engineering proceedings

MDPI

Proceeding Paper

Backward Causal Inference on Nuclear Knowledge Graphs: Domain-Specific Entity Resolution, Fault Path Discovery, and Blind Accident Prediction †

David Tom Foss

Independent Researcher, 02826 Görlitz, Germany; d.foss@ieee.org

† Presented at the 8th International Conference on Nuclear and Renewable Energy Resources (NURER 2026), Almaty, Kazakhstan, 10–12 September 2026.

Abstract

Probabilistic risk assessment (PRA) relies on manually constructed fault trees that cannot discover unknown root causes from observed failures. We present a causal knowledge graph approach for nuclear safety that addresses this limitation through backward transitive inference. The knowledge graph encodes 240 nuclear domain triplets spanning decay chains, reactor physics, material degradation, and safety systems with 11 domain-specific inference rules. A nuclear entity resolver replaces generic Jaro–Winkler string matching, achieving 97.8% accuracy on 1190 isotope notation pairs versus 19.9% for Jaro–Winkler, while reducing inference complexity from $O(N^2)$ to $O(N^{1.3})$—a 373× speedup with higher precision. Applied to the Davis–Besse reactor vessel head degradation, backward inference from 19 triplets enumerates 71 causal paths with 17 unique root causes. In a blind prediction test using only general pressurized water reactor (PWR) domain knowledge without accident-specific information, the system recovers four of five known Three Mile Island Unit 2 (TMI-2) root causes (80%). Validation against 1512 NUBASE2020 nuclides confirms 497/497 inferred decay chains with zero false positives. Convergence hub analysis identifies 33 common-cause failure candidates, and a hybrid confidence model—multiplicative for sequential chains, Möbius velocity addition for hub aggregation—is validated experimentally.

Keywords: causal knowledge graph; nuclear safety; backward inference; entity resolution; fault tree; common-cause failure; isotope notation; confidence propagation; TMI-2; Davis-Besse

1. Introduction

Nuclear safety analysis depends on probabilistic risk assessment (PRA) in which analysts manually construct fault trees tracing initiating events to consequences [1,2]. This forward-only approach cannot discover unknown root causes from observed failures. The Davis–Besse reactor vessel head degradation (2002) [3] and the Fukushima Daiichi accident (2011) [4] both involved failure chains that existing fault trees did not fully enumerate.

We encode nuclear domain knowledge as typed causal triplets—trigger, mechanism, outcome—whose inference rules expose implicit failure chains that PRA fault trees miss. The causal format [5,6] demonstrated this approach for the biomedical literature. Applying it to nuclear engineering raises three challenges: (1) entity resolution [7]—the Jaro–Winkler (JW) string similarity [8] assigns the pair U-235 ↔ U-238 a similarity of 0.92, confusing fissile and fertile uranium; (2) scalability—fuzzy matching scales as $O(N^2)$ in the number of triplets N; and (3) confidence propagation—multiplicative decay collapses to near zero at

Check for updates

Academic Editors: Sümer Şahin, Hacı Mehmet Şahin and Bekir Dursun

Published: 30 October 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

convergence hubs. Our solutions to these three problems enable backward causal inference that, given only general pressurized water reactor (PWR) knowledge, recovers 80% of the known Three Mile Island Unit 2 (TMI-2) root causes and surfaces 33 common-cause failure (CCF) candidates absent from existing fault trees.

2. Materials and Methods

2.1. Knowledge Graph Construction

The knowledge graph consists of 240 causal triplets across six domains: decay chains, fission products, material degradation, radiation/shielding, fuel cycle, and safety physics. Each triplet carries a confidence value $c \in [0,1]$ and provenance metadata. We define 11 inference rules—for example, decay transitivity (if A decays to B and B decays to C, then A indirectly produces C), breeding chains (neutron capture followed by beta decay), and degradation cascades, in which one material failure triggers a downstream failure.

2.2. Nuclear Entity Resolver

We implement an isotope parser that decomposes any notation into a canonical (element, mass number, metastable flag) tuple across five input forms, e.g., Cs-137, Cesium-137, Cs137, 55-Cs-137, and Tc-99m. Matching is exact: U-235 ↔ U-238 returns 0.0, whereas U-235 ↔ Uranium-235 returns 1.0. Because canonical tuples can be hashed, isotope lookup becomes an $O(1)$ operation instead of a pairwise $O(N^2)$ string comparison over the entity set.

