Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 59-10 (2025) 2415–2420

Network-based Root Cause Identification to Improve OEE in High-Precision Manufacturing

Felix Saretzky * **Thomas Engel** ** **Fazel Ansari** ***

* *University of Luxembourg, Department of Engineering (e-mail: felix.saretzky@uni.lu)*
** *University of Luxembourg, Department of Computer Science (e-mail: thomas.engel@uni.lu)*
*** *Technische Universität Wien, Chair of Production and Maintenance Management (e-mail: fazel.ansari@tuwien.ac.at)*

Abstract:
Small deviations in the production cycle can cause expensive downtime or quality deviations in high-volume, high-precision production lines. If no precise root cause can be identified, only the symptoms are eliminated, resulting in a pattern of repetitive failures and temporary remedial measures. This paper presents a knowledge-based framework and algorithm that combines network science and graph theory to detect anomalies and identify root causes. The approach converts multivariate time series data into temporal multiplex recurrence networks and uses eigenvalue-based anomaly detection in addition to causal process graphs (CPGs). The framework is evaluated on a simulated pick-and-place production line with four failure scenarios. This contributes to a causal and transparent identification of root causes, which will be benchmarked against other methods in future work using real manufacturing data.

Keywords: root cause analysis, fault detection and diagnosis, sensor networks, causal graphs, anomaly detection, intelligent decision support systems in manufacturing

1. INTRODUCTION

The emergence of advanced maintenance models, such as the Prescriptive Maintenance Model proposed by Ansari et al. (2019), has emphasised the importance of integrating multimodal data and advanced analytics to improve decision-making processes in industrial maintenance. One of the main objectives of such maintenance models is to reduce unplanned downtime. Small and medium sized enterprises (SMEs) are particularly affected by unplanned downtimes, as the wellbeing of many SMEs depends heavily on the supply of larger manufacturers. Industry surveys indicate that the Fast Moving Consumer Goods (FMCG) sector has seen the most significant increase in the total cost of downtime. It has doubled annually since 2019, see Siemens AG (2024).

Especially in high-volume and high-precision production lines, small deviations in the production cycle can have significant consequences such as defective products or expensive downtime. According to research by Rahm et al. (2018) and A. Schult and Majschak (2015), 70% of machine downtimes are resolved in two minutes or less, suggesting that maintenance personnel typically implement temporary fixes rather than addressing underlying causes. In today's networked systems, faults propagate, so the cause of a fault causes different symptoms in different locations. These unidentified causes of faults lead to a loss in Overall Equipment Effectiveness (OEE), as the Mean Time To Repair (MTTR) increases due to time-consuming manual diagnostic tasks, see Gram et al. (2024). According to Niggemann et al. (2014), such relationships are too extensive and too complex for heuristic approaches. This complexity highlights the need for more advanced methods to handle fault diagnosis and prediction. In contrast, the predictive maintenance methods presented in many maintenance models focus on answering the question 'What will happen when?', see Oliveira et al. (2022) and Papageorgiou et al. (2022). However, in order to be able to intervene intelligently and correct faults, the question of 'Why?' must be answered with a Cause-and-Effect principle.

To address this gap, the missing link is a Root Cause Analysis (RCA) driven by observational data, which can identify the cause-and-effect principle of production deviations using causality and expert knowledge. Consequently, the following research question arises from the discussion: How to discover cause-effect relationships of quality deviations in high-precision manufacturing machines based on observational data using advanced analytical methods to improve OEE?

This research question investigates the lack of explainable algorithms to reduce MTTR based on observational data and expert knowledge. In this paper, a novel knowledge-

10.1016/j.ifacol.2025.09.406

based framework combining methods from network science and graph theory is proposed to detect global and local anomalies and identify the root causes. The application example in this paper is the simulation of a simple pick-and-place production line consisting of two robots, three conveyor belts and cameras for quality control.

