
Industrial-Grade Time-Dependent Counterfactual Root Cause Analysis through the Unanticipated Point of Incipient Failure: a Proof of Concept

Alexandre Trilla¹ Rajesh Rajendran² Ossee Yibo³ Quentin Possamai⁴ Nenad Mijatovic³ Jordi Vitrià⁵

¹Alstom, Barcelona (Spain)

²Alstom, Bengaluru (India)

³Alstom, Saint Ouen (France)

⁴Alstom, Villeurbanne (France)

⁵Universitat de Barcelona, Barcelona (Spain)

Abstract

This paper describes the development of a counterfactual Root Cause Analysis diagnosis approach for an industrial multivariate time series environment. It drives the attention toward the Point of Incipient Failure, which is the moment in time when the anomalous behavior is first observed, and where the root cause is assumed to be found before the issue propagates. The paper presents the elementary but essential concepts of the solution and illustrates them experimentally on a simulated setting. Finally, it discusses avenues of improvement for the maturity of the causal technology to meet the robustness challenges of increasingly complex environments in the industry.

1 INTRODUCTION

The degradation of complex industrial assets is a multifaceted problem that can be explained by different factors. For instance, in the Reliability Engineering field, assets are most expected to fail either prematurely (early) during their break-in period, or late by the end of their remaining useful life (wear-out) [Dersin, P., 2023]. These failure types can be anticipated because their modes and mechanisms are well known. Moreover, their impact can be mitigated by introducing quality checks in the manufacturing process and inspection actions in their (more-or-less conservative) preventive maintenance schedule. However, for as long as the machines operate, failures can randomly appear at any point in time. This is especially challenging for dependable assets while they transit the middle region, when the failure rate is relatively low, but uniform/constant.

In this uncertain setting, the field of Predictive Maintenance tackles the problem by introducing the data as a means to closely follow the actual degradation of each asset and make better informed and timely decisions [Fink, O., Wang,

Q., Svensén, M., Dersin, P., Lee, W.-J., and Ducoffe, M., 2020]. In this sense, the detection of anomalous behaviors and the capacity to *diagnose* their *root causes* become increasingly important to guarantee the availability of the machines. Since these failures appear abruptly, they cannot be anticipated, and their evolution is not smooth as they undergo various stages of severe degradation. What is more, the available operational-service data always comes from the field, and thus it is regarded as observational time series data, where the value of the variables is always determined by their causes, not through experimentation.

The literature on causal approaches for identifying the root causes dealing with such type of data typically assume that the influences between observed processes change smoothly over time, which is introduced as a general confounding variable [Huang, B., Zhang, K., and Schölkopf, B., 2015], and modeled probabilistically to logically reason on the likelihood of an event within a certain interval [Van Houdt, G., Depaire, B., and Martin, N., 2022]. The most common approach, though, is to frame the problem around the topic of anomaly detection by modeling the normal operational regime. Specifically, Assaad, C. K., Ez-zejjari, I., and Zan, L. [2023] develop a summary graph and apply a decomposition into abstract causal relations, Budhathoki, K., Minorics, L., Blobaum, P., and Janzing, D. [2022] assess the contribution of each variable to the target outlier score using counterfactuals, Strelnikoff, S., Jammalamadaka, A., and Lu, T.-C. [2023] develop a flexible neural graph for edge attribution, Yang, W., Zhang, K., and Hoi, S. C.H. [2023] focus on abnormal data points that do not follow the regular data-generating process, and for Han, X., Zhang, L., Wu, Y., and Yuan, S. [2023], anomalies are caused by external interventions on the normal causal mechanism and therefore find the algorithmic recourse via counterfactuals to revert them.

This workshop paper exploits the fact that root causes can be found directly from the causal graph and from the time of appearance of anomalies [Assaad, C. K., Ez-zejjari, I., and Zan, L., 2023]. We follow previous works on developing a

complete causal approach, from the structured model building stage to its exploitation through probabilistic counterfactual analysis, considering the idiosyncrasies of industrial multivariate time series environments, and addressing the research gap to learn the system condition *transition* instead of detecting anomalies. The paper is organized as follows: Section 2 reviews the fundamentals of Root Cause Analysis and Causal Inference, Section 3 describes the proposed causal diagnosis method from the standpoint of Predictive Maintenance, Section 4 illustrates the implementation of the method on a synthetic system, Section 5 discusses avenues of improvement to increase the robustness of the approach, and Section 6 concludes the work.

2 BACKGROUND

2.1 ROOT CAUSE ANALYSIS

Root Cause Analysis (RCA) is a troubleshooting method of problem solving used for identifying the root causes of faults or failures [Wilson, P. F., Dell, L. D., and Anderson, G. F., 1993]. RCA is a form of deductive inference since it requires an understanding of the underlying causal mechanisms for the potential root causes and the problem, i.e., what is typically found in the context of Predictive Maintenance. RCA can be decomposed into four steps:

1. *Identify and describe* the problem clearly.
2. Establish a *timeline* from the normal situation until the failure finally occurs, through the *Point of Incipient Failure*.
3. Distinguish between the *root cause* and other causal factors.
4. Establish a *causal graph* between the root cause and the observed problem.

