INL/CON-24-77805-Revision-0

Model-Based Approaches to Generate Knowledge from Data in a Plant Reliability Context

August 2024

Changing the World's Energy Future

Diego Mandelli, Congjian Wang, Vivek Agarwal, Joshua J Cogliati, Chaitee Milind Godbole

Idaho National Laboratory

INL/CON-24-77805-Revision-0

Model-Based Approaches to Generate Knowledge from Data in a Plant Reliability Context

Diego Mandelli, Congjian Wang, Vivek Agarwal, Joshua J Cogliati, Chaitee Milind Godbole

August 2024

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

Model-Based Approaches to Generating Knowledge from Data in a Plant Reliability Context

D. Mandelli[*], C. Wang, V. Agarwal, C. Godbole, and J. Cogliati

Idaho National Laboratory, Idaho Falls (ID), USA
[leave space for DOI, which will be inserted by ANS]

ABSTRACT

One challenge that nuclear power plant system engineers are facing is continuous generation of an extremely large amount of equipment reliability (ER) data. These data elements come in textual (e.g., condition reports) and numeric (e.g., generated by monitoring systems) forms. They provide system engineers with valuable insights and information by discovering anomalous behaviors or degradation trends, identifying possible causes behind such behaviors and trends, and predicting their direct consequences. This paper directly targets the knowledge generation from ER data by putting "data into context." We employ model-based system engineering (MBSE) of systems and assets to represent and capture their architecture and functional (i.e., cause-effect) relations. ER data elements are processed by first identifying which of the developed MBSE elements they are referring to. This task is harder for textual data since the information contained in issue or maintenance reports needs to be "understood" by a computational tool. We called this process "knowledge extraction" since our methods extract knowledge from textual data. Last, once numeric and textual ER data elements have been processed and "understood," we discover possible cause-effect relations among them. This is performed by observing whether a logical connection through the MBSE models exists, and if there is a temporal relationship among them. The logic and temporal are the two main ingredients to perform "machine reasoning" from ER data.

Keywords: reliability, nlp, mbse

1. INTRODUCTION

For decades, existing nuclear power plants (NPPs) have been transitioning from corrective and periodic maintenance to predictive maintenance strategies to reduce operation and maintenance costs. While corrective maintenance is performed only when the asset fails (with high costs due to asset replacement and unexpected system and plant unavailability, e.g., loss of generation), periodic maintenance is performed at specific time intervals based on reliability factors and past operational experience (with high costs due to continuous maintenance operations that may not be warranted). On the other hand, predictive (i.e., performance-based) maintenance operations are only performed when the asset under consideration requires it. The approach requires advanced prognostic and health management (PHM) techniques [1] [2] and this can be achieved by constantly monitoring asset status and performances and processing such data (through anomaly detection, diagnostic, and prognostic computational algorithms) to identify asset degradation trends and faulty states [3].

The transition from periodic or corrective maintenance to predictive maintenance is designed so that maintenance occurs only when the asset requires it (i.e., before its imminent failure). This guarantees that asset availability is maximized, and operation and maintenance costs are minimized [4]. These benefits can

[*] Corresponding Author: Diego.Mandelli@inl.gov.

be achieved by employing monitoring sensors, automated data acquisition systems, data analysis methods, and improved decision processes. When combined, they can provide precise information about the health of an asset, track its degradation trends, and provide information about its expected failure time. With such information, maintenance operations can be scheduled and performed for each asset only when needed.

Activities around NPP PHM produce a large amount of equipment reliability (ER) data about the status of component, assets, and systems. Such data can come in several forms, such as online monitoring data (e.g., pump vibration data, pump mass flowrate), surveillance and testing data (performed by plant operators at regular intervals), condition reports (which typically contain anomalous conditions), or maintenance reports (which indicate operations performed to restore component or asset health. These data elements precisely record assets and systems performance and health throughout their life cycle. In addition, such data have the potential to provide insights to system engineers about the presence of anomalous behaviors or degradation trends, the possible causes of such behaviors and trends, and identify in advance their direct consequences. However, several challenges have proven to be roadblocks. While some are technical in nature (i.e., data are often distributed over several physical servers and databases), others are conceptual since data elements have different formats (numeric versus textual) and measured values have different scales (vibration spectra versus oil temperature). This article directly focuses on the integration of numeric and textual data elements to assist plant system engineers in analyzing ER data.