The rest of the paper is structured as follows: Section 2 presents related work to RCA. Section 3 explains the background and motivation. Section 4 presents the framework for the effective detection of anomalies and root causes. Section 5 shows the application and results of the RCA algorithm applied to four interventions. Section 6 discusses the assumptions of the framework and pathways for future research.

2. RELATED WORK

Knowledge-Based Maintenance (KBM) has become an important approach that uses data analysis, machine learning and knowledge engineering to improve the efficiency and effectiveness of maintenance processes, according to Nemeth et al. (2019). Models such as PriMa demonstrate that diagnostic maintenance techniques can be used to understand cause-and-effect relationships, which increases knowledge transparency, see Ansari et al. (2019). In addition to process and product data, these models also utilise semantic data such as expert knowledge, which highlights that a multimodal data approach is essential for RCA. A large part of the research work for data-driven RCA relates to the area of microservices. Hardt et al. (2023) introduced a comprehensive benchmark dataset for evaluating RCA algorithms in microservice systems. They evaluated six RCA approaches - three that require causal graph information (traversal, CIRCA and counterfactual attribution) and three that do not (ϵ-diagnosis, RCD and correlation-based ranking). Their experiments showed that while causal graph-based methods generally performed better, a simple correlation-based approach without graph information provides a strong baseline, particularly when working with limited data. Janzing et al. (2019) present a method for finding the root causes of an outlier event based on functional causal models and Shapley values. An end-to-end approach for RCA is proposed by Pham et al. (2024), which combines multivariate Bayesian online change point detection for anomaly detection with a non-parametric statistical hypothesis testing technique called RobustScorer. This scorer is trained on the pre-anomaly data to learn the median and interquartile range of that distribution. RobustScorer ranks the causes in descending order based on the magnitude of the scores. Current research on data-driven RCA in manufacturing remains limited, as identified by Oliveira et al. (2022). Notably, most existing publications employ artificial intelligence methodologies, which introduces significant concerns regarding interpretability and explainability, as highlighted by Papageorgiou et al. (2022). Rehak et al. (2024) propose a temporal-informed framework for RCA that combines causal graphs with anomaly detection by annotating edges with process step information. The framework evaluates potential root causes by calculating a modified Jaccard similarity score between the set of detected anomalies and the expected anomaly propagation patterns through the causal graph. Gram et al. (2024) introduced an ensemble machine learning approach for RCA. The authors use cyclic multivariate time series data from PLC sensors to obtain productivity losses in manufacturing systems. They combine an Incremental AutoEncoder for anomaly detection, PCA for structural embedding, XGBoost for cycle behaviour, and statistical methods (Pearson's Correlation and Mutual Information) into an ensemble framework that assigns weighted scores to identify the most probable root causes of productivity losses. To fill in the missing link from Section 1, the most important theoretical and practical gaps as well as challenges of data-driven RCA are summarised below. Manufacturing data is more complex compared to microservice data, as only four types of metrics such as traffic, saturation, latency and errors are typically used in these applications, see Pham et al. (2024). Manufacturing data, on the other hand, can be categorised into process data, product data and machine data. Understanding causal knowledge to diagnose the actual cause of a production deviation is necessary, as Rehak et al. (2024) argue, because if variables are only correlated, a deviation can occur by chance. Furthermore, many RCA methods are based on machine learning approaches, which often lack interpretability and causality, see Oliveira et al. (2022) and Papageorgiou et al. (2022), which is essential for a decision-making process in manufacturing. For this reason, this paper uses a combination of process data and expert knowledge with explainable methods such as network science and graph theory to identify causes of production deviation.