The trigger signal of the RCA is given by the failure timestamp (i.e., the point in time when the failure variable is observed). Then, RCA yields a list of potential root cause variables along with their probabilities, which aligns with the way complex systems fail [Cook, R. I., 2000]. The variables that comprise the data are required to be representative enough to help the developers and engineers pinpoint the source of the observed problems through the root causes and their effects [Weidl, G., Madsen, A. L., and Dahlquist, E., 2008].

2.2 STRUCTURAL CAUSAL MODEL

The causal links among the variables X that build the model of a system are assumed to be most effectively represented using the tools from the field of Causality. In this sense, the Structural Causal Model (SCM) is the framework that can most generally capture such directed associations [Pearl, J.,

2019]. The SCM defines a set of assignments governing their specific functional associations f , along with some independent noise N that accounts for everything that is not explicitly included in the model:

$$X_j := f_j(PA_j, N_j), \quad (1)$$

where PA_j represents the direct causes of the X_j variable.

If enough knowledge and experience from the field is available from the subject matter experts, i.e., strictly complying with the RCA requirements, then a complete SCM may be developed right from the start. However, this is not the typical use-case scenario in complex industrial settings, and data generally needs to be carefully leveraged to drive the development of the causal model.

2.2.1 Causal Discovery

Whenever the structure of the model is to be inferred from the observed variables, i.e., the Causal Discovery task, assumptions need to be made about the data generating process, constraints need to be applied, and usually the statistical methods of the algorithms yield different graphs that explain the same factual data [Glymour, C., Zhang, K., and Spirtes, P., 2019].

In a multivariate environment, the most straightforward approach is led by the so-called “constraint-based” discovery methods. These traditional approaches iteratively build the causal graph by utilizing a score such as the p -value of conditional independence tests. As a general technique, the Peter-Clark (PC) algorithm is described [Spirtes, P., Glymour, C., and Scheines, R., 2001]. PC is a causal network structure learning algorithm that copes well with high dimensionality and can often also identify the direction of contemporaneous links [Runge, J., Bathiany, S., Bollt, E. *et al.*, 2019]. It is consistent under i.i.d. sampling assuming no latent confounders, i.e., all relevant variables need to be observed in the data. Its outcome is a Markov Equivalence Class, and thus it is likely to have different graphical representations that explain the same observed data. The PC algorithm is especially suited to discover causality in combination with the Fisher-Z independence test because it requires less constraints for the input data [Kobayashi, S., Otomo, K., Fukuda, K., and Esaki, H., 2017].

2.2.2 Causal Bayesian Network

Once the structural graph that binds the variables is determined, the functional associations of the SCM may be learned, and this work adopts a stochastic interpretation of the world. Therefore, it treats all X as random variables, and the resulting SCM statistically describes their (conditional) probability distributions.

Considering n random variables X_1, X_2, \dots, X_n and a directed acyclic graph that relates them causally, a Causal Bayesian Network (CBN) is a generative model that has the following factorized joint probability distribution:

$$P(\mathbf{X}) = P(X_1, \dots, X_n) = \prod_{j=1}^n P(X_j | PA_j, N_j) . \quad (2)$$

The graphical nature of Bayesian networks allows seeing relationships among different variables, and their conditional dependencies enable performing probabilistic inference [Alaeddini, A., and Dogan, I., 2011]. Specifically, CBN are powerful tools for knowledge representation and inference under *uncertainty* [Pourret, O., 2008].

2.3 CAUSAL INFERENCE

Beyond probabilistic inference, Causal Inference provides the tools that allow estimating causal conclusions even in the absence of a true experiment, given that certain assumptions are fulfilled. These assumptions increase in strength as is defined in Pearl’s Causal Hierarchy (PCH) abstraction [Bareinboim, E., Correa, J. D., Ibeling, D., and Icard, T., 2022], which is summarized as follows.

2.3.1 PCH Rung 1: Associational

Describes the observational distribution of the factual data through their joint probability function $P(X)$. From this point forward, interesting quantities, i.e., the queries X_Q , can be directly computed given some evidence X_E , through their conditional probability:

$$P(X_Q | X_E) = \frac{P(X_Q, X_E)}{P(X_E)} . \quad (3)$$

This level of analysis displays a degree sophistication akin to classical (un)supervised Machine Learning techniques. As such, it is subject to confounding bias.