2.3. Backward Inference

Given a top event, backward inference performs breadth-first search on the reversed causal graph, enumerating all paths that terminate at the target; the chain confidence of each path is computed as the product of its edge confidences (Section 2.4). Paths through inferred edges represent failure modes that are invisible to direct analysis.

2.4. Hybrid Confidence Model

For sequential chains, multiplicative propagation is appropriate because each inference step adds epistemic uncertainty:

$$c_{\text{chain}} = \gamma^{n-1} \prod_{i=1}^{n} c_i, \tag{1}$$

where $c_i \in [0,1]$ is the confidence of the i-th edge, n is the number of edges in the chain, and $\gamma = 0.85$ is a damping factor. For convergence hubs, where k independent paths support the same outcome, multiplicative aggregation collapses to $c \approx 0$—implying that more evidence yields less certainty. We therefore propose Möbius aggregation,

$$f(c_1, c_2) = \frac{c_1 + c_2}{1 + c_1 c_2}, \tag{2}$$

a well-known bounded combination rule [9] with $f(1,x) = 1$ and $f(0,x) = x$, applied iteratively over the k incoming path confidences. This formula arises naturally in hyperbolic geometry as relativistic velocity composition [10].

3. Results

3.1. Entity Resolution Accuracy

On 1190 NUBASE2020 isotope pairs, the nuclear resolver achieves 97.8% accuracy with zero false positives, versus 19.9% for JW [8] (33 false positives; 920 false negatives). Figure 1 shows representative failure cases.

https://doi.org/10.3390/engproc2026160051

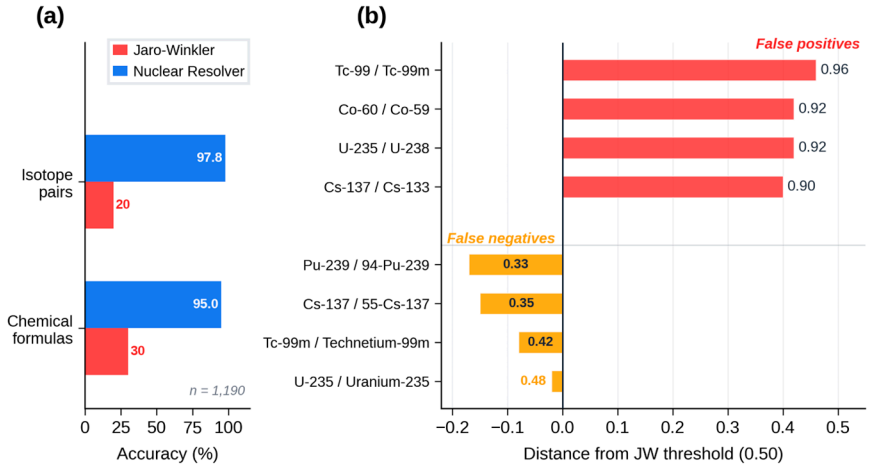


Figure 1. Entity resolution comparison. (**a**) Accuracy on isotope and chemical formula pairs ($n = 1190$). (**b**) Jaro–Winkler (JW) failure cases: red bars are false positives (different entities matched); orange bars are false negatives (same entity missed); the horizontal axis gives the distance of the JW similarity from the 0.50 decision threshold.

3.2. Inference Scalability

Replacing JW with the nuclear resolver (NR) reduces the fitted scaling exponent of the inference time from 2.0 to 1.3, i.e., from $O(N^2)$ to $O(N^{1.3})$ (Figure 2). At $N = 1584$ triplets, the inference time drops from 8.80 s to 0.024 s (373×). Projected to $N = 100{,}000$, this corresponds to 9.2 h versus 4.1 s. Precision increases from ~33% to ~100% because the removed inferences were JW false positives.