The task starts with extracting knowledge from textual data using natural language processing (NLP) methods [5] and quantifying system, asset, and component health from numeric data. We then employ model-based system engineering (MBSE) models [6] of systems and assets to identify their architecture and functional, or cause-and-effect, relations. Based on their nature, the ER data elements are associated with a single MBSE entity. This bonding of MBSE models and ER data elements constitutes the first-of-its-kind knowledge graph of an NPP system. At this point, data elements are organized in a structured way so that system engineers can identify cause-effect patterns between data elements and act accordingly.

2. DIGITAL REPRESENTATION OF SYSTEMS AND ASSETS

The ability of NPP system engineers to analyze ER data relies on their knowledge about system architecture and the physical and logical interdependencies between the assets that are part of such a system. Current ER data analysis tools rely only on available data, and they are blind on the actual operating context that have generated such data. The term context here refers to the actual physical element being monitored and observed, the function(s) supported by such a physical element, and the other elements directly linked to it.

To address this limitation, we have developed a set of methods that are based not solely on data but also models. The objective of these models is to emulate system engineer knowledge and capture system architecture and the physical and logical interdependencies between the assets that are part of such a system. Here, we are employing state-of-the-art MBSE methods [6], which provide several solutions to represent systems, assets, and components from both form (which elements are part of the structures, systems, and components) and functional (how systems and assets interact with each other, and which functions they support) points of view. These solutions are based on MBSE languages that represent system and asset form and function elements through a set of diagrams. The most common languages are object process methodology (OPM) [7] , unified modeling language [8], and systems modeling language [9]. For the scope of this project, we have chosen the OPM language because it provides the basic modeling elements we are looking for, and more importantly, it is possible to automatically generate digital data structures (i.e., graphs) from OPM diagrams. Each element of an OPM diagram indicates either a function or form element. Links between OPM elements have precise meaning [7].

Figure 1 shows the OPM diagram of a centrifugal pump. In our work, each OPM diagram is translated into a graph structure, a database for capturing the different types of OPM elements and edge. The resulting graph made it possible to query all the textual IDs that describe each form or function element of the diagram. These textual IDs can be used in our NLP methods to identify and recognize each element based on a given set of sentences (see Section 3). More importantly, if multiple textual IDs are identified, the directed graph can identify the links between them. Note that such links are implicitly causal in nature. In summary, the obtained graph structure is used to accomplish three main objectives: identify OPM entities mentioned in ER textual data elements, identify OPM entities that refer to ER numeric data elements, and identify logical connections between the ER data elements by determining whether a direct causal path exists between the OPM entities associated with each ER data element.

Note that the OPM diagrams can be used for more than simply putting "text into context" via the NLP methods described below in Section 3. These diagrams can, in fact, be employed to identify which OPM elements are being monitored by condition-based, diagnostic, and prognostic systems (e.g., pump shaft vibrations or the rate of mass flow exiting the pump itself) (see Section 4).

