root Cause Diagnosis Often Remains Unsolved

- **Correlation ≠ Causation**
  Downstream symptoms look like root causes in the data
- **Non-stationarity**
  Machine behavior shifts across operating regimes
- **Observational data only**
  No controlled experiments in safety-critical plants
- **Overclaiming is common**
  Models force answers even when evidence is too weak

The Root Cause Gap:
Correlation Is Not Enough for Action

Level 1

What?

Association

Variables move together

xmv(3) correlates with quality drop

Most PdM
stops here.

Level 2

What if?

Intervention

What if I change X?

If I set xmv(3)=52, does quality recover?

SCOPE OF THIS RESEARCH

Level 3

Why?

Counterfactual

Why did this happen?

Would failure have occurred without this fault?

Most sophisticated level
(We are not there yet).

The core challenge: derive intervention-level root-cause answers from purely observational, nonstationary industrial data — without experiments.

Artifact: Three-Stage Causal Diagnostic Protocol

STAGE I

Regime Segmentation

MDL joint optimization

Detect structural mechanism changes in nonstationary time series — without pre-specifying how many regimes exist.

→ **Regime boundaries + skeleton per regime**

STAGE II

Causal Representation

CPDAG construction

Convert regime-specific skeletons into honest causal graphs (CPDAGs) that preserve Markov-equivalence ambiguity.

→ **Equivalence class per regime**

STAGE III

Intervention Verification

do-intervention + support gating

Verify whether a candidate lever produces meaningful change via $P(Y|do(X=x))$ — and refuse when evidence is insufficient.

Refusal is part of the design.

→ **Ranked levers + auditable report**

STAGE I

Regime Segmentation & Skeleton Discovery

The Challenge

Segmentation and structure learning depend on each other — a chicken-and-egg problem. You cannot learn regime boundaries without knowing the causal structure, and vice versa.

The MDL Solution

Solve both simultaneously via Minimum Description Length (MDL) optimization. A new regime is admitted only when added complexity is justified by improved data compression.

- ✓ Nonparametric (Gaussian Processes) — no restrictive parametric assumptions
- ✓ Greedy search + random restarts — scalable to industrial datasets
- ✓ Each regime gets its own causal skeleton — not a blended global model

Output → Regime boundaries + regime-specific causal skeleton

STAGE II

Causal Representation — CPDAG Construction

What is a CPDAG?

A Completed Partially Directed Acyclic Graph (CPDAG) represents the full Markov equivalence class supported by the data:

→ **Directed edge:** where evidence is compelled

↔ **Undirected edge:** where ambiguity remains

Key design choice: preserve uncertainty — don't force one DAG when the data don't justify it.

Statistical Tools

- HSIC-based independence testing — nonlinear dependence
- KCIT — identifies collider structures
- Cross-fitted residual tests — directional asymmetry
- Block bootstrap calibration — accounts for serial dependence
- Benjamini–Yekutieli FDR control — multiple testing
- Meek rules — propagate admissible orientations

Output → One CPDAG per regime — honest, ambiguity-preserving causal representation

STAGE III

Interventional Root-Cause Verification

Root cause definition: A variable is a root cause if a feasible intervention on it — $P(Y \mid do(X = x))$ — produces a meaningful outcome change.

① Abduction

Fit the regime-specific model to the observed sensor history.

② Action

Replace the structural equation of the candidate variable with an assigned feasible value.

③ Prediction

Propagate the modified model forward to compute the downstream outcome distribution.

Guardrails: Support gating flags extrapolations. Structural averaging across all CPDAG-consistent DAGs propagates uncertainty. **Refuses when the evidence is insufficient.**
Three validity conditions must hold: causal ancestry, identifiability, and empirical support.