2.3.2 PCH Rung 2: Interventional

Describes an actionable distribution, which endows causal information at the population level. This level of analysis can be achieved through actual experimentation via Randomized Control Trials, or through statistical adjustments that smartly combine observed conditional probabilities to reduce spurious associations in the estimation. Pearl’s *do*-calculus is likely to be the most effective approach to determine the identifiability of causal effects by applying the following three rules: 1) insertion/deletion of observations, 2) action/observation exchange, and 3) insertion/deletion of actions [Pearl, J., 2012].

2.3.3 PCH Rung 3: Counterfactual

Describes a potential distribution at the individual level driven by hypothetical speculations over data that may contradict the facts. Conducting this estimation requires the following three steps [Pearl, J., Glymour, M., and Jewell, N. P., 2016]:

1. **Abduction:** Beliefs about the world are initially updated by taking into account all the evidence E given in the context. Formally, the exogenous noise probability distributions $P(U)$ are updated to $P(U|E)$.
2. **Action:** Interventions are then conducted to reflect the counterfactual assumptions, and a new causal model is thus created.
3. **Prediction:** Finally, counterfactual reasoning occurs over the new model using the updated knowledge.

3 METHOD

Since the applied industrial environment belongs to the area of Predictive Maintenance, the observation of a common development standard such as the ISO 13374 is recommended [ISO, 2003]. This specification breaks down the complexity of a problem into small modules that may be developed in isolation, thus increasing the chances of project success while improving the interpretability and explainability of the technical solution, and also help to reduce the technical debt. This section describes the Data Manipulation, State Detection, and Health Assessment processing blocks.

3.1 DATA MANIPULATION

Causality is an emergent property of complex industrial systems [Yuan, B., Zhang, J., *et al.*, 2024]. In this setting, *event* variables constitute high level, nominal, time-stamped, qualitative data records that group functions into categories and hierarchies. In fact, the causal relation is a relation among events (not properties or states) [Bunge, M., 2009]. To preprocess these collected event logs [Van Houdt, G., Depaire, B., and Martin, N., 2022], a message template extractor is typically used [Chuah, E., Kuo, S.-h., *et al.*, 2010].

3.1.1 Event Transformation

To progress with their analysis, the event variables need to be standardized into a time-series format through a transformation [Hu, X., Eklund, N., and Goebel, K., 2007]. A common approach is to utilize a counting function, which adds up the number of logged messages within a given time-slot, therefore yielding an integer-valued representation for all the variables. Other details such as the sampling rate

need to match the speed of change of the variables. Finally, the data need to be windowed on the (recent) past of the observed failure, considering a limited history in time, which needs to be sufficient to observe the evolution of the health condition of the system (from normal to failure). Once the failure has occurred, the state of the system is assumed to change as operators may take actions to mitigate its impact and prevent further damage [Li, M., Li, Z., Yin, K., Nie, X., Zhang, W., Sui, K., Pei, D., 2022].

3.1.2 Relevance Filter

Since the event-variable space may be large at this stage, it is advised to reduce it by filtering the relevant variables only. In this sense, the proposed data strategy consists of first using a robust measure of Mutual Information between any two variables [Reshef, Y. A., Reshef, D. N., Finucane, H. K., Sabeti, P. C., and Mitzenmacher, M., 2016], then discarding the non-significant variable relationships via independence testing, and finally ranking the remaining variables. At last, it is also advised to remove periodic events such as timers that are unrelated to troubleshooting, e.g., using Fourier analysis and regressions [Kobayashi, S., Otomo, K., Fukuda, K., and Esaki, H., 2017]. As a result, a collection of significant integer-valued time series variables representing the evolution of event counts is obtained.

3.2 STATE DETECTION

This module builds the data-driven SCM and conducts a *coarse-grained* diagnosis by determining if the asset under test shows a normal or abnormal working condition. However, if the resolution is not adequate in the sampled data, the causal precedence may not be observed, leading to cycles and unobserved confounding. Therefore, what makes this approach especially suited for dynamic data is the explicit consideration of time in the causal model, which is especially required to break cycles and resolve race conditions.

3.2.1 Structure Learning

The proposed approach initially infers the causal relations from observational time series event data. Nevertheless, no family or method for causal discovery in time series stands out in all situations with different characteristics [Assaad, C. K., Devijver, E., and Gaussier, E., 2022].

An initial baseline is obtained with the PC algorithm (using the Fisher-Z independence test) on data augmented with time lags. The count-based transformation described in Section 3.1.1 naturally lends itself to the application of this technique as long as the counts approximate a Gaussian distribution, which can be asserted using the Lilliefors normality test [Lilliefors, H. W., 1967].

However, the direct application of PC discovery may not be advised for certain time series cases, and other more involved methods using more powerful statistical tests (also with time lags) should be explored on top of it. In consequence, the Momentary Conditional Independence (PC-MCI) test is considered [Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., and Sejdinovic, D., 2019], which has a stronger causal detection ability based on partial correlation tests.