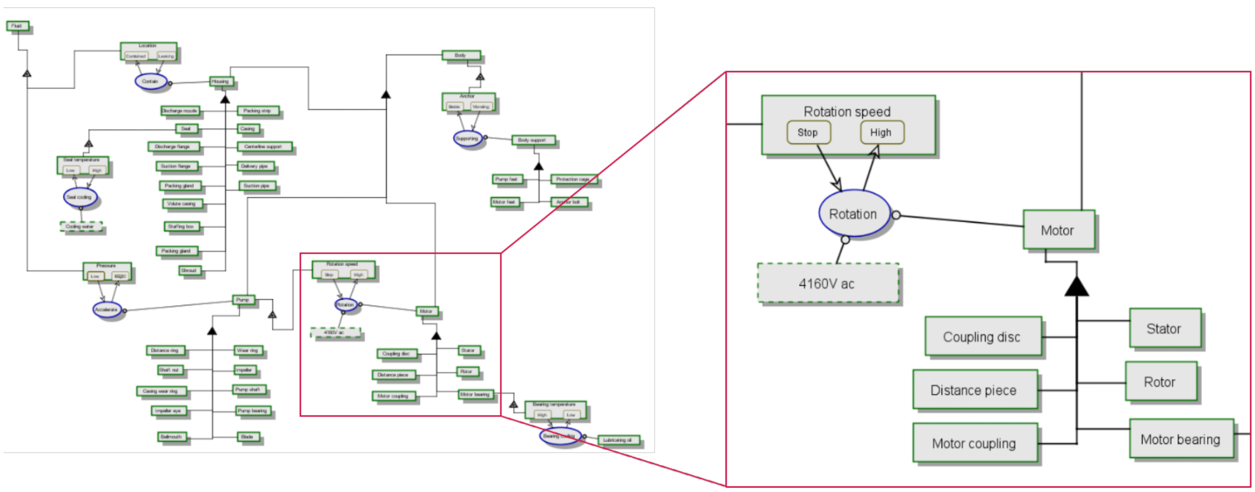

Figure 1. OPM representation of a centrifugal pump.

3. ANALYSIS OF TEXTUAL DATA

ER textual data such as issue reports (IRs) and work orders (WOs) are valuable data sources for tracking asset health histories, identifying health trends, and performing root-cause analyses. These data sources, typically obtained in text form, are usually available in digital repositories. Methods have been developed over the past two decades to enable machine learning (ML) models [10] to analyze textual data and classify textual elements based on their nature (e.g., safety-related versus non-safety-related). In the context of the present work, we are not interested in solving classification problem, but rather in extracting actual knowledge from textual data. This is a harder task, as it requires the development of context-dependent models and vocabularies. The medical field is leading the way in this area by developing methods to extract knowledge from textual data (e.g., for diagnostic purposes or to estimate the performance of specific treatments). When applied to the nuclear field, knowledge extraction consists of several tasks, including identifying:

- Plant-specific entities such as systems, assets, and components (e.g., centrifugal pump, accumulator system, and pump shaft)
- Temporal attributes that characterize events such as the occurrence, duration, and order of events
- Measured quantities with numeric values followed by units of measure
- Phenomena like material degradation or asset functional failure
- Causal relations between events.

Knowledge extraction is enabled by a series of data, models, and methods. The developed series of methods was designed to identify the elements listed above using a mixture of rule-based and ML algorithms. These methods rely heavily on data dictionaries and plant, system, or asset models. Data dictionaries containing a large number of keywords related to the nuclear field were partitioned into several classes (e.g., materials, chemical elements and compounds, degradation phenomena, and electrical/hydraulic/mechanical components).

The ability of system engineers to analyze textual data is enabled by their knowledge of the system's architectural scheme comprised of the components and assets that comprise the system. In simpler terms, they know what physical elements comprise a given asset or system, along with their functional relations and dependencies. Without such information, knowledge extraction from textual data is very difficult. For the present study, our methods were designed to check whether OPM entities (see Section 2, above) are mentioned in ER textual data elements.