Output → Ranked actionable levers + uncertainty bounds + 'refuse' flags + auditable evidence trail

✓ AFFIRMATIVE CASE

Fault 6 — A-Feed Loss: Protocol Confirms Root Cause

STAGE I — 3 Regimes Detected

R1 Pre-fault	Mean 4.20
t = 0–180	Std 0.56
R2 Onset	Mean 3.45
t = 180–360	Std 0.05
R3 Fault steady	Mean 3.46
t=360–500	Std 0.06

Sharp change at t=180

Sharp shift at t=180: mean drops 0.75, variance collapses 10×

STAGE II — Causal Structure

xmv(3)

✓ Admissible upstream actuator

Directed ancestor of xmeas(35) across ALL valid DAG extensions

Other xmv(*k*)

✕ Excluded

No directed path to xmeas(35) in any valid DAG

xmeas(·) vars

✕ Downstream only

Correctly labeled as symptoms, not root cause candidates

STAGE III — Intervention Result

Dose → Response: xmeas(35)

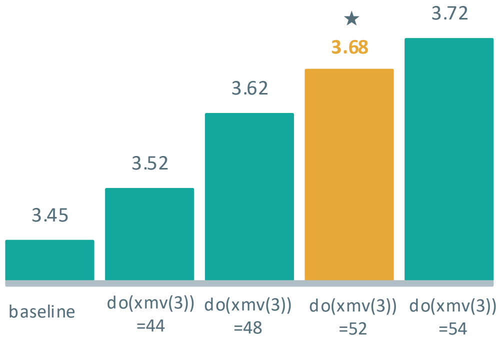

+6.7% improvement · 95% CI excludes zero

Fault 6 is a direct actuator fault. One valve fails. The signal is clear. The protocol found the root cause.

VERDICT: ✓ Root cause identified and verified — xmv(3) is the confirmed actionable lever.

CONSERVATIVE CASE

Fault 1 — A/C-Feed Ratio: Protocol Refuses to Overclaim

STAGE I — Regimes Detected

3 regimes found
t = 0–180, 180–360, 360–500

Shifts are subtle
Statistically weaker than Fault 6

Boundaries less sharp
Mechanism change is ambiguous

No sharp structural break identified at mechanism level

STAGE II — Structural Finding

No actuator xmv(*k*)
× **No directed ancestor of xmeas(35)**
in any valid DAG extension

Other xmv(k)
× Also excluded — no valid path

xmeas(·) vars
× Downstream only — correctly excluded

Graph structurally indeterminate for action

STAGE III — Intervention Estimate

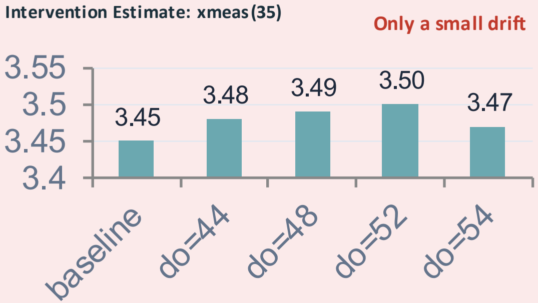

CI crosses zero · Support gating flags insufficient evidence

× Diagnosis explicitly withheld

Fault 1 is a ratio and composition disturbance (subtle). No single actuator is directly responsible. The signal is not clear. The protocol found no root cause.

Refusal is not failure. Withholding a diagnosis when evidence is too weak is the scientifically honest and safety-critical correct output. **Overclaiming is the real failure.**

Scientific Maturity: Knowing When NOT to Claim a Root Cause

Fault 1 is a stress test, and the protocol's correct refusal is a scientific contribution, not a failure.

Fault 6 — Affirmative Case

Stage I	3 clear regimes — sharp mechanism shifts confirmed
Stage II	xmv(3) directed ancestor of xmeas(35) in all valid DAGs
Stage III	do(xmv(3)=52) → +6.7% improvement, CI excludes zero
Verdict	✓ **Root cause identified and verified**

Fault 1 — Conservative Case

Stage I	3 regimes found — but shifts are subtle, statistically weaker
Stage II	No actuator with directed path to xmeas(35) in any valid DAG
Stage III	Intervention effect: modest drift, on-support doses only, withheld
Verdict	✓ **Overclaiming prevented — withheld, with explanation**

Both outcomes are success conditions. The protocol addresses all identified gaps: regime awareness, uncertainty preservation, and validity-guarded diagnosis.