3.2.2 Dynamic Networks

What follows is the construction of the probabilistic model from the learned time-dependent causal structure. This approach is agnostic to any specific (non/linear) parametric functional relationship, and also provides natural access to its inherent uncertainty (even at the individual instance level).

In this sense, the Dynamic CBN (DCBN) yield a factorized representation of a stochastic process. They extend the standard causal Bayesian network formalism by providing explicit discrete temporal dimensions. DCBN represent a probability distribution over the possible histories of a time-invariant process; their advantage with respect to classical probabilistic temporal models like a Markov chain is that a DCBN is a stochastic transition model factored over a number of random variables, over which a set of conditional dependency assumptions is defined [Bobbio, A., Codetta-Raiteri, D., Montani, S., and Portinale, L., 2008].

Considering n time-dependent discrete random variables $X_1^t, X_2^t, \dots, X_n^t$, a DCBN is essentially their replication over time slices $t - \Delta$ (creating the so-called discretization steps), with the addition of a set of arcs in the graph representing the transition model, which is defined through the distribution $P(X_i^t | X_j^{t-\Delta})$, for all time-related variables i and j . Arcs connecting nodes at different time-slices ($\Delta > 0$) are called interslice edges, while arcs connecting nodes at the same slice ($\Delta = 0$) are called intraslice edges. The joint probability distribution of the DCBN is shown as follows:

$$P(\mathbf{X}) = \prod_{\forall j} P(X_j^{t-\Delta} | PA_j^{t-\Delta}) . \quad (4)$$

3.2.3 Failure Prediction

The learned DCBN shall be used to estimate the probability of the Failure variable X_F in time $P_F(t)$, which is the sink node in the model that represents the eventual system crash, given the observed data (i.e., the root causes and their effects):

$$P_F(t) = P(X_F^t | PA_F^t) . \quad (5)$$

Ideally, the probability of observing a high count of failure events $X_F^t = H$ should be a monotonically increasing function (in time) until the moment of system failure.

To detect if an anomaly is present, the Point of Incipient Failure T should be determined. This is the moment in time when the system *starts* developing an abnormal behavior that will eventually lead to the crash. Also, this is where the root cause of the observed anomaly is reasonably expected to be found. A possible strategy to determine this instant can be defined by the minimum-time significant-second-derivative of the probability of Failure, as this is the *first* inflection point with a *minimally relevant* increase Θ of risk (it may not be the greatest absolute increase, but it shall be one with the precedence in time):

$$T = \min_t \left(\frac{\partial^2}{\partial t^2} P_F(t) > \Theta \right). \quad (6)$$

While there may be many different ways to express this criterion, by using a Θ threshold parameter the subject matter experts can be easily involved in the design of the solution.

3.3 HEALTH ASSESSMENT

This module exploits the probabilistic SCM and conducts a *fine-grained* diagnosis by determining the root cause of the observed anomaly.

3.3.1 Path Finding

The hypothesis of isolation is a methodological requirement of the sciences for research; hence, the useful fiction of the isolated “causal chain” or “singled-out path” in the structure will work to the extent to which such an isolation takes place, and this is often the case in definite respects during limited intervals of time. Moreover, since every isolable process is causal, anomalies can emerge solely as a result of external perturbations [Bunge, M., 2009].

Concerning the analysis of a DCBN for RCA, estimating the most likely time-sequence chain of variables for the observed anomaly event adds explanatory value in an industrial environment. In the DCBN, each node represents an event count or state change of a variable, and the arcs represent causal–temporal relationships between the nodes. In this setting, probabilistic temporal logic determines that causes and effects are steady state formulas, the properties of which hold for the system at a certain point in time [Van Houdt, G., Depaire, B., and Martin, N., 2022], and this allows for each formula to be a path formula too where multiple variables are involved. Therefore, the causal paths shall be given by the structure of the graph: a search algorithm shall be used to traverse it and find all the routes S from the different root nodes to the sink Failure node.

For the Point of Incipient Failure T , the most likely causal path S^* that explains the anomaly data can be determined after the exhaustive search among all the potential paths S and their respective probabilities:

$$S^* = \max_{s \in S} P(s|\hat{s}); t = T, \quad (7)$$

where s represents a structural path from a source node to the sink node (i.e., the failure event variable).

Conditioning on the variables not in the path under analysis (\hat{s}) is important to block spurious associations. This is especially relevant in the case of descendants, because in the event of an anomaly, the parent/ancestor variables are preferred as precedents [Li, M., Li, Z., Yin, K., Nie, X., Zhang, W., Sui, K., Pei, D., 2022].

Finally, in addition to putting the focus on the most expected behavior, one could argue that the root cause may also have occurred in the most unexpected/irregular setting [Yang, W., Zhang, K., and Hoi, S. C.H., 2023], assuming that the most commonly experienced issues will have already been solved. This alternative perspective may also be covered in the proposed approach by minimizing the path likelihood probability.