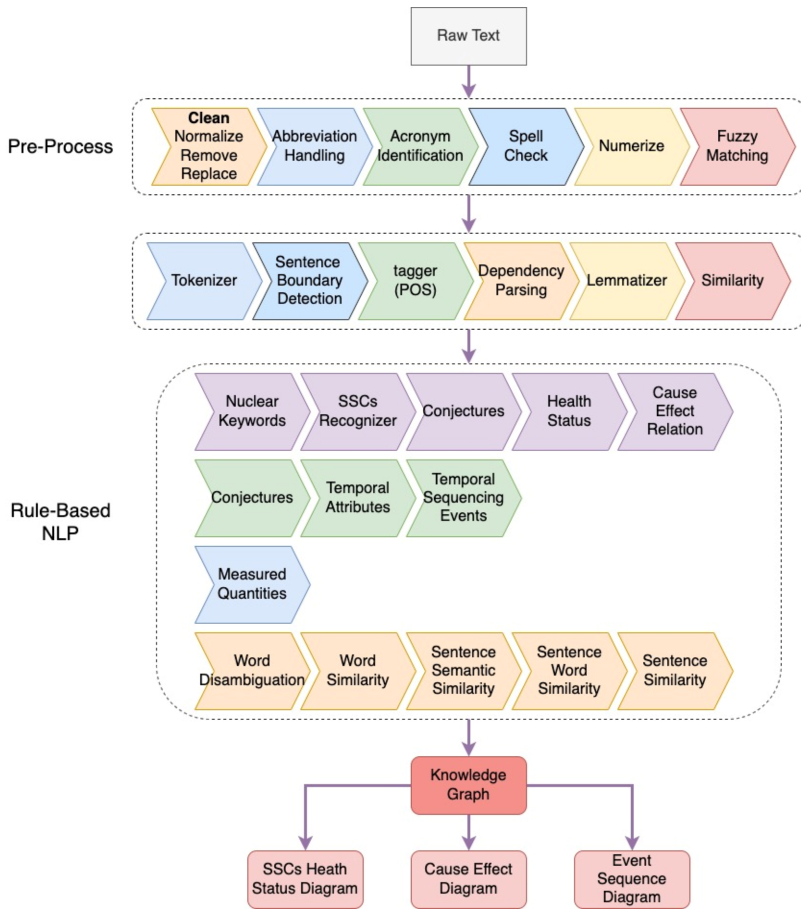

Figure 2. Graphical representation of the methods employed to analyze ER textual data elements.

In this work, we integrated SpaCy (https://github.com/explosion/spaCy), PySBD (https://github.com/nipunsadvilkar/pySBD), Coreferee (https://github.com/msg-systems/coreferee) and a few other libraries for text data analyses. Using the NLP capabilities provided by these libraries (dependency parsing, part-of-speech [POS] tagging, tokenizing), our methods focus on extracting health status, cause-effect information, and temporal relationships from ER data to support robust decision-making within a plant operation context. The presented package consists of many functionalities/pipelines covering the actions needed to perform ER data analyses (see Figure 2 and Table 1).

Table 1. Capabilities leveraged from external libraries.

ID	NLP Step	NLP Pipeline	Note
1	Tokenization	tokenizer (SpaCy)	Segmenting text into words, punctuation marks, etc.

ID	NLP Step	NLP Pipeline	Note
2	Sentence segmentation	pysbdSentenceBoundaries (PySBD)	Finding and segmenting individual sentences
3	POS	tagger (SpaCy)	Assigning word types (e.g., verbs, nouns) to tokens
4	Dependency parsing	parser (SpaCy)	Assigning syntactic dependency labels and describing the relationships between individual tokens (e.g., subject or object)
5	Lemmatization	lemmatizer (SpaCy)	Assigning the base forms of words (e.g., the lemma of "was" is "be" and that of "pumps" is "pump")
6	Similarity	tok2vec (SpaCy)	Comparing words, text spans, and documents in terms of their similarity
7	Rule-based entity recognition	entity_ruler (SpaCy)	Finding sequences of tokens based on their texts and linguistic annotations, and labeling named structures, systems, and components (SSCs)
8	Coreference	Coreferee (Coreferee)	Resolving coreference situations in which two or more words within a text refer to the same entity

Figure 3 provides an example of knowledge extraction from an ER textual data element. Based on the developed libraries, the asset (i.e., pump) and reactions (i.e., cracking and failure) mentioned in the text are identified, along with a specific pump OPM entity (i.e., shaft). Furthermore, additional elements are captured: the existence of a conjecture and the temporal attribute associated with pump failure.

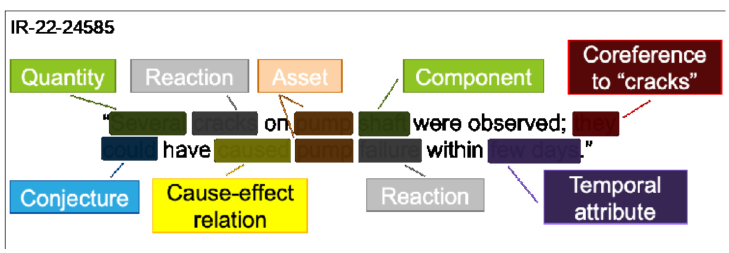

Figure 3. Example of NLP knowledge extraction from an ER textual data element.