3. BACKGROUND AND MOTIVATION

In order to increase the OEE through data-driven RCA, it is necessary to optimise availability and quality. According to Elsayed (1996), the steady state availability A of a machine is calculated from Mean Time Between Failure (MTBF) and MTTR as follows

$$A = \frac{MTBF}{(MTBF + MTTR)} \tag{1}$$

Equation 1 shows that the response time during the diagnostic analysis is of crucial importance. Since the correction of a production deviation often requires manual work, an accurate and context-related root cause identification is the most important step to get the system operational again. For this reason, the knowledge-based method shown in Figure 1b is proposed, which combines traditional expertise with data-driven approaches. For this, a pick-and-place process in CoppeliaSim V4.6.0 rev18 Edu was implemented as an experimental setup. The experimental setup in Figure 1a shows all components of the simulation. Figure 1b shows the knowledge-based root cause identification framework. First, a causal process graph (CPG) as shown in Figure 2 is created 'offline' algorithmically or with the help of experts. In the 'online' phase, near-real-time process and product data is used to determine deviations in production. These deviations are used together with the CPG to identify possible error propagations and root causes.

4. THE DEVELOPED FRAMEWORK

This Section builds on the proposed framework using methods from network science and graph theory.

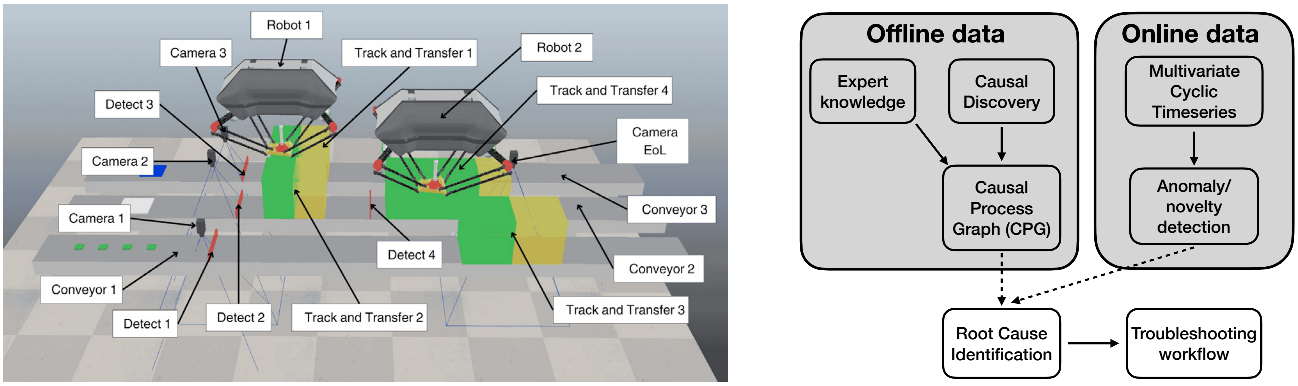

(a) Technical overview of the physical simulation

(b) Knowledge-based root cause identification framework

Fig. 1. From simulation data to root cause identification

4.1 Expert-derived Causal Process Graph

Figure 2 visualises the CPG, in which each node represents a sensor of the discrete production simulation from Figure 1a. The CPG is divided into five groups, which are highlighted in colour. Blue nodes belong to the measured values of the camera systems, which measure part dimensions and are responsible for recognising products. The grey nodes represent the measured speed of the conveyor belts. The brown nodes indicate robot joint speeds and pressure sensors of the grippers for two robots. The green nodes show the end-of-line measurements of four parts, one insert and one tray. The red node in the directed graph stands for the product evaluation (OK, NOK). The individual components and the finished product flow in the direction of the arrow from top to bottom (blue to red).

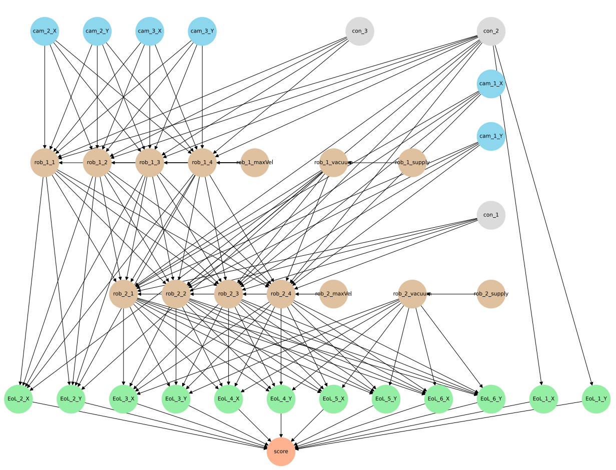

Fig. 2. A CPG represented as directed graph.