3.3.2 Algorithmic Recourse

So far, the main focus of the analysis has been on the observed factual data. However, these data represent only one of the many potential outcomes the system could have experienced: had things been different, an alternative outcome may have been observed. Algorithmic Recourse explores these counterfactual worlds [Karimi, A.-H., Barthe, G., Schölkopf, B., and Valera, I., 2022]. Such environments are simulated via inference through (atomic) interventions α in time on a specific abnormal instance in order to revert the anomaly [Han, X., Zhang, L., Wu, Y., and Yuan, S., 2023], i.e., to lower the risk of failure X_F . This is expected to help in the recognition and understanding of the general root causes that lead to the system failure [Li, M., Li, Z., Yin, K., Nie, X., Zhang, W., Sui, K., Pei, D., 2022].

Formally, the specific retrospective reasoning that these counterfactuals explore on the anomaly, i.e., the Point of Incipient Failure at $t = T$, can be stated as:

$$P(X_F^{t=T} = L | do(X^{t=T} = \alpha), X^{t=T}, X_F^{t=T} = H). \quad (8)$$

Given that an anomaly was factually recorded in the data, i.e., through observing a high risk of failure $X_F^{t=T} = H$, i.e., a high count of failure events, Equation (8) estimates the probability that the risk would have been low at the Point of Incipient Failure $X_F^{t=T} = L$, had the root cause $X^{t=T}$ had the value α , instead of the value it actually had when the anomaly was triggered. Note that this formula does not involve regular probabilistic conditioning, but the application of the Abduction-Action-Prediction process described in Section 2.3.3.

4 RESULTS

This section elaborates on the experimental work. For further details, the code along with the description of the required software tools is available here.¹

4.1 SYSTEM DATA DESCRIPTION

For illustrative purposes, the system considered in this work is synthetic. It is comprised of 4 integer-valued time-series variables that could describe a 2-out-of-3 redundant system as follows: three full-duplex data channels that exchange messages among the devices X , and three simplex alarm channels from X to Y , which checks that the system is in good working condition, see Figure 1.




Figure 1: Time-implicit summary graph for a synthetic system that could describe a 2-out-of-3 redundancy. Note the time-confounded associations among the X channels.

The variables represent the counts of event messages over time. The window of analysis frames the timeline evolution of the system condition, from a normal operation to the crash, through the Point of Incipient Failure, where the root cause is likely to be found, see Figure 2 for a specific example of an anomaly on Channel 1. The failure is simulated as an injection of exogenous noise into one device X , which then propagates to the rest of the devices. This simplified environment should account for all of the data collection and manipulation steps.

4.2 CAUSAL MODEL LEARNING

The structure of the model is learned with time-dependent discovery algorithms to deal with the implicit confounding issue. Figure 3 shows the ground truth of the time-explicit graph of the system under analysis. In this structure, the lagged terms correspond to the channel data, and the contemporaneous terms correspond to the alarm signal.

To statistically train the causal model, a dataset of 100 instances is generated, barely over the required minimum amount of failure examples for the three potential root cause variables Lejeune, M. [2010]. Table 1 charts the causal discovery performance of PC (stable version, manually augmented with time lags) and PC-MCI algorithms over differ-




Figure 2: Timeline of system condition evolution showing a failure on Channel 1.




Figure 3: Time-explicit ground truth graph.

ent p -value configurations, using the Structural Hamming Distance, which compares the resulting graphs by computing the difference between their (binary) adjacency matrices. According to these results, the simplicity of the PC algorithm augmented with manual time lags like a buffer, which is a necessary step to fairly compare the results (PC-MCI automatically implements this), seems to be more effective to discover the structure that binds the variables of the synthetic system (all results seemed reasonable, e.g., there were no oddities such as empty graphs).

This performance analysis on learning the structural information has been conducted for research informative pur-

Table 1: Performance of causal discovery algorithms through the Structural Hamming Distance (SHD).

Algorithm (p -value)	SHD
PC-MCI (0.01)	20
PC-MCI (0.03)	20
PC-MCI (0.05)	21
PC-stable (0.01)	9
PC-stable (0.03)	9
PC-stable (0.05)	10

¹Code repository: <https://github.com/atrilla/ci4ts>

Table 2: Likelihood-ranked paths that explain the observed anomaly right after the Point of Incipient Failure $t = T+$.

