INL/CON-23-72314-Revision-0

Natural Language Processing-Enhanced Nuclear Industry Operating Experience Data Analysis: Aggregation and Interpretation of Multi-Report Analysis Results

Changing the World's Energy Future

Sai Zhang, Fei Xu, Zhegang Ma, Min Xian

Idaho National Laboratory

INL/CON-23-72314-Revision-0

Natural Language Processing-Enhanced Nuclear Industry Operating Experience Data Analysis: Aggregation and Interpretation of Multi-Report Analysis Results

Sai Zhang, Fei Xu, Zhegang Ma, Min Xian

May 2023

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Natural Language Processing-Enhanced Nuclear Industry Operating Experience Data Analysis: Aggregation and Interpretation of Multi-Report Analysis Results

Sai Zhang[1], Fei Xu[1], Zhegang Ma[1], Min Xian[2]**

[1]Idaho National Laboratory, Idaho Falls, ID
[2]University of Idaho, Idaho Falls, ID

ABSTRACT

Industry-wide operating experience is a critical source of raw data for reliability and risk model parameter estimations for nuclear power plants. A large portion of operating experience data are failure events stored as reports that contain unstructured data, such as narratives. In current practice, a failure report is usually reviewed and manually coded by analysts. The coding is based on extracting several event characteristics such as system name, component type, sub-part type, failure mode, and failure cause. Event narratives are mostly used to help understand events and extract their characteristics. In this line of research, we aim to maximize the usage of event narratives by leveraging natural language processing (NLP) methods to automatically convert an event narrative to a causal graph. This research has promise to improve physical understanding of failure initiation and propagation and to facilitate use of non-failure data (e.g., near-misses and degradations) to complement the limited data pool of failures. In our previous work, we developed an NLP tool and applied it to analyze a number of licensee event reports submitted by U.S. nuclear power plants to the Nuclear Regulatory Commission. In this paper, we will report our recent research progress in aggregating the results of multiple reports, developing network model(s), and drawing statistical insights.

Keywords: Nuclear power plant; event narrative; natural language processing; probabilistic risk assessment; operating experience data; causal learning

1. INTRODUCTION

1.1. Background and Research Focus

Idaho National Laboratory (INL) has been providing technical assistance to the United States Nuclear Regulatory Commission (NRC) in the areas of reliability and risk analysis since the 1980s. The parameters (i.e., initiating event frequencies, component reliabilities, and common-cause failure parameters) used in the standardized plant analysis risk models have been estimated from industrywide operating experience data. The estimations have been based upon classical statistical methods [1,2]; however, the emergence of artificial intelligence (AI), machine learning (ML), natural language processing (NLP), and other advanced tools and techniques may provide new capabilities. Report NUREG/CR-7294, prepared by some of the authors of this paper for the NRC, discusses five aspects that advanced tools and techniques could contribute to an increased understanding of safety and risk [3]. The first aspect is that they are capable of processing unstructured data, and can thus facilitate the usage of more data sources and provide a larger quantity of raw data for probabilistic risk assessment (PRA) parameter estimation. This aspect is closely related to the nuclear operating experience data, a large portion of which are failure events stored as reports that contain unstructured data, such as narratives. In current practice, a failure report is usually reviewed and manually

** Email: Sai.Zhang@inl.gov

coded by analysts. The coding is based on extracting several event characteristics such as component type, sub-part type, failure mode, and failure cause. Event narratives are mostly used to help understand events and extract characteristics.

In this study, we aim to maximize the usage of event narratives by leveraging NLP for causal learning and eventually developing a model for causal inference. This study has promise to improve physical understanding of failure initiation and propagation and to facilitate use of non-failure data (e.g., near-misses and degradations) to complement the limited data pool of failures. To achieve the goal, in this study, firstly we develop a knowledge-based tool to extract the causal information and generate the causal graph from single event narratives. Expert confirmation is introduced to correct and validate the NLP tool-generated results. The validated results will be used as annotation data for training the NLP model in the future. Moreover, this study provides the causal graph based on multiple reports for the first time, which explicitly unveils the relationship among failure initiations and propagations of multiple failure events. Results from this causal learning study will be a foundation for future causal inference study.

