



UNSW
SYDNEY



Where Root Cause Analysis Fails: A Retrieval-Reranking Decomposition

Hada Melino M[1] · Luan Pham[1] · Laure Barrière[2] · Sachin Shetty[2] · Leonardo Pulga[2] · Flora Salim[1]

[1]University of New South Wales · [2]Baker Hughes

Root Cause Analysis Ranks Metrics by Their Likelihood of Being the Fault Origin

- Complex monitored systems, from cyber-physical systems (CPS) such as water treatment and building management to cloud microservices, expose hundreds to thousands of metrics
- Given a detected anomaly, RCA produces a ranking π over the p metrics
- Methods are evaluated by top@k, the fraction of faults whose root cause appears in the first k positions

Root Cause Identification Reduces to Retrieval and Reranking

- With an oracle detector, the root cause is the *earliest* deviating metric at t_{fault}

$$m^* = \arg\min_{s_i \in \mathcal{S}} \tau(s_i)$$

- In practice, no oracle provides t_{fault}, anomaly detectors observe t_{det}, the time at which a deviation is detected, which is a noisy and delayed proxy for t_{fault}

$$t_{\text{det}}(s_i) < t_{\text{det}}(s_j) \not\Rightarrow \tau(s_i) < \tau(s_j)$$

Root Cause Identification Reduces to Retrieval and Reranking

- Since τ is unobservable, methods score metrics with an observable proxy φ and retain $\mathcal{C}$ = top@k(φ)
- Examples of φ in prior work:
  - Deviation magnitude from the baseline window (BARO)
  - Distribution shift between baseline and fault windows (ε-Diagnosis)
  - Conditional dependence on the fault indicator (RCD)
  - Node centrality in a causal graph learned per scenario (PageRank, RandomWalk over PC/FCI graphs)

top@k Conflates Retrieval Failure and Reranking Failure

- Retrieval failure: the root cause is absent from the candidate set, m* $\notin \mathcal{C}$
- Reranking failure: the root cause is in $\mathcal{C}$ but ranked too low
- A method that **never** retrieves m* and one that ranks it **sixth** receive **identical top@5**

Datasets and Experiment Setup

Datasets:

- CPS: WADI, SWaT, HVAC (LBNL FDD, rooftop unit)
- Microservices: RCAEval suites

Baselines:

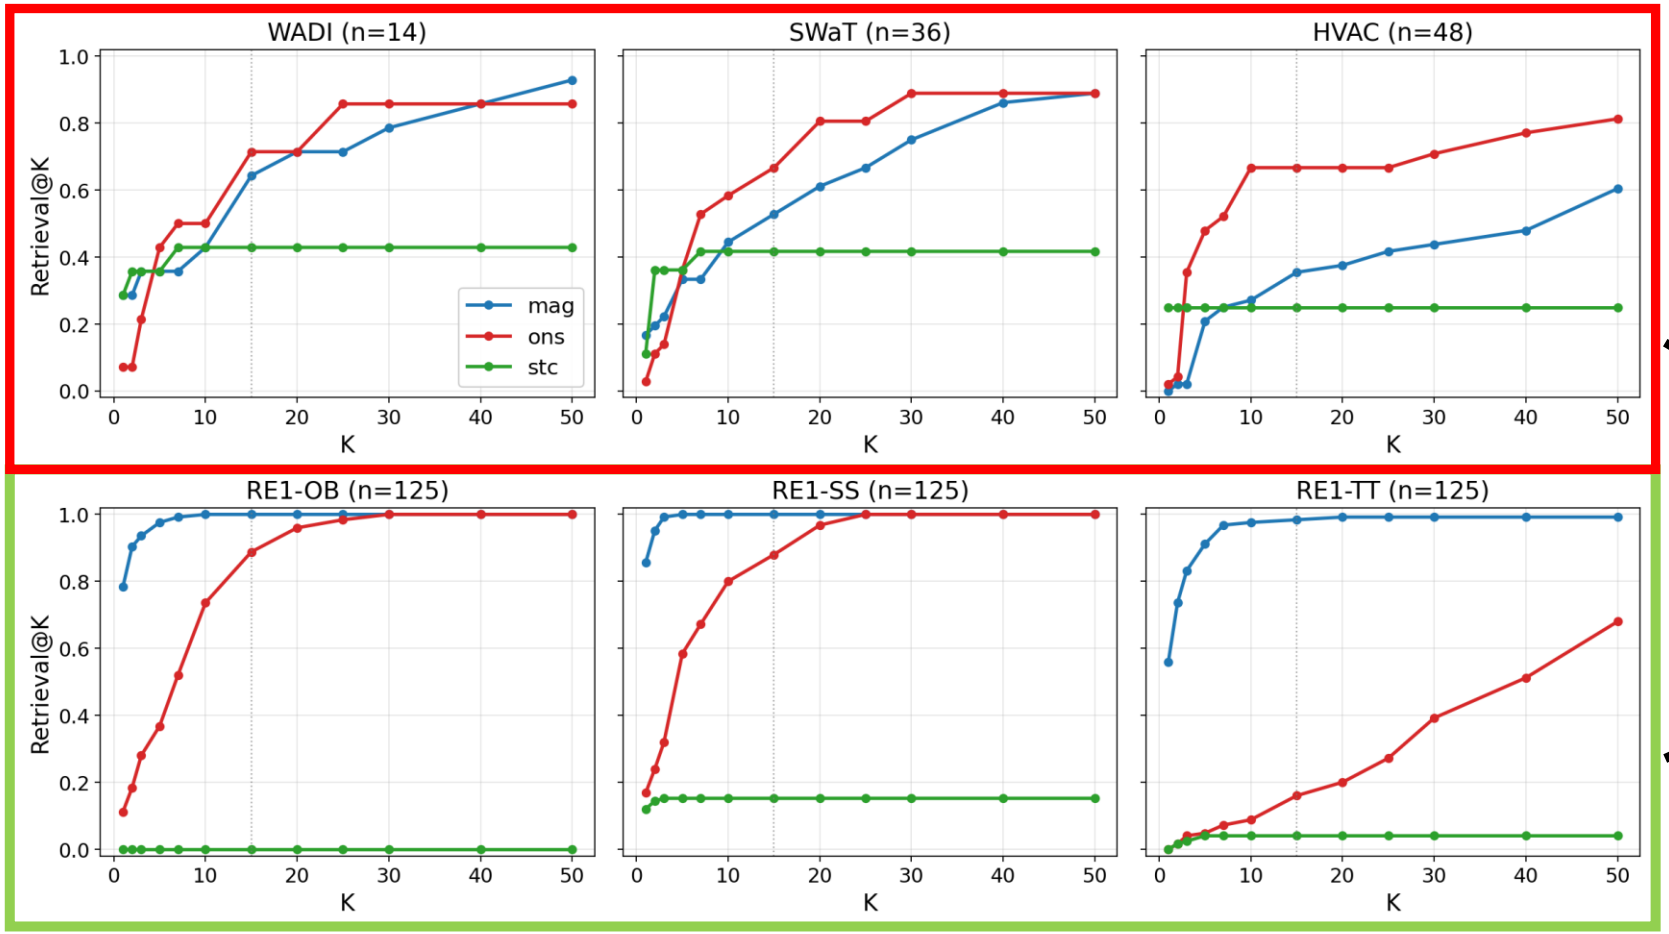
- Statistical: BARO, RCD, ε-Diagnosis
- Graph-based: PageRank, RandomWalk and CIRCA on PC/FCI graphs

LLM reranker: gpt-oss-120b (three runs), full prompts on github

Retrieval Failure: Near Zero on Microservices and Substantial on CPS

- We combine three retrieval signals, each targeting a distinct fault regime
  - Deviation magnitude (mag): the root cause shows the largest amplitude change, as in microservice faults
  - Earliest onset (ons): the root cause deviates first with suppressed amplitude, as in gradual drifts (HVAC)
  - Discrete state-change (stc): the root cause is a binary actuator or mode switch (SWaT)

Retrieval Failure: Near Zero on Microservices and Substantial on CPS



CPS

Microservices

Designed signals (φ):

mag: magnitude

ons: onset

stc: state change

Reranking Failure: When the Candidate Set is Guaranteed to Contain the Root Cause, the LLM Reranker is Never Worse Than Any Baseline

	Retriever pool (K=15)				All candidates		
Group	BB	LLM (no DK)	LLM (with DK)	Δ	BB	LLM (no DK)	Δ
WADI	0.308	0.410±0.044	**0.436**±0.089	+12.8	0.231	**0.385**±0.077	+15.4
SWaT	0.200	0.229±0.029	**0.267**±0.017	+6.7	0.200	**0.210**±0.017	+1.0
HVAC	0.188	0.188±0.055	**0.326**±0.012	+13.9	0.146	**0.299**±0.032	+15.3
RE1-OB	0.784	0.883±0.005	**0.891**±0.012	+10.7	0.784	**0.896**±0.014	+11.2
RE1-SS	0.856	0.859±0.009	**0.928**±0.000	+7.2	0.856	**0.877**±0.005	+2.1
RE1-TT	0.560	0.688±0.024	**0.741**±0.009	+18.1	0.560	**0.795**±0.020	+23.5
Mean	0.483	0.543	**0.598**	+11.6	0.463	**0.577**	+11.4