Causal Path	Likelihood
$X_2^{t-1} \rightarrow X_1^t \rightarrow Y^t$	0.0505
$X_2^{t-1} \rightarrow X_3^t \rightarrow Y^t$	0.0293
$X_3^{t-2} \rightarrow X_1^t \rightarrow Y^t$	0.0208
$X_1^{t-2} \rightarrow X_3^t \rightarrow Y^t$	0.0124
$X_3^{t-1} \rightarrow X_2^t \rightarrow Y^t$	0.0080
$X_1^{t-1} \rightarrow X_2^t \rightarrow Y^t$	0.0076

poses only. The ground truth graph is used from this point forward to train the probability distributions of the DCBN model. The performance of this predictive model is evaluated using the root mean square error metric compared with the alarm variable, and it yields a score of 0.1247. In the illustrative timeline example shown in Figure 2, the predicted alarm is almost indistinguishable from the actual alarm signal.

4.3 CAUSAL DIAGNOSIS

Having a model that is able to accurately predict the observed alarm level, and therefore, the anomaly at the Point of Incipient Failure T before the issue propagates, Table 2 shows the list of paths that traverse its graph from a source (channel) node to the sink (alarm) node, ranked by their likelihood scores. Note how the most likely path explains the observed anomaly at $t = T+$, immediately after the incipient failure has occurred on Channel 1 (this diagnosis is necessarily reactive since by definition it cannot be anticipated).

Finally, to further assert the blame for the root cause, the following counterfactual is evaluated $P(Y_*^{t=T+} | do(X_1^{t=T+}), X_1^{t=T+}, Y^{t=T+})$. The resulting estimand shall adjust for the anticausal backdoor path introduced by the X_2^{t-1} confounder. Figure 4 shows the results for a range of potential alarm outcomes through their distributions. Note how the alarm level would have stayed low had the root cause X_1^t kept the values it had before the Point of Incipient Failure. Also note how the region around those values is the one that displays the least amount of uncertainty. In retrospect, the diagram also shows how this one single variable, i.e. the root cause, was sufficient to cause the anomaly that would end up with the system failure.

5 DISCUSSION

Up to this point, the solution presented in this workshop paper has described the basic principles of its causal RCA technology and an initial experimental proof of concept has




Figure 4: Distributions of potential alarm outcomes for a range of counterfactuals on the root cause channel. The descriptive averages in terms of means and standard deviations for the (discrete) failure random variable are shown.

been shown. This early stage of maturity corresponds to a standard ISO 16290 Technology Readiness Level (TRL) of 3 [ISO, 2013]. This section brainstorms some avenues of improvement to increase this robustness indicator up to higher quality standards, considering the specific challenges of complex industrial environments, especially towards pinpointing the incipient failure, which is where the root cause is most likely to be found.

5.1 VALIDATE THE TECHNOLOGY IN A RELEVANT ENVIRONMENT (TRL 4-5)

One first idea could be to improve the learning of the structure of the model. In line with similar constraint-based approaches, the consideration of tiers that heuristically stratify groups of variables can be advantageous [Andrews, B., Spirtes, P., and Cooper, G. F., 2020]. Alternatively, score-based approaches, where multiple candidate models are fit and checked, can also be explored [Glymour, C., Zhang, K., and Spirtes, P., 2019]. Finally, the usage of reductionist Functional Causal Models (FCM) may be especially helpful. A FCM represents a pairwise (or bivariate) interaction of the effect as an *analytic function* of a direct cause and some unmeasurable noise. Several forms of the FCM have been shown to be able to produce unique causal directions and have received practical applications [Huang, B., Zhang, K., and Schölkopf, B., 2015]. In specific scenarios such as multivariate time series, FCM can even improve the performance of traditional constraint-based approaches [Runge, J., Bathiany, S., Bollt, E. *et al.*, 2019]. Additionally, different independence tests can be introduced in the FCM-based discovery process to tackle heterogeneous data settings.

Second, in industrial settings where *only* the crash event at the end of the timeline is available (instead of an evolving alarm signal), a different approach shall be adopted. If

a similar solution based on regression is still pursued, a transform based on artificially prepending the failure with a “rise-time” pattern driven by prior experience could be introduced [Hu, X., Eklund, N., and Goebel, K., 2007]. Alternatively, a classification approach based on logistic regression could be adopted by treating the single crash event as a binary target variable, but data imbalance may be a concern. Finally, the kink discontinuity at the Point of Incipient Failure or the sharp discontinuity at the Failure could be exploited with a regression-based difference-in-differences technique [Abadie, A., 2005], provided that data are available after the critical event, which may not always be the case.

Finally, assessing the capacity to scale of the described solution is necessary. This goal should consider the impact of the number of variables, the length of the records, the nature of data, etc.