1.2. Definitions, Literature Review, and Research Novelty

In this study, we use "causal learning" and "causal inference" differently. We define "causal learning" as learning from natural language text, extracting causal information from it, and aggregating all the causal relations into one network. We define "causal inference" as inferring or predicting outcome(s) based on an established causal model, which can be constructed from causal learning. In short, constructing a model from causal learning and using the constructed model for causal inference. Causal learning will help our first motivation (i.e., to help improve physical understanding of failure initiation and propagation). Causal inference will contribute to our second motivation (i.e., to facilitate use of non-failure data to complement the limited pool of complete failures). In other words, by knowing the likelihood of failure precursors, the likelihood of failures can be inferred, so that we can push the level of observed data collection from failures to precursors, which will be a much larger data pool. Causal learning from natural language text is an emerging application of NLP. Yang et al., categorized techniques of causality extraction into three types: 1) knowledge-based (e.g., rule-based, pattern-based); 2) statistical ML-based (e.g., using decision tree, Naïve Bayes, logistic regression); and 3) deep learning-based approaches (e.g., convolutional neural networks) [4]. Causal learning has been explored for several domains such as extracting causal relations from medical databases [5], pipeline incident databases [6], and disaster-related news [7]. Literature moving on to causal inference have not been seen yet.

NLP has been applied to nuclear industry and a good review can be found in Smith et al. [8] and Siu et al. [9]. NLP for the nuclear industry is mostly used for classification and clustering. Causal learning has seldom been explored. Two threads of known existing works include Mandelli [10,11,12] and Zhao et al. [13]. Common aspects shared by Mandelli, Zhao et al., and our work include: 1) All are about causal learning and haven't moved on to inference and 2) all adopt knowledge-based approaches for extracting causalities. The differences in all three studies are discussed below. (1) *Dealing with different data sources.* Zhao et al. and our study look at licensee event reports (LERs) and Mandelli et al. analyze day-to-day plant records. LERs document accidents involving reactor scrams and contain thorough event narratives, which can be as long as several pages. Plant records are usually shorter and less thorough than LERs, but can be of a greater variety (like component health documentation, condition report or notes taken during plant walkdown) and contain a more diverse source of information (e.g., not only failures, but success data as well). (2) *Different scopes of result aggregation.* Zhao et al. extract causalities from each sentence and do not aggregate the results. Mandelli et al. extract and aggregate the results from a single report. Our study extracts and aggregates the results from multiple reports, which usually are of the same type, such as LERs related to motor-operated valves. (3) *Different directions for causal inference.* Zhao et al. do not mention their plan of causal inference. Mandelli et al. and our work both plan for component reliability evaluation but differ in the evaluation methods: margin-based method (i.e., how far to failure) based on condition monitoring data versus a statistical method (i.e., how likely is failure) based on operating experience data. The appli cation direction is expected to significantly impact the shape of the causal inference model. For instance, our work needs to generate statistical insights. That is why we put so much effort into establishing the capability of processing a large number of reports and aggregating the results. It can also be foreseen that, with the expansion of data inventory, the rule-based method may not accommodate our needs well and we will have to switch to ML-based approach to extract causalities.

2. METHOD

This section introduces an NLP-enhanced approach we developed [14] for causal learning from descriptive incident and accident narratives. The result from causal learning is a network graph, with nodes representing events and directed arrows showing causal relationships. The network graph contains qualitative information only. Our next-step work is to develop a quantitative, graphical model for causal inference based on the qualitative network graph. Model candidates include Bayesian belief network and graph neural network.

2.1. Causal Learning Approach Comparison: Strategy-Based vs. Machine-Based Learning

In this section, we propose an NLP-enhanced approach to extract and visualize causal relations from portable document format (PDF) reports. The LERs that domestic operating nuclear power plant licensees submit to the NRC are used in this work. A template of an LER is provided in NRC Form 366 [15]. For an LER, the proposed approach selects the following twelve fields for processing: facility name, title, event date, report date, cause, system, component, manufacturer, reportability to Industry Reporting and Information System, abstract, event description, and cause description. Note that only the information in three fields (i.e., abstract, event description, and cause description) are used for learning causal relations. The information in each of the other fields is extracted as a placeholder for potential future studies, such as report classification, report-to-report relevancy, and so on.

ML approaches are widely adopted in NLP; however, they are not yet applied to our LER analysis because no annotated datasets exist yet to train ML models. Therefore, in this study, we used a heuristic strategy-based (or rule-based) NLP approach for extracting the causal relations from LERs. The strategy-based approach can be applied to assist in the annotation of causal relations in LERs data. The strategy-based approach is introduced in more detail in Section 2.2. We are now transitioning from the current strategy-based method into an ML-based method and are establishing of a benchmark dataset as the first step. Preparing a benchmark dataset requires analyst reviews to identify causal relationships, which can be extremely expensive and time-consuming. To expedite this process, we developed a web-based tool for automating some of manual annotating steps. The tool has three major functionalities: 1) it only requires simple user interactions to efficiently annotate cause-and-effect events; 2) it enables online document management and sharing, records editing, and can export results online; and 3) it enables multiple-user collaboration to annotate same datasets.