4. ANALYSIS OF NUMERIC DATA

ER data generated in numeric format are very common in existing NPPs, in which many assets are continuously monitored (e.g., vibration data, oil temperature, and outlet water pressure) via advanced monitoring and PHM systems to identify data trends that may inform system engineers of degraded performance or failure of the considered asset. One challenge with complex systems (not only nuclear) is quantifying the health of each asset, then propagating the assets' health values to the system level. Both [11] and [12] provide a margin-based approach that address these challenges. Such approaches enable the data generated by condition assessment, diagnostic, and prognostic systems to be converted into margin values that serve as a quantitative measure of asset health. For the scope of the present article, we are interested in tracking the health history of an asset (if ER data are available) and exploring how reported events (e.g., IRs) might causally relate to health trends. The advantage of margin-based approaches is that complex monitoring data generated by condition assessment, diagnostic, and prognostic systems can be integrated into a common quantitative health measure. Rather than focusing on the likelihood of a given event (in probabilistic terms), we think in terms of how far off the event is from occurring. This new interpretation of reliability shifts the focus away from probability of occurrence and toward assessments of how near assets are to failure, or at least reaching an unacceptable level of performance (see Figure 3). Note that two data elements are required for this assessment: the estimated actual health condition of the asset, which can be acquired by the asset-monitoring system or through diagnostic methods, and the limiting conditions that must be avoided, which can be acquired from past operational experience (e.g., monitoring data generated by similar assets under failure conditions).

An asset's margin value M is defined over the $[0,1]$ interval, where $M = 1$ corresponds to a perfectly healthy asset (requiring minimal to no maintenance attention) and $M = 0$ corresponds to a faulty asset (requiring maintenance attention). Note that margin quantification is impacted by the availability of monitoring data and can be defined over heterogenous variables such as pressure, vibration spectra, and time. For example, when dealing with condition-based monitoring data (both current and archived), margin M is here defined as the distance between actual and past conditions (e.g., oil temperature and vibration spectrum) that lead to failure. Hence, margin-based reliability modeling provides a unified approach to dealing with heterogeneous monitoring data elements.

Here, each numeric data element (i.e., margin temporal profile of a specific monitored variable) can be matched to an exact OPM entity. The example taken from [12] and reflected in Figure 5 shows the margin temporal profile for an asset of an existing NPP. A margin value can be calculated using a reduced-order model trained on both healthy and faulty data (i.e., a classifier).

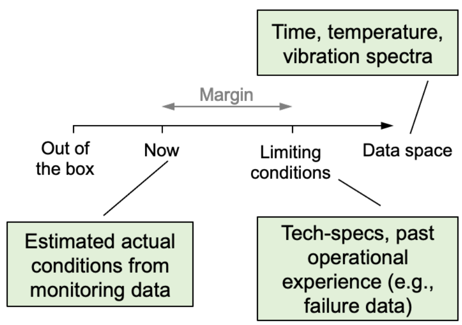

Figure 4. Graphical representation of margin (Mandelli, 2023).

5. CAUSAL REASONING FROM ER DATA

Sections 3 and 4 presented methods of analyzing numeric and textual ER data elements, and we explained how OPM diagrams (Section 2) can be employed to identify possible causal relationships between ER data elements. The word "possible" is intended to indicate that two events sharing an OPM-based direct relation may in fact exist independently from each other. The first step in testing such dependence is to observe their temporal correlation. Our work extends that presented in [13], in which the temporal correlation between time series and events is formulated in terms of a two-sample problem [14]. Our extension includes three relevant items: a modification to the testing process structure, a different two-sample testing algorithm, and the handling of events defined over an interval (as opposed to a time instant).

In its original formulation in [13], the temporal correlation was measured between a set of identical events and the time series. In the scope of the present work, we often deal with single events (e.g., abnormal behavior of an asset) rather than sets of events. The algorithm presented in [13] was based on measuring the statistical difference between the portions of the time series pertaining to both before and after the occurrence of an event (as defined over a temporal instant). Our extension, which enables dealing with events defined over a temporal interval requires the additional time series portion that corresponds to the duration of the event itself. We employed the Maximum mean discrepancy (MMD) [14] to perform such testing.

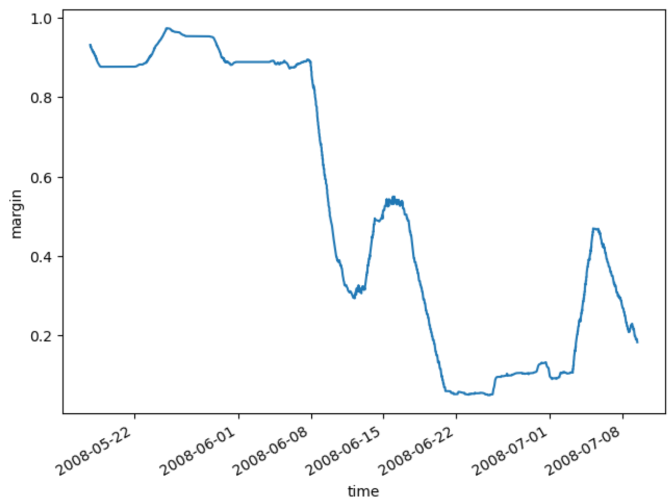

Figure 5. Example of margin temporal profile, as determined with a trained classifier [12](Mandelli, 2024).