A CPG is a directed graph $G = (V, E)$ where $V = \{v_\alpha | \alpha \in \{1, \ldots, n\}\}$ is a set of sensors or process parameters that are measurable along the discrete production. E is a set of directed edges representing causal relations. If one edge $(a, b) \in E$ is pointing from $sensor_1$ to $sensor_2$, then $sensor_2$ is caused by $sensor_1$. In manufacturing, there are causal relationships based on physically connected sensors and causal relationships based on the transported product. If product 1 is transported from $sensor_1$ (folding torque) to $sensor_2$ (product length), $sensor_1$ does not directly influence $sensor_2$. The causal relationship is established via the product itself as, for example, incorrect folding can influence the length.

Based on the production flow, a causal graph can be derived using expert knowledge. The simplest form is the connection of the sensors according to the material and product flow. In the application example, it is assumed that there are no loops in the graph, as no reworking of defective products is carried out in intermediate steps. Instead of relying strictly on expert knowledge, one can use causal discovery algorithms as described in Runge et al. (2017) and Pamfil et al. (2020) to support CPG construction.

In the presented framework, the CPG of Figure 2 is used as a validation structure and for reasoning. Anomalies in individual sensors can be validated causally and temporally using the CPG in order to make a statement about the root cause.

4.2 Transforming Time Series into Temporal Multiplex Networks

In this work, recurrence-based approaches are used, based on the fact that recurrences serve as key indicators of the fundamental properties of dynamical systems, see Eroglu et al. (2018). Since discrete production repeats its production cycles, time series data is transformed into a multiplex recurrence network. To construct this network, we employ Equations 2-5, which are derived from the work of Eroglu et al. (2018). First, each time series $\{u_i\}_{i=1}^{N}$ is reconstructed as a trajectory in its phase space with a time delay.

$$x_i = (u_i, u_{i+\tau}, ..., u_{i+\tau(M-1)}), \tag{2}$$

where M is the embedding dimension, τ the embedding delay and u_i denotes $u(t_i)$. Two state vectors of the reconstructed time series can be considered recurrent as soon as the second vector falls 'close' to the first vector. The neighbourhood is defined by the threshold value ϵ. For the trajectory $\vec{x_i}$ $i = 1, ..., N, \vec{x_i} \in \mathbb{R}^M$, the adjacency matrix is defined as

$$A_{i,j} = \Theta(\epsilon - \parallel \vec{x_i} - \vec{x_j} \parallel) - \delta_{i,j}, \quad i, j = 1, ..., N, \tag{3}$$

where N is the trajectory length, Θ is the Heaviside function and $\delta_{i,j}$ is the Kronecker delta.

Applied to our usecase, each sensor v_α measures a time

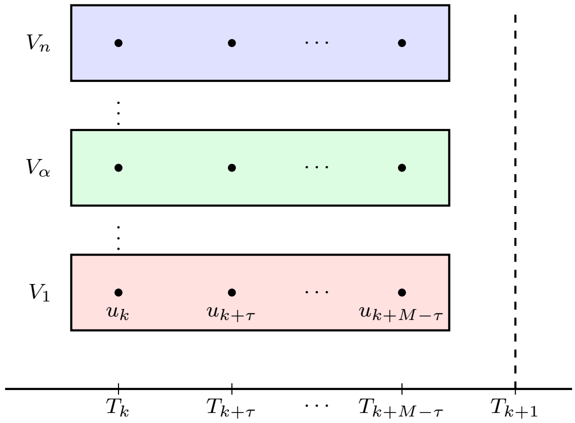

Fig. 3. The figure illustrates a temporal multiplex recurrence network. Each rectangular box reflects a recurrence network with u time steps. These are connected to each other and thus form a coupled network structure, which serves as a fingerprint for this time step.