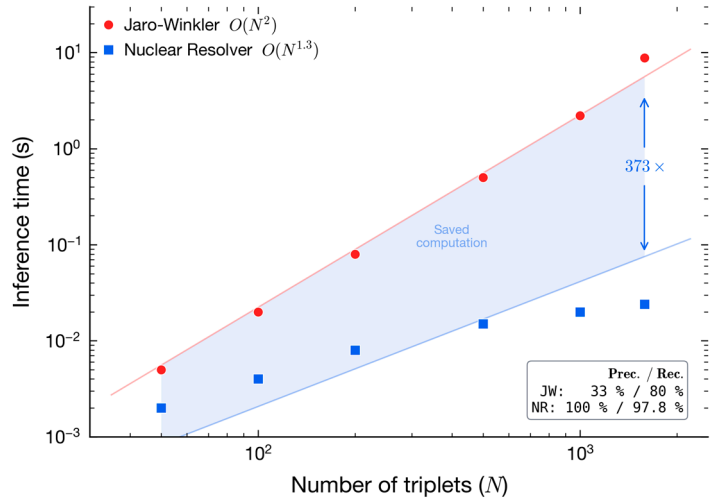


Figure 2. Inference scalability. Log–log plot of inference time versus triplet count N, showing the complexity reduction from $O(N^2)$ to $O(N^{1.3})$ and the 373× speedup at $N = 1584$. Inset: precision (Prec.) and recall (Rec.) of the Jaro–Winkler (JW) and nuclear resolver (NR) configurations confirm that the speedup comes with higher quality rather than a precision–recall trade-off.

3.3. Ground Truth Validation

The full NUBASE2020 [11] decay graph (1512 nuclides, 1601 triplets) produces 1278 inferred chains. All 497 "indirectly produced" inferences were validated against the radioactivedecay Python package [12], which implements the International Commission on Radiological Protection (ICRP) Publication 107 decay data [13]: 497/497 correct, with zero false positives. Cross-validation against the independent International Atomic En-

ergy Agency (IAEA) Live Chart of Nuclides (3386 entries) confirms 78% of the sampled parent nuclides.

3.4. Backward Inference: Accident Reconstruction

3.4.1. Davis–Besse Fault Enumeration

From 19 explicit triplets describing the Davis–Besse event [3], the inference rules derive 61 additional (inferred) triplets, and backward inference over the resulting graph enumerates 71 causal paths with 17 unique root causes—a 3.7× expansion over the 19 explicit triplets (Figure 3). All six root causes identified by the U.S. Nuclear Regulatory Commission (NRC) appear among them.

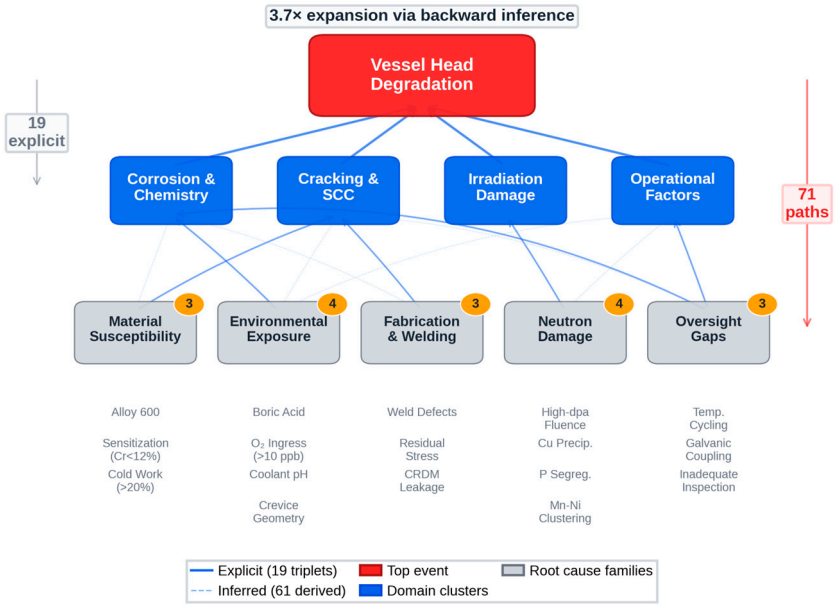


Figure 3. Davis–Besse backward fault enumeration. From 19 explicit triplets (solid edges), the inference engine derives 61 inferred triplets (dashed edges) and enumerates 71 backward causal paths from the top event to 17 unique root causes—a 3.7× expansion over the explicit knowledge base. The numbers in the circles give the number of individual root causes in each root cause family (3 + 4 + 3 + 4 + 3 = 17). SCC: stress corrosion cracking; CRDM: control rod drive mechanism.