2.2. Strategy-Based NLP-Enhanced Approach for Causal Learning from Event Narratives

This section introduces the strategy-based approach, which can learn from event narratives, identify causal relationships, and aggregate the relationships into a causal network. The approach first processes a single sentence, then multiple sentences in a single report, and eventually all the sentences in multiple reports. The aggregation involves identifying and resolving co-references (i.e., two or more expressions refer to the same person or thing). The final output from the approach is a graphical causal network. The approach includes five steps as shown in Figure 1.

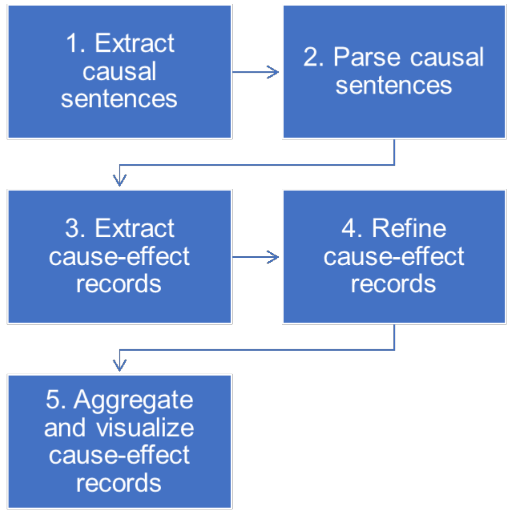

Figure 1. Steps of strategy-based NLP-enhanced method for causal learning.

Step 1: **Extract causal sentences.** First of all, the tool narrows down analysis scope by identifying and extracting the sentences containing causal relationships using a set of pre-defined keywords such as {'*cause,' 'causing,' 'result,' 'attributed to,' 'because,' 'due to,' 'in response to,' 'lead to,' 'disable,' 'render,' 'as well as'*}) compiled in [11]. An example of this step is shown in Figure 2a.

Step 2: **Parse causal sentences**. In this step, the tuple information, along with the sentence number, the keyword phrase, the index of the phrase starting in the sentence, and the index of the ending, will be extracted, shown in Figure 2b. If there are *n* keyword phrases existing in the sentence, *n* tuples will be created. Two natural language processing tools, Stanford CoreNLP and Allen NLP, are evaluated to parse the causal dependencies in sentences. Both tools obtain accurate results on simple sentences. However, Allen NLP cannot generate accurate results for complex sentences. As shown in Figure 3, Allen NLP tags 'as well as' in sentence S_2 incorrectly and generates the dependencies of 'function,' which should be tagged as punctuation; but its dependency is 'disabled' in Stanford CoreNLP's result. Therefore, we utilize the Stanford CoreNLP to perform text tagging and dependency parsing.

Step 3: **Extract cause-effect records**. The results of sentence parsing tools could reveal the sentence structure by using trees as shown in Figure 2c, but they cannot identify the cause-effect relation between events. We initialize the proposed approach using the tags of the verb keyword phrase and special causal phrases. If the causal keyword is a non-passive verb (e.g., 'lead to'), the verb's left subtree should be the parent node (cause), and the right one should be the child (effect); if it is a passive verb (e.g., 'led by'), flip the parent and child nodes. If the keyword phrase is a causal relator (e.g., 'therefore'), the dependency subject from main clause should be the parent node, the whole sentence following causal relator phrase should be the child. If the keyword phrase is an effect-noun (e.g., 'consequence'), find the first verb dependency and label the verb's left subtree as the child node and the right subtree as the parent node. For effect-relator keyword phrase, such as 'due to,' 'because,' and 'attributable to,' the left subtree of the first dependent verb is labeled as the child node, and the right subsentence after the keyword should be the parent node. For special keywords, such as 'as well as,' the left subtree of the first dependent verb is labeled as the parent node, and the right subsentence after the keyword should be the child node. Finally, a sentence with *n* tuples generates *n* causal records with three elements (i.e., cause event, keyword, and effect event). The causal relation is presented by a directed graph with two nodes, named parent node (cause) and child node (effect). Forty-five causal relations were extracted for a sample report.