series $\{u_{\alpha,i}\}_{i=1}^{N}$. The corresponding adjacency matrix is defined by

$$A_{i,j}^{\alpha} = \Theta(\epsilon - \parallel \overrightarrow{u_{\alpha,i}} - \overrightarrow{u_{\alpha,j}} \parallel) - \delta_{i,j}, \quad i,j = 1,...,N \quad (4)$$

and a layer graph $G_{\alpha,i}$ is provided by

$$G_{\alpha,i} = (\{u_{\alpha,i}\}_{i=1}^{N}, E_{\alpha,i}) \quad (5)$$

with $E_{\alpha,i}$ being the edge set induced by the adjacency matrix $A_{i,j}^{\alpha}$. Each node in a recurrence network represents a time step and these are connected if the measured values are within the threshold value ϵ. The superimposition of the layer graphs with identical sampling times forms a multiplex recurrence network as depicted in Figure 3.

4.3 Eigenvalue Anomaly Detection

After transforming the time series into a suitable network, it is analysed temporally on a global and layer-by-layer level. For this purpose, the approach of Huang et al. (2024) is used at the global level (Equations 6-9), where a strong change in the Laplacian spectrum of the network is considered as an anomaly or change point. To detect such a change, a low-dimensional embedding of the complex multiplex recurrence network is required at each time step. Therefore, characteristic vectors are created to help distinguish the 'normal' behaviour of the underlying production from an 'anomalous' production. First, the unnormalised graph Laplacian matrix is calculated for each layer graph $G_{\alpha,i}$ in the multiplex recurrence network by

$$L_{\alpha,i} = D_{\alpha,i} - A_{\alpha,i} \quad i = 1,...,N, \quad (6)$$

Here $D_{\alpha,i}$ is the diagonal degree matrix and $A_{\alpha,i}$ is the adjacency matrix from Equation 4, see Huang et al. (2024).

Based on $L_{\alpha,i}$, the normalised graph Laplacian matrix L_n is then calculated as follows:

$$L_{n\alpha,i} = D_{\alpha,i}^{-\frac{1}{2}} L_{\alpha,i} D_{\alpha,i}^{-\frac{1}{2}} = I_i - D_{\alpha,i}^{-\frac{1}{2}} A_{\alpha,i} D_{\alpha,i}^{-\frac{1}{2}} \quad (7)$$

Here, I_i is the identity matrix. To obtain a characteristic vector or fingerprint, the eigenvalues $\lambda_{\alpha,i}^{(0)}, ..., \lambda_{\alpha,i}^{(n-1)}$ per layer are calculated. These are then aggregated across layers using the scalar power mean operation proposed by Huang et al. (2024) as follows

$$s_p(y_1, ..., y_m) = \left(\frac{1}{m} \sum_{i=1}^{m} x_i^p \right)^{\frac{1}{p}} \quad (8)$$

The eigenvalues of the graph Laplacian matrix contain important structural information. The second smallest eigenvalue, also called algebraic connectivity or Fiedler value, can provide information about the connectivity and thus about the robustness and synchronisation. The resulting fingerprint representation

$$\Lambda_t = s_p(\lambda_{1,t}, ..., \lambda_{m,t}) \quad (9)$$

can be used to determine changes in the underlying production dynamics.
Graph measures such as Interlayer Mutual Information or Average Edge Overlap can be used to quantify information such as similarity between layers or coherence, as explained in Eroglu et al. (2018). These measures can help to detect anomalies locally per layer. In this work, we use a simple measure such as the Local Recurrence rate (LRR)

$$LRR_i = \frac{1}{N} \sum_{j=1}^{N} R_{i,j} \quad (10)$$

to determine how often the system visits states within the threshold value ϵ. A higher LRR indicates a more periodic behaviour of the dynamic system, which we expect in discrete production. The LRR is proportional to the graph measure Degree of Centrality. The error-free production cycle serves as a reference for the detection of global and local anomalies. Thus, a cyclic and dynamic system can be characterised by simple similarity measures (global) and thresholds (local) allowing anomalies to be determined.