3.4.2. TMI-2 Blind Prediction

Using the general knowledge graph (240 triplets, no TMI-2-specific information), backward inference from the top event "vessel melt-through risk" recovers four of the five known TMI-2 root causes (80%): loss-of-coolant accident (LOCA), core uncovery, fuel temperature excursion, and corium formation. Hydrogen generation is missed because the corresponding graph path terminates at "hydrogen explosion risk" without connecting to vessel failure. This is a blind prediction—the graph was designed for general PWR safety, not for TMI-2.

3.5. Common-Cause Failure Discovery

Convergence hub analysis identifies 33 nodes at which ≥3 independent causal paths converge, each a candidate CCF point [14]. Safety-critical hubs include the ductile-to-brittle transition temperature (DBTT) shift (11 incoming edges, 5 independent degradation mechanisms: Cu precipitation, P segregation, Mn–Ni clustering, hardening, and intergranular embrittlement), cladding failure risk (10 incoming edges, 4 mechanisms), and fission product release (11 incoming edges, 4 barriers). The DNA damage hub aggregates 17 paths spanning ionizing radiation, oxidative stress, and chemical mutagenesis—the same con-

vergence structure as reactor material degradation—confirming that the inference engine generalizes beyond reactor physics.

3.6. Hybrid Confidence Validation

Figure 4 compares the two models at convergence hubs. For hubs with 10–18 incoming edges, the multiplicative confidence collapses to 10^{-3}–10^{-4}, whereas the Möbius confidence remains at ~1.0. On 413 sequential NUBASE2020 chains, the multiplicative model achieves a lower mean absolute error (MAE) than the Möbius model (0.119 vs. 0.149), confirming the hybrid rule: multiplicative propagation for chains and Möbius aggregation for hubs.

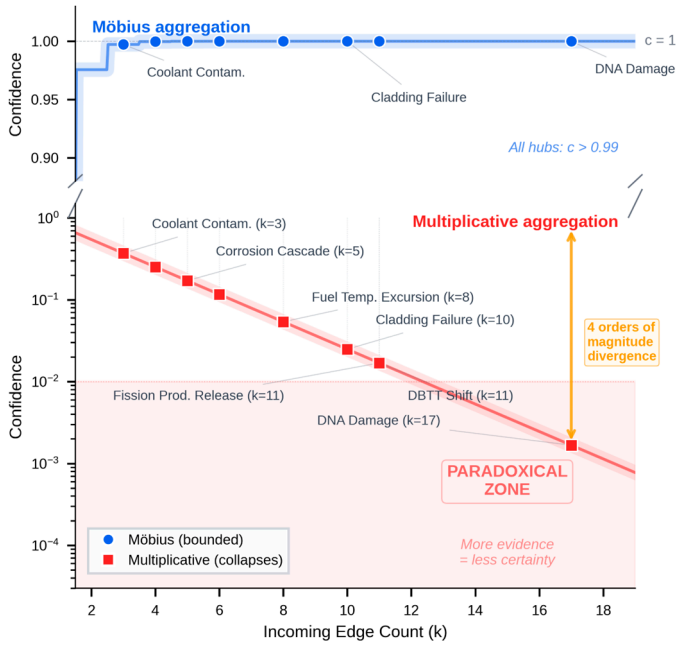


Figure 4. Confidence aggregation at convergence hubs as a function of the incoming edge count k. Multiplicative confidence collapses with increasing evidence (paradoxical zone, bottom panel, logarithmic scale), whereas Möbius aggregation preserves high confidence (top panel) when multiple independent paths support the same conclusion. Hub labels give the incoming edge count in parentheses; DBTT: ductile-to-brittle transition temperature.

4. Discussion

Backward causal inference on a nuclear knowledge graph achieves (1) 97.8% versus 19.9% entity resolution accuracy with zero false positives; (2) a 373× speedup ($O(N^2)$ → $O(N^{1.3})$); (3) 80% blind prediction of the TMI-2 root causes; (4) 100% agreement on 497 validated decay chains; and (5) 33 CCF candidates. The entity resolution failure of generic string similarity generalizes to chemical formulas [7], so the domain-specific resolver is a prerequisite for any nuclear knowledge graph rather than an optimization. The main limitations are the 240 hand-crafted triplets and the uncalibrated confidence values. The first limitation can be removed by populating the triplet base with LLM-extracted causal triplets under deterministic validation [15], which yields byte-level reproducible, validated triplets.