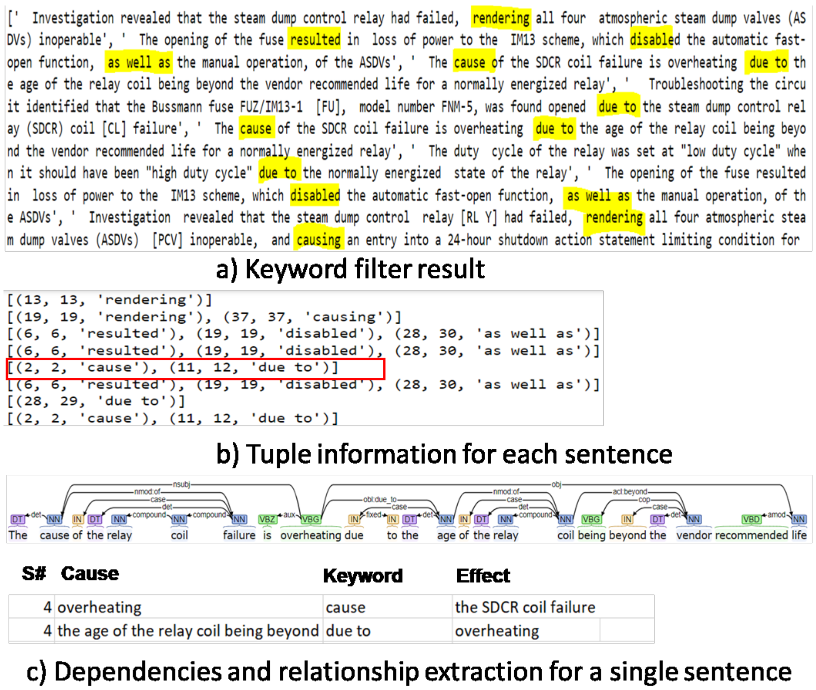

Figure 2. An example of identifying and parsing causal sentences.

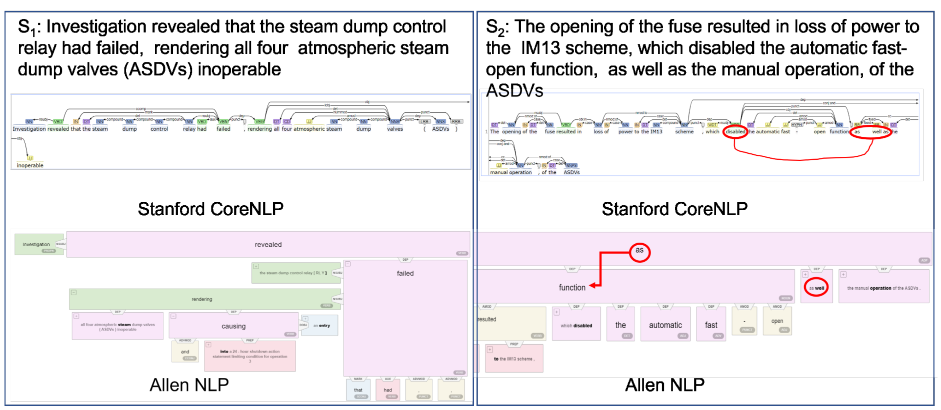

Figure 3. Comparison of Stanford CoreNLP and Allen NLP.

Step 4: **Refine cause-effect records**. This step deals with co-reference resolution. As shown in Figure 4, the records generated from Step 3 may have the same or similar causes or effects. We developed an algorithm to remove duplicates. Given two records and , remove the stop words from the descriptions, and calculate the Cosine similarity of the four event pairs , , , and which are denoted as , , , and , respectively. If and , the two records are labeled as duplicated. If , or , or , or , the corresponding events are treated as the same, and their descriptions are updated to the same. In Figure 4, the green-marked descriptions are replaced with 'the SDCR coil relay failure.' The abbreviation would be replaced by its name (e.g., 'SDCR' being steam dump control relay). Eventually, the number of refined causal records was reduced from 45 to 13 for the sample report.