End-to-End: The LLM Reranker Matches or Exceeds the Best Baselines on All Six Benchmarks

	Industrial CPS			Microservices (RCAEval)		
Method	WADI	SWaT	HVAC	RE1-OB	RE1-SS	RE1-TT
BARO	0.214	0.194	0.000	0.784	0.856	0.560
RCD	0.071	0.083	0.083	0.304	0.248	0.144
ϵ-Diagnosis	0.000	0.028	0.146	0.072	0.232	0.008
CIRCA (PC)	0.214	0.083	0.062	0.536	0.632	0.320
CIRCA (FCI)	0.214	0.111	0.104	0.576	0.624	0.328
PageRank (PC)	0.000	0.028	0.021	0.112	0.136	0.008
PageRank (FCI)	0.071	0.000	0.000	0.096	0.184	0.048
RandomWalk (PC)	0.071	0.000	0.000	0.080	0.112	0.064
RandomWalk (FCI)	0.000	0.000	0.000	0.112	0.144	0.072
LLM (no DK)	**0.333**±0.082	**0.213**±0.016	0.153±0.043	0.875±0.009	0.872±0.024	0.653±0.024
LLM (with DK)	0.310±0.041	0.148±0.016	**0.278**±0.012	**0.888**±0.008	**0.944**±0.008	**0.685**±0.009

Baker Hughes

UNSW SYDNEY

top@k Can Hide a Component's Benefit: Domain Knowledge on SWaT

- End to end, DK lowers SWaT top@1 by 6.5 pp (0.213 -> 0.148)
- With retrieval guaranteed, DK raises SWaT top@1 (0.229 -> 0.267)
- DK helps on subtle causes that the retriever misses, which score zero end to end regardless, so only its cost on pronounced-evidence causes is visible