4.4 Root Cause Identification

The root cause identification module from Figure 1b identifies potential causes of production deviations in the form of faulty sensors. To do this, the algorithm requires the 'offline' CPG, introduced in Subsection 4.1, in the form of an adjacency matrix. Also required are the points in time at which either the machine, sensors or measured values behave abnormally, which form the 'online' part of the framework. For this purpose, the anomalous nodes are stored with a time stamp in a matrix. To determine the root cause of a production deviation, the first point in time at which each anomalous sensor was classified as 'anomalous' is first searched for. Then the CPG backward traversal is applied, allowing all possible paths for each node to be determined by its predecessors. This is important to narrow down the problem and ignore anomalies that correlate with each other but have no causal relationship. Accordingly, a sensor is classified as a potential root cause if none of its predecessors in the CPG showed an anomaly at an earlier point in time when an anomaly is detected. If no such predecessors are found, the sensor is considered a potential root cause. This query is performed iteratively for all anomalous nodes and points in time to identify the root causes. The assumptions are that a directed acyclic graph is given and that the CPG is constant over time.

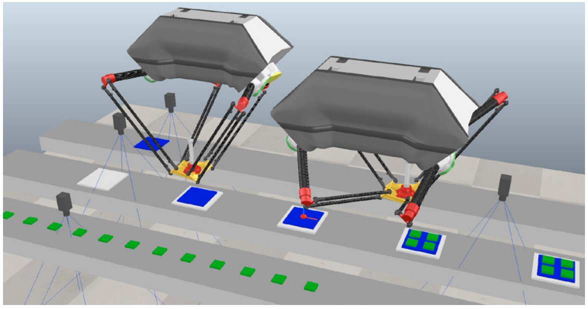

(a) Process steps of pick-and-place assembly

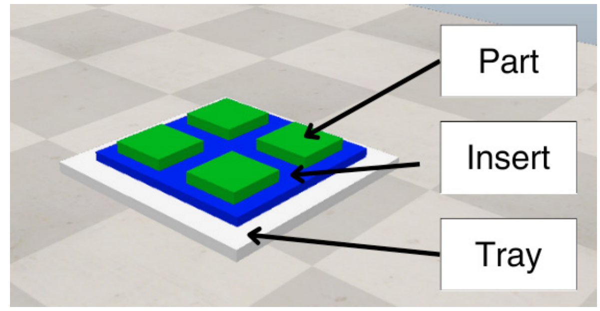

(b) Complete assembly

Fig. 4. The pick-and-place process

5. APPLICATION OF THE FRAMEWORK

In this Section it is explained how the simulation of the pick-and-place assembly is structured. A detailed description can be found in the technical report by Loureiro et al. (2024) and the simulation is also publicly available at GitHub [1].

5.1 The Assembly Process and Interventions

The assembly of the three different components (part, insert and tray) in Figure 4b is fault-free and is used as a reference assembly for detecting anomalies and analysing the root causes. Figure 1a and Figure 4a illustrate the pick-and-place process. First, three different parts are distributed to three conveyor belts. At the beginning, each conveyor belt is equipped with a camera that measures the dimensions of the parts at the entrance. In the first assembly step, the parallel robot (1) places the insert of conveyor belt (3) on the tray of conveyor belt (2). Then the parallel robot (2) places four parts of conveyor belt (1) on the insert of conveyor belt (2). At the end of the conveyor belt (2), a quality check is carried out in which all dimensions and alignments are measured.
In order to generate realistic faults, four scenarios were developed, which are explained in Table 1.