s#	cause	key word	effect
0	Investigation revealed that the steam dump control relay had failed,	rendering	all four atmospheric steam dump valves (ASDVs) inoperable
1	Investigation revealed that the steam dump control relay [RL Y] had failed,	rendering	all four atmospheric steam dump valves (ASDVs) [PCV] inoperable, and
1	Investigation revealed that the steam dump control relay [RL Y] had failed,	causing	an entry into a 24-hour shutdown action statement limiting condition f
2	The opening of the fuse	resulted	in loss of power to the IM13 scheme, which
2	in loss of power to the IM13 scheme, which	disabled	the automatic fast-open function,
2	in loss of power to the IM13 scheme, which	disabled	the manual operation, of the ASDVs
3	The opening of the fuse	resulted	in loss of power to the IM13 scheme, which
3	in loss of power to the IM13 scheme, which	disabled	the automatic fast-open function,
3	in loss of power to the IM13 scheme, which	disabled	the manual operation, of the ASDVs
4	overheating	cause	the SDCR coil failure
4	the age of the relay coil being beyond the vendor recommended life for a norm	due to	overheating
5	The opening of the fuse	resulted	in loss of power to the IM13 scheme, which
5	in loss of power to the IM13 scheme, which	disabled	the automatic fast-open function,
5	in loss of power to the IM13 scheme, which	disabled	the manual operation, of the ASDVs

a) Causal relationships before coreference resolution

s#	cause	key word	effect
0	the steam dump control relay failure	rendering	all four atmospheric steam dump valves (ASDVs) inoperable
1	overheating	cause	the steam dump control relay failure
1	the age of the relay coil being beyond the vend	due to	overheating
2	The opening of the fuse	resulted	in loss of power to the IM13 scheme
2	in loss of power to the IM13 scheme	disabled	the automatic fast-open function,
2	in loss of power to the IM13 scheme	disabled	the manual operation, of the ASDVs
4	the steam dump control relay failure	causing	an entry into a 24-hour shutdown action statement limiting condition for operation 3
7	the steam dump control relay failure	due to	The opening of the fuse
9	the normally energized state of the relay	due to	the duty cycle of the relay was set at "low duty cycle" when it should have been "high duty cycle"
10	an internal failure of the electrical coil	due to	the steam dump control relay failure
12	the duty cycle of the relay was set at "low duty	because	the age of the relay coil being beyond the vendor recommended life

b) Causal relationships after coreference resolution

Figure 4. Causal relationships before and after coreference resolution.

Step 5: Aggregate and visualize cause-effect records. A visualized causal graph is produced by aggregating all the cause and effect events using the Graphviz Python package. Each cause or effect event produces one graph node, and the same events are combined to only generate one node. For each record, a directional link is produced from the node of the cause event and the node of the effect event.

3. RESULTS

This section introduces the causal graph generated from aggregating the causalities extracted from twenty LERs related to motor-operated valve failures shown in Table I. LERs are publicly available and can be searched and accessed at https://lersearch.inl.gov/LERSearchCriteria.aspx.

Table I. A list of the analyzed licensee event reports related to motor-operated valve failures.

LER Number				
3412022004	4402022001	3212021002	3412016002	3412017003
4832020005	3162020003	4162020004	4142015001	3732017006
2812019001	2472018003	4612018003	4832015001	3732017005
2512018001	3522018002	3732018004	2962014003	2592015002

3.2. Graph Structures

Using the NLP tool, an event report can be converted to a causal graph visualizing the event initiation and propagation. As can be seen from Figure 5, a causal graph in this study consists of two graphical elements—nodes and directed associations. A node represents an event describing what happens. A directed association (i.e., a directed arrow) represents the causal relationship between child node(s) and parent node(s). The direction of an arrow from a parent node to a child node means that the occurrence of the parent node causes the occurrence of the child node.

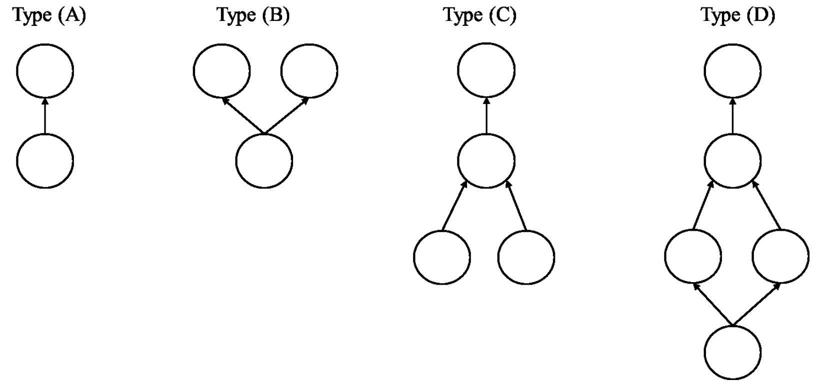

Figure 5. Observed graph structures from the analyzed reports.