When dealing with two events detected only by textual ER data elements (or only by numeric ER data elements), causal analysis can be performed simply by directly employing MBSE models to check the logical and temporal relations. On the other hand, when dealing with the mixed case, numeric and textual data elements, slightly different thinking applies from a temporal standpoint. More precisely, we need to understand if the occurrence of an event (e.g., reported as an IR) has triggered a change in the numeric counterpart or vice versa. After NLP knowledge extraction, it is possible to temporally characterize it (i.e., define event time of occurrence and its duration, if available). The following step is to assess the behavior of the time series prior to, after, and during the occurrence of such an event; in our work, such an assessment is performed through a classical two-sample testing algorithm.

An example is shown in Figure 6 where two events (indicated as E_1 and E_2) are analyzed along with a time series TS (shown in blue). Events E_1 and E_2 (provided in textual form) are analyzed to capture the nature of the event along with their temporal attributes (i.e., time of occurrence and duration, if available). Then, the provided numeric time series is analyzed to identify if there is a temporal correlation with each event. For E_1, the developed two-sample testing flags the existence of such a correlation; for example, *after* the occurrence of E_1 the time series behavior changes (which is indicated as $E_1 \rightarrow TS$). Similarly, we obtain a temporal correlation between E_2 and the time series, which is confirmed by the behavior of the time series *while* event E_2 is occurring (which is indicated as $E_2 \leftrightarrow TS$).

Provided the set of processed ER data elements—either textual (see Section 3) or numeric (see Section 4)—the goal becomes to organize each element into a graph structure that captures the cause-effect relations (logical and temporal). Our approach began with the OPM graph structure of the system/asset under consideration (see Section 2), then progressed according to the following steps:

1. Associate an ER textual data element with an OPM entity.
2. Identify ER numeric data elements that have a logical path to the ER textual data element identified in Step 1.
3. Determine whether there is a temporal relation between the ER textual data element identified in Step 1 and the ER numeric data elements identified in Step 2 (see Section 6).
4. If both the temporal and logical relation have been identified in Step 3:
   i. Link the portion of the ER numeric data element to its corresponding OPM element.
   ii. Link the data element identified in Step i to the ER textual data element identified in Step 1.
5. Repeat Steps 1–4 for each ER textual data element.

The resulting relational database will take the form of a graph structure reflecting the links between the data elements associated with a particular OPM entity. Again, the actual skeleton of the graph structure is directly derived from the OPM diagram of the system/assets under consideration.

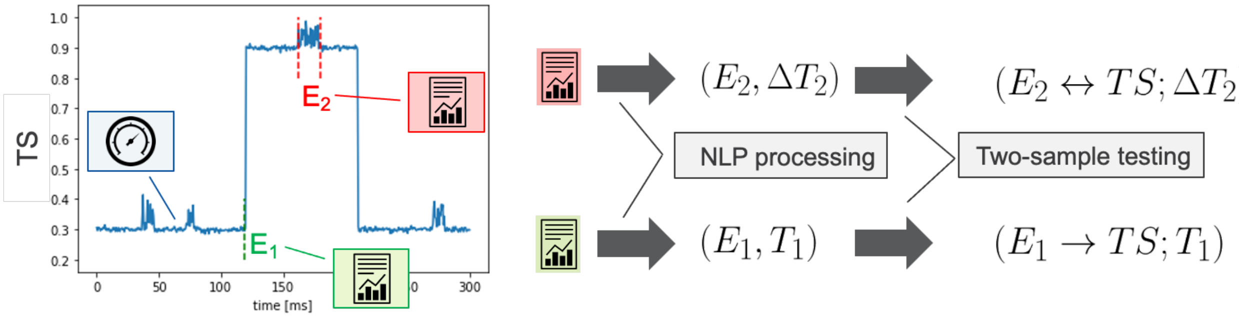

Figure 6. Identification of the temporal relations between numeric and textual events (adapted from [13][Luo, 2014]).