Table 1. Abnormal Scenarios

Intervention	**Scenario & root cause**
Gripper 1	Air leak. Power reduction from 100% to 30% in node rob_1_supply
Gripper 2	Air leak. Power reduction from 100% to 30% in node rob_2_supply
Part Size	Change in part dimension and weight of nodes: cam_1_x and cam_1_y
Insert not available	Change raw material feeder from 1 to 0 for nodes: cam_3_x and cam_3_y

Each intervention has either one or two sensors as the root cause. The intervention triggers a function, e.g. the pressure sensor in the robot arm. As a result, the robot arm cannot grip a part, which leads to incorrect assembly. This intervention results in fault propagation in the case of the 'air supply' for grippers 1 and 2 and a simple fault in the case of 'part size NOK'. As described in Section 4, the root cause identification algorithm can identify the error chain in addition to the original anomalous sensor or process.

[1] `https://github.com/FelixSaretzky/Pick_and_Place_Simulation`

Table 2. RCA Algorithm Results

Intervention	Top-3 recall	Top-1 recall
Gripper 1	1.00	1.00
Gripper 2	1.00	1.00
Part Size	1.00	1.00
Insert not available	0.00	0.00

5.2 Results

The Top-1 and Top-3 recall were used as evaluation metrics. Assuming that there is a unique root cause, the Top-1 recall measures the accuracy of the method to identify the correct root cause as the highest scoring suggestion. This means that the Top-1 recall indicates the proportion of cases in which the correct root cause is at the top of the list generated by the method. The Top-3 recall, on the other hand, measures the accuracy in identifying the correct root cause within the three highest-rated suggestions. Table 2 shows that the framework is able to correctly identify all interventions except for 'Insert not available'.

6. DISCUSSION AND OUTLOOK

The main objective of the paper is to introduce a knowledge-based framework and an algorithm that transparently identifies anomalies and root causes of production deviations based on time-series sensor data. This framework includes the topology of a production line represented as a CPG. This supports the proposed change of approach from correlation-based to causal, through the use of a directed acyclic graph. The basic concept of the framework can be applied in all complexity levels of maintenance defined by Ansari et al. (2019) to make the decision criteria more explainable. In general, it is assumed that anomalies that have a negative impact on OEE should be visible in the observed data. Causes of faults that cannot be explained by the observational data and the CPG, such as external influences, can therefore not be identified. While this work focused on simple point anomalies, as they are easy to detect, such anomalies are relatively rare in real-life scenarios. In future work, conceptual and theoretical work should be done on individual parts of the framework so that complex anomalies such as subsequences or complete sequences can be modelled in addition to point anomalies. Besides the pure identification of anomalies, the magnitude of the anomaly should also be taken into account. For this purpose, an outlier scorer can be introduced, as in the work of Janzing et al. (2019), which can make anomalies quantifiable and comparable. The CPG in this paper was created using expert knowledge specifically for the simulation. In the future, a creation guideline or an algorithmic method should be developed for this. Finally, this paper shows that the question of 'Why' in high-precision manufacturing can only be answered through a combination of data-driven approaches and causal modelling to fully understand both the identification and the causes of production deviations.

REFERENCES

A. Schult, E.B. and Majschak, J.P. (2015). Hin zum optimum. steigerung der effizienz von verarbeitungs- und verpackungsanlagen durch wirkungsgradanalysen. Pharma+Food, 66–68.

Ansari, F., Glawar, R., and Nemeth, T. (2019). Prima: a prescriptive maintenance model for cyber-physical production systems. volume 32, 482–503. Taylor & Francis. doi:10.1080/0951192X.2019.1571236.

Elsayed, E.A. (1996). Reliability engineering. doi:10.1002/9781119665946.

Eroglu, D., Marwan, N., Stebich, M., and Kurths, J. (2018). Multiplex recurrence networks. *Phys. Rev. E*, 97, 012312. doi:10.1103/PhysRevE.97.012312.