As shown in Figure 5, four types of causal graph structures are observed from the twenty analyzed event reports. The structures are simplified versions for a categorization purpose and the real structures may include a larger number of nodes. Type (A) represents a chain-pattern, one-to-one correspondence type of causal relationship in which one node leads to another node only. Type (B) represents a fork-pattern type of causal relationship in which one node leads to more than one node. Type (C) represents a collider-plus-chain relationship in which multiple nodes contribute to the occurrence of the same node. Type (D) has a more complicated structure (i.e., collider plus fork plus chain) and can be understood as a combination of type (B) and type (C) structures. There can be cases in which multiple child nodes lead to the same parent node, such as those in the type (C) graph structure in Figure 5. It is to be noted that, unlike the fault tree analysis, the causal graph developed for this study does not show the relationships (i.e., OR [the occurrence of any child node can lead to the occurrence of the parent node] or AND [the occurrence of the parent node requires the occurrences of all the child nodes]) between the child nodes leading to the same parent node. This capability may be added if necessary as this research is further developed.

Finally, of the twenty analyzed reports, Type (A) is the most common structure with thirteen reports falling in this type; four reports belong to (C), two to (B), and one to (D). With future efforts in analyzing more event reports, it can be expected that more graph structure types will be observed. It is also expected that all the potential structures, regardless of their complication levels, will be combinations of the three basic patterns (i.e., chain, fork, and collider).

3.3. Aggregated Results from Multiple Reports

The causal records learned from the 20 reports were aggregated into one graph shown in Figure 5 with the nodes illustrated in Table II. In Figure 6, top nodes represent failure modes (e.g., valve failed to open) and bottom nodes represent root cause (e.g., design deficiency). These two types of information are already included in the current NRC database; however, a big portion of nodes in between, representing different levels of failure causes (or "failure precursors"), haven't been fully utilized. This study, using NLP techniques, made it possible to identify these in-between nodes and incorporate them as part of model. The aggregated graph also reveals a lot more potential failure paths besides the observed failure paths.

The aggregated graph is very informative; however, as of now it is still at a "raw" stage since it provides qualitative information rather than quantitative insights. To pack up this "raw" graph as a model for inference use, we are now working on the following two topics: 1) how to address logical relationships (i.e., the AND and/or OR relationships among causal relations); 2) how to estimate quantitative metrics based on the graph (e.g., the conditional probability of a precursor evolves into a failure).

Table II. Causal events extracted from 20 reports.

No.	Event	No.	Event	No.	Event
1	Valve wrongly opened	22	Torque switch failed to close	43	Inappropriate manipulation
2	Valve failed to open	23	Shaft disengaged	44	Loose fuse connection
3	Valve wrongly closed	24	Fatigue	45	Vibrations
4	Valve status unstable	25	Foreign material in a contactor	46	Wrong bolt installed
5	Valve failed to close	26	Weld failed	47	Foreign materials left
6	Valve leakage	27	Flow switch failed	48	Valve unable to be closed at normal force
7	Load driver wrongly activated	28	Relay wires lifted and incorrectly landed	49	Valve coasting further opened
8	Positioner failed	29	One phase of power inhibited, and other phases experienced high amps	50	Steam leak from nearby
9	Valve stem rotation	30	Torque switch corrosion	51	Lack of filtration device
10	Relay failed	31	Roll/shear pin broken	52	Inadequate preventive maintenance frequency
11	Power supply interrupted	32	Overstress	53	Contacts between parts
12	Motor operator control failed	33	Suppressors failed	54	Binding
13	Valve grounding	34	Pilot valve stuck	55	Subpart displacement
14	Debris from control air	35	Wear	56	Random failure
15	Valve actuator failed	36	Contactor susceptible to foreign material	57	Manufacturing deficiency
16	Vent line crack	37	Inadequate inspection procedure	58	Design deficiency
17	Circuits energized	38	Incomplete preventive maintenance procedure	59	Human error
18	Contactor stuck	39	Inadequate performance monitoring	60	Degradation
19	Torque switch arm disengaged from stem key	40	Ineffective troubleshooting	61	Environmental impact
20	Anti-rotation key dropped out	41	Loose sliding link	62	Organizational deficiency
21	Thermal overload device tripped	42	Internal water leakage	63	Installation error

4. CONCLUSIONS

This paper introduces an NLP-enhanced method to learn causal relations from textual event narratives and aggregate the causal records from multiple reports into an integrated causal graph. This paper also present s the results of analyzing 20 motor-operated valve-related failure event reports. At this stage, the aggregat ed graph is still "raw" and provides qualitative information only. This paper discusses the path forward to develop a quantitative model for causal inference.