5. Conclusions

Typed causal triplets with domain-specific inference rules and an exact isotope resolver turn a static nuclear knowledge base into a backward-inference engine that enumerates failure paths that PRA fault trees do not contain. Combined with the hybrid confidence model—multiplicative for sequential chains and Möbius for convergence hubs—the ap-

proach provides a computationally tractable complement to fault tree analysis for root cause discovery and common-cause failure screening.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The nuclear knowledge graph, the entity resolver, and the inference engine source code are available from the corresponding author upon reasonable request. NUBASE2020 decay data were accessed through the radioactivedecay Python package [12]. IAEA Live Chart of Nuclides data are publicly available at https://nds.iaea.org/relnsd/, accessed on 19 October 2026.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Vesely, W.E.; Goldberg, F.F.; Roberts, N.H.; Haasl, D.F. *Fault Tree Handbook*; NUREG-0492; U.S. Nuclear Regulatory Commission: Washington, DC, USA, 1981.
2. U.S. Nuclear Regulatory Commission. *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants*; NUREG-1150; U.S. Nuclear Regulatory Commission: Washington, DC, USA, 1990.
3. U.S. Nuclear Regulatory Commission. *Davis-Besse Reactor Pressure Vessel Head Degradation: Overview, Lessons Learned, and NRC Actions Based on Lessons Learned*; NUREG/BR-0353, Rev. 1; U.S. Nuclear Regulatory Commission: Washington, DC, USA, 2008.
4. Institute of Nuclear Power Operations. *Lessons Learned from the Nuclear Accident at the Fukushima Daiichi Nuclear Power Station*; INPO 11-005 Addendum; Institute of Nuclear Power Operations: Atlanta, GA, USA, 2012.
5. Foss, D.T. The causal Format: Embedded Deterministic Inference for Domain-Agnostic Knowledge Graph Amplification. In Proceedings of the 2026 IEEE 27th International Conference on Information Reuse and Integration for Data Science (IRI), Seattle, WA, USA, 31 July–2 August 2026.
6. Christophides, V.; Efthymiou, V.; Palpanas, T.; Papadakis, G.; Stefanidis, K. An Overview of End-to-End Entity Resolution for Big Data. *ACM Comput. Surv.* **2020**, *53*, 127. [CrossRef]
7. Jaro, M.A. Advances in Record-Linkage Methodology as Applied to Matching the 1985 Census of Tampa, Florida. *J. Am. Stat. Assoc.* **1989**, *84*, 414–420. [CrossRef]
8. Winkler, W.E. String Comparator Metrics and Enhanced Decision Rules in the Fellegi-Sunter Model of Record Linkage. In *Proceedings of the Section on Survey Research Methods*; American Statistical Association: Alexandria, VA, USA, 1990; pp. 354–359.
9. Ungar, A.A. From Möbius to Gyrogroups. *Am. Math. Mon.* **2008**, *115*, 138–144. [CrossRef]
10. Foss, D.T. From Markov Chains to Minkowski Space: Lorentz Invariance, Quantum Measurement, and Gravitational Analogs. *SSRN* **2026**. [CrossRef]
11. Kondev, F.G.; Wang, M.; Huang, W.J.; Naimi, S.; Audi, G. The NUBASE2020 Evaluation of Nuclear Physics Properties. *Chin. Phys. C* **2021**, *45*, 030001. [CrossRef]
12. Malins, A.; Lemoine, T. radioactivedecay: A Python Package for Radioactive Decay Calculations. *J. Open Source Softw.* **2022**, *7*, 3318. [CrossRef]
13. ICRP. Nuclear Decay Data for Dosimetric Calculations. ICRP Publication 107. *Ann. ICRP* **2008**, *38*, 7–96. [CrossRef]
14. International Atomic Energy Agency. *LiveChart of Nuclides*. Available online: https://www-nds.iaea.org/livechart (accessed on 19 October 2026).
15. Mosleh, A.; Rasmuson, D.M.; Marshall, F.M. *Guidelines on Modeling Common-Cause Failures in Probabilistic Risk Assessment*; NUREG/CR-5485; U.S. Nuclear Regulatory Commission: Washington, DC, USA, 1998.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.