Gram, J., Sai, B.K., and Bauernhansl, T. (2024). Root cause analysis of productivity losses in manufacturing systems utilizing ensemble machine learning. *ArXiv*, abs/2407.21503. doi:10.15488/17728.

Hardt, M., Orchard, W., Blöbaum, P., Kasiviswanathan, S.P., and Kirschbaum, E. (2023). The petshop dataset - finding causes of performance issues across microservices. URL `https://api.semanticscholar.org/CorpusID:265051104`.

Huang, S., Coulombe, S., Hitti, Y., Rabbany, R., and Rabusseau, G. (2024). Laplacian change point detection for single and multi-view dynamic graphs. *ACM Trans. Knowl. Discov. Data*, 18(3). doi:10.1145/3631609.

Janzing, D., Budhathoki, K., Minorics, L., and Blöbaum, P. (2019). Causal structure based root cause analysis of outliers. *ArXiv*, abs/1912.02724.

Loureiro, J.F., Engel, T., and Saretzky, F. (2024). Technical report - pick place assembly simulation. URL `https://orbilu.uni.lu/handle/10993/62665`.

Nemeth, T., Ansari, F., and Sihn, W. (2019). A maturity assessment procedure model for realizing knowledge-based maintenance strategies in smart manufacturing enterprises. *Procedia Manufacturing*. URL `https://api.semanticscholar.org/CorpusID:213450206`.

Niggemann, O., Jasperneite, J., and Vodencarevic, A. (2014). *Konzepte und Anwendungsfälle für die intelligente Fabrik*, 173–190. Springer Fachmedien Wiesbaden, Wiesbaden. doi:10.1007/978-3-658-04682-8_9.

Oliveira, E., Miguéis, V.L., and Borges, J.L. (2022). Automatic root cause analysis in manufacturing: an overview & conceptualization. *Journal of Intelligent Manufacturing*, 34, 2061–2078. URL `https://api.semanticscholar.org/CorpusID:246627911`.

Pamfil, R., Sriwattanaworachai, N., Desai, S., Pilgerstorfer, P., Beaumont, P., Georgatzis, K., and Aragam, B. (2020). DYNOTEARS: Structure Learning from Time-Series Data. *arXiv e-prints*, arXiv:2002.00498.

Papageorgiou, K., Theodosiou, T.C., Rapti, A., Papageorgiou, E.I., Dimitriou, N.M., Tzovaras, D., and Margetis, G. (2022). A systematic review on machine learning methods for root cause analysis towards zero-defect manufacturing. In *Frontiers in Manufacturing Technology*. URL `https://api.semanticscholar.org/CorpusID:253164831`.

Pham, L., Ha, H., and Zhang, H. (2024). Baro: Robust root cause analysis for microservices via multivariate bayesian online change point detection. *ArXiv*, abs/2405.09330. URL `https://api.semanticscholar.org/CorpusID:269772866`.

Rahm, J., Graube, M., Müller, R., Klaeger, T., Schegner, L., Schult, A., Bonse, R., Carsch, S., Oehm, L., and Urbas, L. (2018). Kommdia: Dialogue-driven assistance system for fault diagnosis and correction in cyber-physical production systems. *2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA)*, 1, 999–1006. URL `https://api.semanticscholar.org/CorpusID:53080381`.

Rehak, J., Youssef, S., and Beyerer, J. (2024). Root cause analysis using anomaly detection and temporal informed causal graphs. In *ML4CPS–Machine Learning for Cyber-Physical Systems*, 1–7. doi:https://doi.org/10.24405/15308.

Runge, J., Nowack, P.J., Kretschmer, M., Flaxman, S., and Sejdinovic, D. (2017). Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5. URL `https://api.semanticscholar.org/CorpusID:88515817`.

Siemens AG (2024). The True Cost of Downtime 2024: How much do leading manufacturers lose through inefficient maintenance? URL `http://siemens.com/senseye-predictive-maintenance`. Article No. DICS-B10146-00-7600 PDF DÖ, Accessed: 2024-11-27.