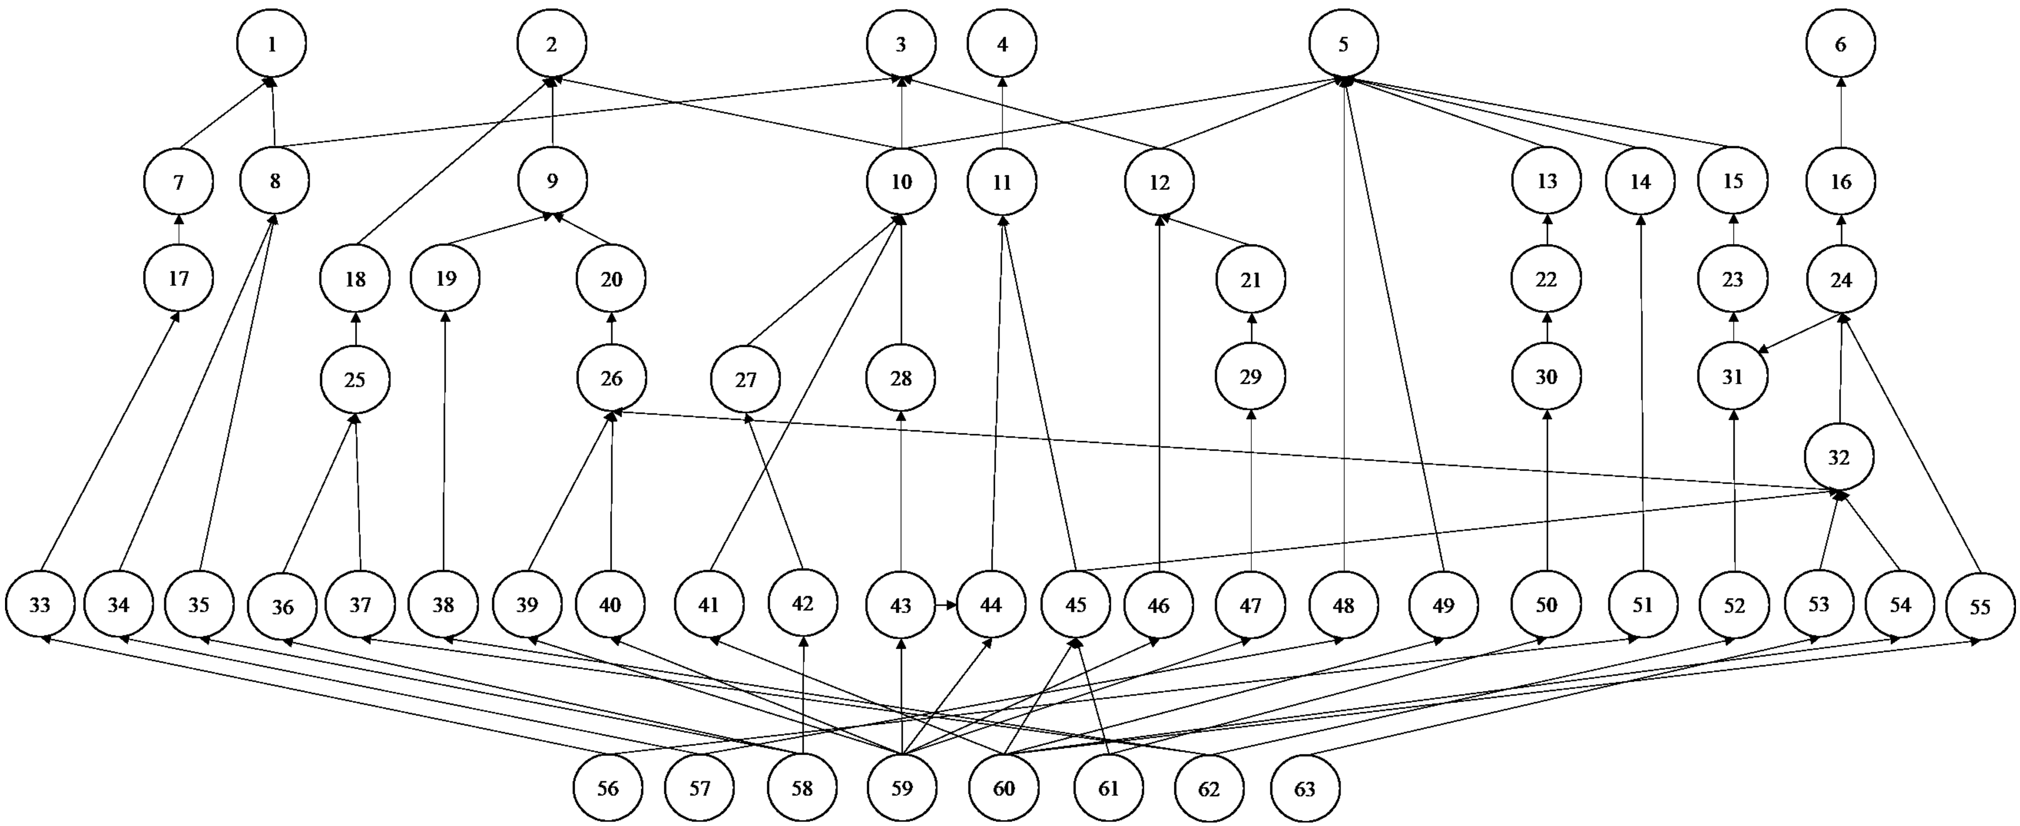

Figure 6. Graph developed by aggregating the causal relations extracted from 20 reports.

ACKNOWLEDGMENTS

The authors would like to thank Michael Calley, Han Bao, and Katie Stokes for their reviews and edits. This paper was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, do not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

REFERENCES

1. C. L. Atwood et al., "Handbook of Parameter Estimation for Probabilistic Risk Assessment, NUREG/CR-6823," United States Nuclear Regulatory Commission (2003).
2. Eide et al., "Industry-Average Performance for Components and Initiating Events at U.S. Commercial Nuclear Power Plants, NUREG/CR-6928," United States Nuclear Regulatory Commission (2007).
3. Z. Ma, H. Bao, S. Zhang, M. Xian, and A. Mack, "Exploring Advanced Computational Tools and Techniques with Artificial Intelligence and Machine Learning in Operating Nuclear Plants, NUREG/CR-7294, INL/EXT-21-61117," United States Nuclear Regulatory Commission (2022).
4. H. Yang, S.C. Han, and J. Poon, "A survey on extraction of causal relations from natural language text," *Knowledge and Information Systems,* **64** (5), pp. 1161–1186 (2022).
5. C.S.G. Khoo, S. Chan, and Y. Niu, "Extracting causal knowledge from a medical database using graphical patterns," *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, Hong Kong, October, pp. 336–343 (2000).
6. G. Liu, M. Boyd, M. Yu, S.Z. Halim, and N. Quddus, "Identifying causality and contributory factors of pipeline incidents by employing natural language processing and text mining techniques," *Process Safety and Environmental Protection*, **152**, pp. 37–46 (2021).
