IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 17 April 2024, accepted 20 April 2024, date of publication 26 April 2024, date of current version 6 May 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3394046

RESEARCH ARTICLE

Bayesian Learning of Causal Networks for Unsupervised Fault Diagnosis in Distributed Energy Systems

FEDERICO CASTELLETTI[1], FABRIZIO NIRO[1], MARCO DENTI[2], DANIELE TESSERA[3], (Member, IEEE), AND ANDREA POZZI[3], (Member, IEEE)

[1]Faculty of Economics, Università Cattolica del Sacro Cuore, 20123 Milan, Italy
[2]Turboden S.p.A., 25124 Brescia, Italy
[3]Faculty of Mathematical, Physical and Natural Sciences, Università Cattolica del Sacro Cuore, 25133 Brescia, Italy

Corresponding author: Andrea Pozzi (andrea.pozzi@unicatt.it)

This work was supported in part by European Union through European Social Fund (FSE) under the Recovery Assistance for Cohesion and the Territories of Europe (REACT-EU) Initiative, within the context of the National Operational Program (PON) on Research and Innovation 2014–2020, pursuant to Decreto Ministeriale [Ministerial Decree (DM)] 1062/2021 under Contract 57-I-999-6. The work of Federico Castelletti was supported in part by UCSC (D1 and 2019-D.3.2 research grants); and in part by MUR-PRIN grant 2022 SMNNKY-CUP, funded by the European Union-Next Generation EU under Grant J53D23003870008. The views and opinions expressed are only those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

ABSTRACT Distributed energy generation systems, key for producing electricity near usage points, are essential to meet the global electricity demand, leveraging diverse sources like renewables, traditional fuels, and industrial waste heat. Despite their high reliability, these systems are not immune to faults and failures. Such incidents can result in considerable downtime and reduced efficiency, underlining the need for effective fault detection and diagnosis techniques. Implementing these strategies is crucial not just for mitigating damage and preventing potential disasters, but also to maintain optimal performance levels. This paper introduces a novel methodology based on Bayesian graphical modeling for unsupervised fault diagnosis, focusing on organic Rankine cycle case study. It employs structural learning to discern unknown intervention points within a directed acyclic graph that models the power plant's operations. By analyzing real-world data, the study demonstrates the effectiveness of this approach, pinpointing a subset of variables that could be implicated in specific faults.

INDEX TERMS Clustering methods, distributed power generation, fault diagnosis, graphical models, machine learning, statistics.

I. INTRODUCTION

In the evolving landscape of energy production, Organic Rankine Cycle (ORC) power plants have emerged as a significant player, particularly in the realm of renewable and waste heat energy utilization [1]. ORC systems operate by converting thermal energy into mechanical and subsequently electrical energy, using organic fluids with lower boiling points than water. This enables them to efficiently harness energy from low-to-medium temperature heat sources, such as geothermal reservoirs, biomass combustion, and industrial waste heat. The adaptability and environmental friendliness of ORC systems make them increasingly relevant in today's pursuit of sustainable and decentralized energy solutions [2]. Nonetheless, despite their impressive reliability, it remains essential to address faults within ORC systems [3]. Even with high availability rates (higher than 98%), the occurrence of faults can still lead to reduced efficiency and increased operational costs. Effective fault management strategies, including fault detection, diagnosis, and control, are thus indispensable to maintain the integrity and efficiency of ORC power [4], [5]. These strategies ensure that ORC systems can continue to operate at their optimal capacity, providing a consistent and sustainable energy supply.

Specifically, fault detection refers to the process of identifying the occurrence of a fault within the system [6].

The associate editor coordinating the review of this manuscript and approving it for publication was Gerard-Andre Capolino.