The Output: An Auditable, Decision-Ready Diagnostic Report

CAUSAL DIAGNOSTIC REPORT · Tennessee Eastman Process · xmeas(35) = Product Quality

SCENARIO CONTEXT

Fault:	Fault 6 (A-feed loss)
Regimes:	3 detected (R1, R2, R3)
Change pts:	t=180, t=360
Outcome var:	xmeas(35)
Nonstationarity:	Confirmed

STRUCTURAL FINDING

Main lever:	xmv(3) ✓ Admissible
Path status:	Directed ancestor (all DAGs)
Regime stable?:	Yes — R1, R2, R3
Competing xmv(k):	Excluded (no path)
xmeas(·) as root?:	Rejected (downstream)

INTERVENTION VERDICT

Action:	do(xmv(3) = 52)
Baseline:	xmeas(35) = 3.45
Post-intervention:	xmeas(35) = 3.68
Effect size:	+6.7% (CI excludes 0)
Validity status:	✓ SUPPORTED

Every conclusion is traceable: regime context → structural evidence → intervention estimate → validity status. Engineers can ask 'Why was this claimed?' or 'Why was this withheld?' — and the report answers.

Contributions That Move the Field Forward

01	**A new diagnostic paradigm**	The first regime-aware, intervention-oriented causal diagnostic protocol for nonstationary industrial time series — unifying temporal, structural, and causal reasoning in one artifact.
02	**Joint segmentation & structure learning**	Breaks the circular dependency between regime boundaries and mechanism discovery by solving both simultaneously via MDL optimization.
03	**Uncertainty-preserving causal representation**	CPDAG output instead of a forced DAG — structural ambiguity is preserved and carried forward, not silently erased.
04	**Intervention-grounded root-cause definition**	A variable is a root cause only if do(X=x) produces a meaningful outcome change under validity constraints — not merely because of correlation.
05	**Principled refusal as a scientific output**	The protocol explicitly withholds judgment when data or structure cannot support a claim — treating epistemic discipline as a design requirement, not a limitation.

Accepted: Wiley Int'l Journal of Intelligent Systems (2026) · IEEE SysCon 2026 · ODU Knowledge & Creativity Expo

Why This Matters: Safer Diagnosis in Safety-Critical Systems

Beyond correlation

Moves predictive maintenance from pattern-matching to intervention-oriented causal reasoning — enabling engineers to act on root causes, not symptoms.

Safer claims

In complex plants, downstream symptoms can easily masquerade as root causes. Explicit guardrails protect engineers from acting on plausible-but-wrong conclusions.

Decision-ready under uncertainty

Outputs are not just predictions — they are ranked levers with uncertainty quantification, refusal flags, and auditable reasoning trails.

Reusable template

The three-stage protocol generalizes to any safety- or reliability-critical system where operating conditions change over time — aerospace, energy, manufacturing.

The Takeaway

This dissertation develops a complete, operationalizable protocol for moving from messy observational data to a causal diagnosis that is actually useful for engineering action.

- Learns regime-specific causal context — no single global model forced onto shifting data
- Preserves structural ambiguity honestly — CPDAGs instead of arbitrary DAGs
- **Verifies root causes through intervention — $P(Y|do(X=x))$, not just correlation**
- Enforces support and identifiability guardrails — refuses to overclaim
- Delivers auditable, decision-ready outputs for maintenance engineers

Questions?

Cansu Yalim
cyali001@odu.edu

Engineering Management & Systems Engineering · Old Dominion University

Solution is a Protocol
3-Stage Causal Diagnostic

Method
MDL + CPDAG + intervention

Benchmark
Tennessee Eastman Process

Accepted
Wiley IJIS (2026)