7. F. Sakahira and U. Hiroi, "Designing cascading disaster networks by means of natural language processing," *International Journal of Disaster Risk Reduction*, **66** (2021).
8. C. Smith, A. Al Rashdan, and V. Agarwal, "Using the New Math: Artificial Intelligence and Machine Learning Applications for the Nuclear Power Industry," *NuclearNews*, **65** (7) (2022).
9. Siu, N., et al. ,"Advanced knowledge engineering tools to support risk-informed decision making: final report (public version)," U.S. Nuclear Regulatory Commission (2016).
10. D. Mandelli and C. Wang, "A Model Based Approach to Extract Health Information from Textual Data," *Annual Conference of the PHM Society*, Vol. **14** (1) (2022).
11. D. Mandelli et al., "Bridging Equipment Reliability Data and Robust Decisions in a Plant Operation Context," INL/RPT-22-67670, Idaho National Laboratory (2022).
12. D. Mandelli and C. Wang, "A Causal Approach to Integrate Component Health Data into System Reliability Models," *Annual Conference of the PHM Society*, Vol. **14** (1) (2022).
13. Y. Zhao et al., "Automated Identification of Causal Relationships in Nuclear Power Plant Event Reports." *Nuclear Technology.* **205** (8), pp. 1021–1034 (2019).
14. S. Zhang et al., "Natural Language Processing-Enhanced Common Cause Failure Data Analysis, " *Proceedings of the 16th Probabilistic Safety Assessment and Management Conference* (2022).
15. United States Nuclear Regulatory Commission. "Form 366, Licensee Event Report Template" (2023).