5.2 DEMONSTRATE THE TECHNOLOGY IN AN OPERATIONAL ENVIRONMENT (TRL 6–7)

The next enhancement idea has to do with the processing of real-world data, and while they may come from the lab (perhaps also using an accelerated degradation testing procedure), what is actually required are data from the field. However, operational data often suffer from imbalance issues, especially showing a shortage of failure instances. In this case, probably the most sensible way forward is to adopt causal approaches that initially tackle the detection of anomalies, in line with the standard pipeline of Predictive Maintenance, including point, contextual, and collective irregularities. Causality-based anomaly detection methods provide at least two significant theoretical benefits over purely statistical methods: 1) improved robustness to non-anomalous out-of-distribution data, which implies a reduction in false-alarms, and 2) a potential for failure localization due to the topological ordering of the causal graph [Strelnikoff, S., Jammalamadaka, A., and Lu, T.-C., 2023].

What is potentially weak in the current state of the art in causal RCA is the specific assessment of heteroskedasticity. The causal relationships are stationary unless an anomaly occurs [Yang, W., Zhang, K., and Hoi, S. C.H., 2023]. Anomalous data are non-stationary, and this violates one of the fundamental assumptions of time series models. Therefore, a stronger emphasis on preprocessing transformations may be necessary. Additionally, the case of time-varying exposure in the presence of time-varying confounders requires special attention [Hernán, M. A., and Robins, J. M., 2023].

Finally, it would also be interesting to relax the assumption that failures can only occur through the path of the root cause, and explore the impact of direct and indirect effects. In this case, mediation analysis is a specific application of counterfactuals that seeks to identify and explain the mech-

anism or process that underlies an observed relationship between an independent variable (i.e., the root cause) and a dependent variable (i.e., the failure effect) via the inclusion of a third variable, known as a mediator variable, an intermediary variable, or an intervening variable [VanderWeele, T. J., 2016, Agler, R., and De Boeck, P., 2017].

6 CONCLUSION

This workshop paper has developed a complete top-down counterfactual Root Cause Analysis approach from first causal inference principles that is also compliant with industrial development guidelines. On the basis of processing multivariate time series data, the focus of this diagnosis challenge has been put on detecting the Point of Incipient Failure, which displays the first anomaly pattern before the issue propagates to the rest of the system. This moment in time is believed to be where the root cause is most likely to be found. This hypothesis has been illustrated on a synthetic system showing how one single counterfactual is sufficient to explain the anomalous behavior, therefore pinpointing the root cause of the eventual failure problem.

References

- Abadie, A. Semiparametric difference-in-differences estimators. *Review of Economic Studies*, 72(1):1–19, 2005.
- Agler, R., and De Boeck, P. On the Interpretation and Use of Mediation: Multiple Perspectives on Mediation Analysis. *Frontiers in Psychology. Sec. Quantitative Psychology and Measurement*, 8, 2017.
- Alaeddini, A., and Dogan, I. Using Bayesian networks for root cause analysis in statistical process control. *Expert Systems with Applications*, 38:11230–11243, 2011.
- Andrews, B., Spirtes, P., and Cooper, G. F. On the Completeness of Causal Discovery in the Presence of Latent Confounding with Tiered Background Knowledge. *Proc. of the 23rd International Conference on Artificial Intelligence and Statistics. Proc. of Machine Learning Research*, 108:4002–4011, 2020.
- Assaad, C. K., Devijver, E., and Gaussier, E. Survey and Evaluation of Causal Discovery Methods for Time Series. *Journal of Artificial Intelligence Research*, 73:767–819, 2022.
- Assaad, C. K., Ez-zejjari, I., and Zan, L. Root Cause Identification for Collective Anomalies in Time Series given an Acyclic Summary Causal Graph with Loops. *Proc. of the 26th International Conference on Artificial Intelligence and Statistics. Proc. of Machine Learning Research*, 206, 2023.