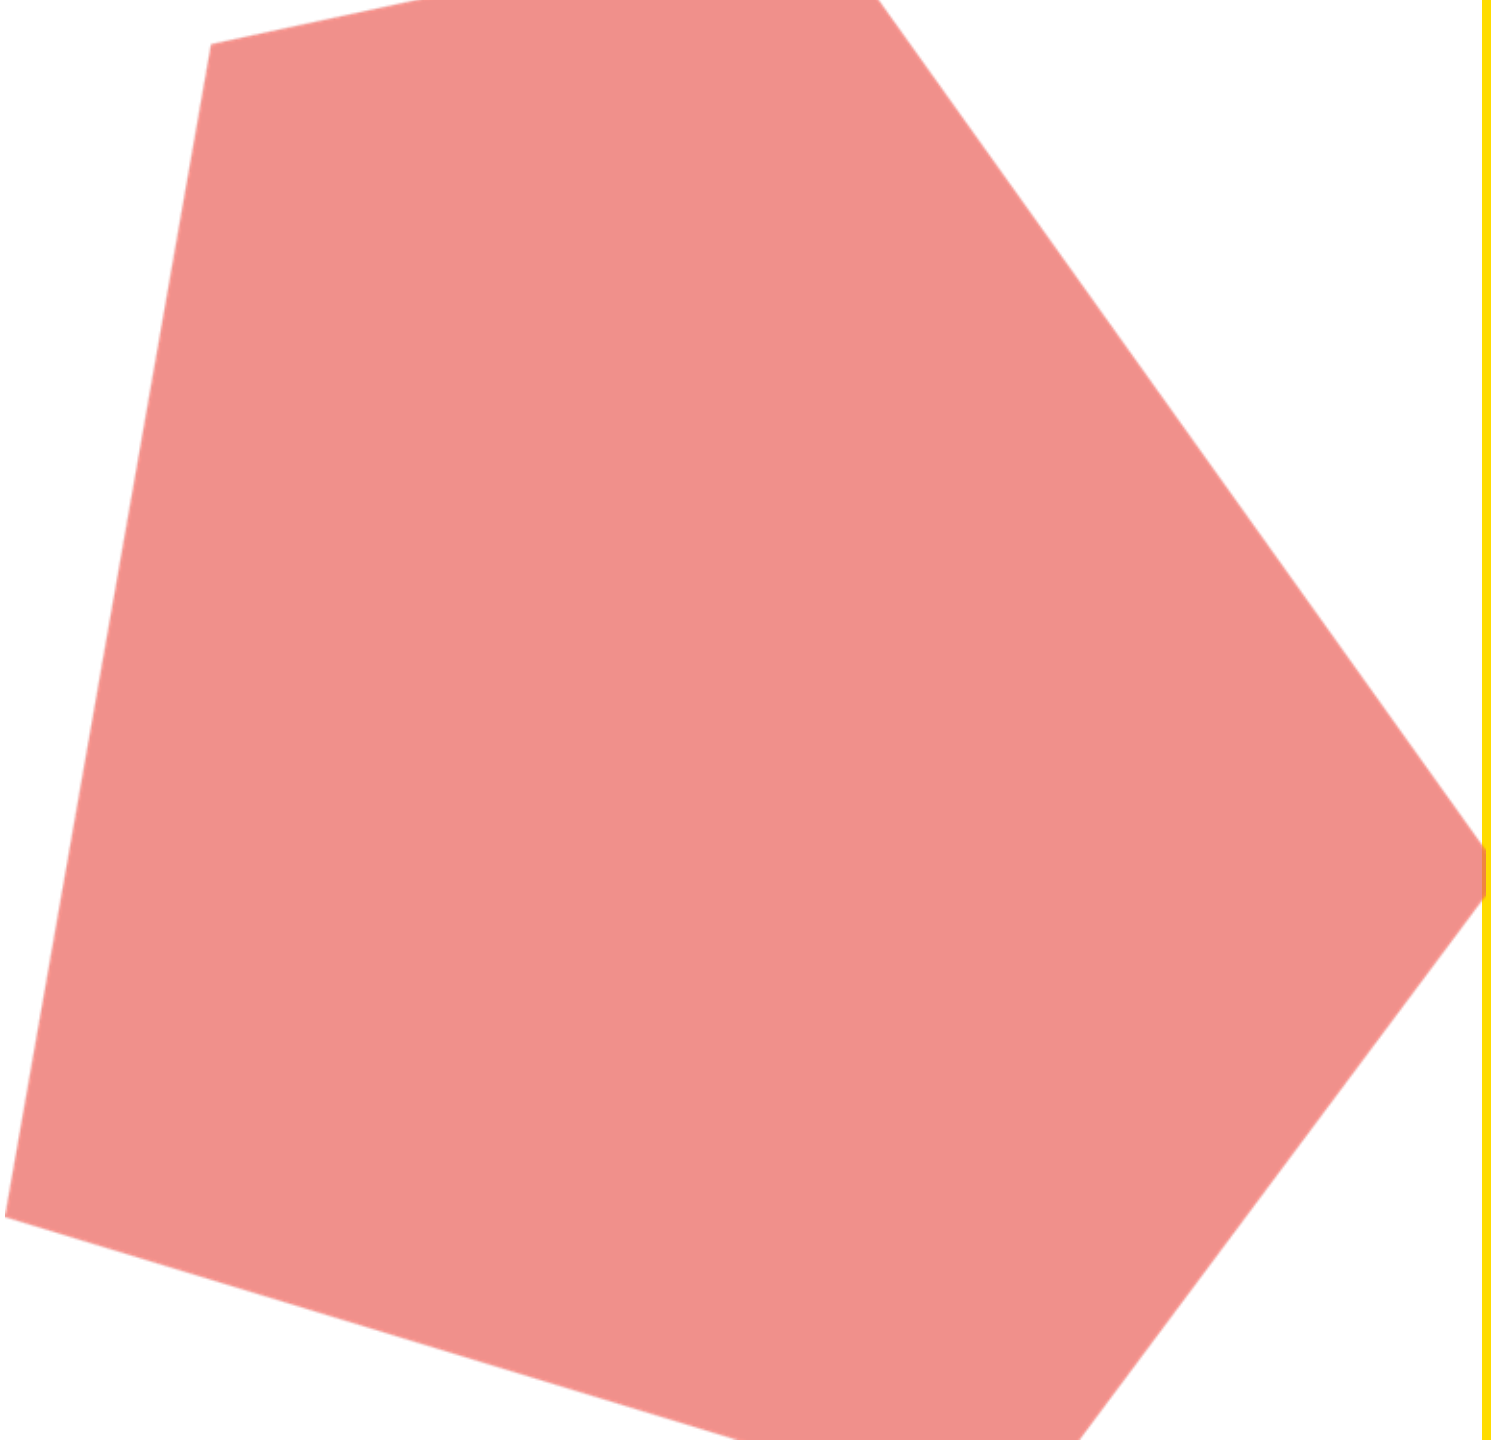
Anomaly Detection and RCA Must Be Co-Designed

- Anomaly detector determines what retrieval can find
- top@k conflates retrieval failure and reranking failure
- On CPS, deviation magnitude alone misses over a third of root causes, and retrieved ones are rarely ranked first
- A multi-signal retriever and an LLM reranker address each failure without a causal graph or labeled traces

Thank You!



Project Page



Baker Hughes

UNSW SYDNEY