6. CONCLUSIONS

This article presented a summary of methods designed to analyze ER data and support system engineers' decision-making. Such ER data sources are very difficult to analyze manually due to their large size (i.e., large number of monitored sensors and data recorded over decades of operation) and their heterogenous formats (e.g., numeric and textual). However, system engineers have the valuable knowledge to interpret ER data elements and identify possible temporal or causal-effect links between them given their knowledge of the systems and assets under consideration. Our work focused on development of both models—designed to emulate system engineer knowledge in a digital form—and methods to put ER data into context. The analyses of numeric and textual ER data started separately (knowledge extraction for textual ER data and margin analysis for numeric ER data) while data integration is finally completed using the MBSE models as a common data skeleton. From there, our ER analysis methods can detect temporal and logical relationships between ER data elements to perform a first-of-its-kind application of "machine reasoning."

REFERENCES

[1] C. Okoh, R. Roy, J. Mehnen, and L. Redding, "Overview of Remaining Useful Life Prediction Techniques in Through-life Engineering Services," Procedia CIRP, vol. 16, pp. 158–163, 2014, https://doi.org/10.1016/j.procir.2014.02.006.

[2] M.G. Pecht and M. Kang, "Introduction to PHM," Prognostics and Health Management of Electronics, pp. 1–37, Aug. 2018, https://doi.org/10.1002/9781119515326.ch1.

[3] Enrico Zio, "Prognostics and Health Management of Industrial Equipment," in *Diagnostics and Prognostics of Engineering Systems: Methods and Techniques*, edited by S. Kadry, pp. 333–356. Hershey, PA, USA: IGI Global. https://doi.org/10.4018/978-1-4666-2095-7.ch017.

[4] X. Zhao et al., "Prognostics and Health Management in Nuclear Power Plants: An Updated Method-Centric Review with Special Focus on Data-Driven Methods," *Frontiers in Energy Research*, vol. 9, Jun. 2021, https://doi.org/10.3389/fenrg.2021.696785.

[5] H. Lane, H. Hapke, and C. Howard. *Natural Language Processing in Action: Understanding, Analyzing, and Generating Text with Python*. Shelter Island, New York: Manning Publications. 2019.

[6] J.M. Borky and T.H. Bradley. *Effective Model-based Systems Engineering*. Springer. 2019. https://doi.org/10.1007/978-3-319-95669-5.

[7] D. Dori and E. Crawley. *Object-Process Methodology: A Holistic Systems Paradigm*. Heidelberg: Springer. 2002. https://doi.org/10.1007/978-3-642-56209-9.

[8] S.W. Ambler, *The Object Primer: Agile Model Driven Development with UML 2.0*. Cambridge, UK: Cambridge University Press, Mar. 2004. https://doi.org/10.1017/CBO9780511584077.

[9] S. Friedenthal, A. Moore, and R. Steiner. *A Practical Guide to SysML: The Systems Modeling Language*. Morgan Kaufmann. 2008.

[10] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. New York: Springer. 2001. https://doi.org/10.1007/978-0-387-84858-7.

[11] D. Mandelli, C. Wang, and S. Hess, "On the Language of Reliability: A System Engineer Perspective," *Nuclear Technology*, vol. 209, no. 11, pp. 1637–1652, Jan. 2023, https://doi.org/10.1080/00295450.2022.2143210.

[12] D. Mandelli, C. Wang, V. Agarwal, L. Lin, and K. A. Manjunatha, "Reliability Modeling in a Predictive Maintenance Context: A Margin-Based Approach," *Reliability Engineering & System Safety*, vol. 243, pp. 109861–109861, Mar. 2024, https://doi.org/10.1016/j.ress.2023.109861.

[13] C. Luo, J.-G. Lou, Q. Lin, Q. Fu, R. Ding, D. Zhang, and Z. Wang. "Correlating events with time series for incident diagnosis," *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1583–1592, Aug. 2014. https://doi.org/10.1145/2623330.2623374.

[14] A. Gretton, K. Borgwardt, M. Rasch, B. Scholkopf, and A. Smola. *A Kernel Method for the Two-Sample-Problem*. Advances in Neural Information Processing Systems, MIT Press, vol. 19. 2006. https://doi.org/10.48550/arXiv.0805.2368.