- Bareinboim, E., Correa, J. D., Ibeling, D., and Icard, T. On Pearl's Hierarchy and the Foundations of Causal Inference. *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 507–556, 2022.
- Bobbio, A., Codetta-Raiteri, D., Montani, S., and Portinale, L. Reliability analysis of systems with dynamic dependencies. *Bayesian Networks: A Practical Guide to Applications*, pages 225–238, 2008.
- Budhathoki, K., Minorics, L., Blobaum, P., and Janzing, D. Causal structure-based root cause analysis of outliers. *Proc. of the 39th International Conference on Machine Learning. Proc. of Machine Learning Research*, 162, 2022.
- Bunge, M. *Causality and Modern Science*. Routledge, 2009.
- Chuah, E., Kuo, S.-h., *et al.* Diagnosing the root-causes of failures from cluster log files. *Proc. of the International Conference on High Performance Computing*, 2010.
- Cook, R. I. How Complex Systems Fail. *Cognitive Technologies Laboratory, University of Chicago*, 2000.
- Dersin, P. *Modeling Remaining Useful Life Dynamics in Reliability Engineering*. CRC Press, 2023.
- Fink, O., Wang, Q., Svensén, M., Dersin, P., Lee, W.-J., and Ducoffe, M. Potential, challenges and future directions for deep learning in prognostics and health management applications. *Engineering Applications of Artificial Intelligence*, 92(103678):1–15, 2020.
- Glymour, C., Zhang, K., and Spirtes, P. Review of Causal Discovery Methods Based on Graphical Models. *Frontiers in Genetics*, 10(524):1–15, 2019.
- Han, X., Zhang, L., Wu, Y., and Yuan, S. On Root Cause Localization and Anomaly Mitigation through Causal Inference. *Proc. of the 32nd ACM International Conference on Information and Knowledge Management*, 2023.
- Hernán, M. A., and Robins, J. M. *Causal Inference: What If*. CRC Press, 2023.
- Hu, X., Eklund, N., and Goebel, K. A Data Fusion Approach for Aircraft Engine Fault Diagnostics. *Proc. of ASME Turbo Expo*, 1(GT2007-27941):767–775, 2007.
- Huang, B., Zhang, K., and Schölkopf, B. Identification of Time-Dependent Causal Model: A Gaussian Process Treatment. *Proc. of the 24th International Joint Conference on Artificial Intelligence*, 2015.
- ISO. Condition monitoring and diagnostics of machine systems: Data processing, communication and presentation. Technical Report 13374-1:2003, International Organization for Standardization, 2003.
- ISO. Space systems - Definition of the Technology Readiness Levels (TRLs) and their criteria of assessment. Technical Report 16290:2013, International Organization for Standardization, 2013.
- Karimi, A.-H., Barthe, G., Schölkopf, B., and Valera, I. A Survey of Algorithmic Recourse: Contrastive Explanations and Consequential Recommendations. *ACM Computing Surveys*, 55(5):95, 2022.
- Kobayashi, S., Otomo, K., Fukuda, K., and Esaki, H. Mining causality of network events in log data. *IEEE Transactions on Network and Service Management*, 15(1):53–67, 2017.
- Lejeune, M. *Statistique, La théorie et ses applications*. Springer Verlag France, 2010.
- Li, M., Li, Z., Yin, K., Nie, X., Zhang, W., Sui, K., Pei, D. Causal Inference-Based Root Cause Analysis for Online Service Systems with Intervention Recognition. *Proc. of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3230–3240, 2022.
- Lilliefors, H. W. On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown. *Journal of the American Statistical Association*, 62(318):399–402, 1967.
- Pearl, J. The Do-Calculus Revisited. *Proc. of the 28th Conference on Uncertainty in Artificial Intelligence*, pages 4–11, 2012.
- Pearl, J. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60, 2019.
- Pearl, J., Glymour, M., and Jewell, N. P. *Causal Inference in Statistics: A Primer*. John Wiley and Sons Ltd, 2016.
- Pourret, O. Introduction to Bayesian networks. *Bayesian Networks: A Practical Guide to Applications*, pages 1–13, 2008.
- Reshef, Y. A., Reshef, D. N., Finucane, H. K., Sabeti, P. C., and Mitzenmacher, M. Measuring Dependence Powerfully and Equitably. *Journal of Machine Learning Research*, 17:1–63, 2016.
- Runge, J., Bathiany, S., Bollt, E. *et al.* Inferring causation from time series in Earth system sciences. *Nature Communications*, 10(2553):1–13, 2019.
- Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., and Sejdinovic, D. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(eaau4996):1–15, 2019.
- Spirtes, P., Glymour, C., and Scheines, R. *Causation, Prediction, and Search*. MIT Press, 2001.

- Strelnikoff, S., Jammalamadaka, A., and Lu, T.-C. Causanom: Anomaly Detection With Flexible Causal Graphs. *The International FLAIRS Conference Proceedings*, 36(1), 2023.
- Van Houdt, G., Depaire, B., and Martin, N. Root Cause Analysis in Process Mining with Probabilistic Temporal Logic. *Munoz-Gama, J., Lu, X. (eds) Process Mining Workshops. ICPM 2021. Lecture Notes in Business Information Processing*, 433:73–84, 2022.
- VanderWeele, T. J. Mediation Analysis: A Practitioner’s Guide. *Annual Review of Public Health*, 37:17–32, 2016.
- Weidl, G., Madsen, A. L., and Dahlquist, E. Decision support on complex industrial process operation. *Bayesian Networks: A Practical Guide to Applications*, pages 313–328, 2008.
- Wilson, P. F., Dell, L. D., and Anderson, G. F. *Root Cause Analysis: A Tool for Total Quality Management*. ASQ Quality Press, 1993.
- Yang, W., Zhang, K., and Hoi, S. C.H. A Causal Approach to Detecting Multivariate Time-series Anomalies and Root Causes. *Proc. of the International Conference on Learning Representations*, 2023.
- Yuan, B., Zhang, J., *et al.* Emergence and Causality in Complex Systems: A Survey of Causal Emergence and Related Quantitative Studies. *Entropy*, 26(2), 2024.