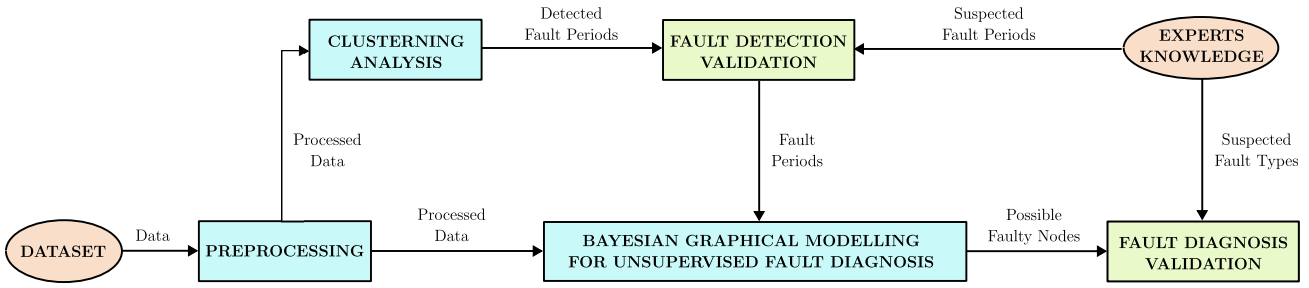

FIGURE 1. Diagram of the proposed fault diagnosis process. The process begins with dataset collection and preprocessing. Subsequent clustering analysis validates and confirms the suspected fault periods as identified by expert engineers. These periods then inform the Bayesian graphical modeling for unsupervised fault diagnosis. Finally, expert knowledge of the fault type is applied to validate the findings from the Bayesian model.

This is often achieved through monitoring system parameters and detecting deviations from normal operation. The ability to detect a fault promptly can prevent further damage and maintain operational continuity [7]. Fault diagnosis goes a step further by determining the nature and cause of the detected fault [8], [9]. It involves pinpointing the exact location or component within the system where the fault has occurred. This step is essential for targeted repairs and maintenance, preventing the unnecessary replacement of unaffected parts, and minimizing downtime [10]. Finally, fault-tolerant control becomes imperative in managing plants under faulty conditions. This approach ensures that the plant continues to operate safely and reliably, even when some components are malfunctioning, thereby enhancing the resilience of the system against faults [11]. While fault-tolerant control is an integral part of fault management, this paper primarily focuses on fault diagnosis and isolation, which are critical for accurately identifying and localizing faults within ORC systems, thereby facilitating effective and efficient interventions.

As the landscape of Organic Rankine Cycle (ORC) technologies continues to advance and broaden in application, the need for sophisticated fault diagnosis strategies becomes essential to ensure the long-term efficiency and sustainability of these crucial energy resources. In this context, the field of data-driven fault diagnosis is emerging as a pivotal area of research, offering significant advantages over traditional model-based approaches. Data-driven methods are particularly adept at handling complex, nonlinear systems where model-based approaches might fall short due to the dynamic nature and uncertainty inherent in ORC systems.

Data-driven fault diagnosis, in general, has been gaining traction across various industrial domains [12], [13], [14], [15]. These techniques leverage the power of data to uncover underlying patterns and anomalies that signal faults. Machine learning (ML) techniques, for instance, have shown considerable promise in fault diagnosis. Approaches such as logistic regression, decision trees, and support vector machines have been widely applied to detect and diagnose faults in diverse industrial settings [16]. These ML models excel at processing large datasets, identifying complex patterns indicative of system abnormalities. Deep learning (DL), an extension of ML, has also been explored for fault diagnosis with encouraging outcomes [17], [18]. DL's ability to process vast amounts of data and its proficiency in feature extraction make it a powerful tool for fault detection and diagnosis. However, both ML and DL are often constrained by their need for extensive training data and their general lack of explainability. This lack of transparency can be a significant hindrance, especially in critical applications where understanding the decision-making process is as important as the diagnosis itself. In contrast, Bayesian networks have emerged as a promising alternative, adept at managing uncertainty and incomplete data. The comprehensive survey by [19] highlights the application of Bayesian networks in fault diagnosis, underscoring their suitability for systems where not all parameters are constantly observable or predictable. Bayesian networks offer a structured approach, integrating probabilistic knowledge and expert input, making them highly relevant for fault diagnosis in complex systems.

Several works in the literature have extensively addressed the topic of fault management in power plants. While model-based techniques remain foundational, exemplified by [20], who proposed a model-based approach for fault detection and isolation in an industrial gas turbine, recent years have seen a surge in the adoption of data-driven methodologies. The authors in [21] provide an insightful survey of these advances in the field. In particular, a comprehensive framework for fault diagnosis in nuclear power plants using multi-source sensor nodes based on a Bayesian network has been presented in [22], while a step further has been taken by the authors in [23], who have introduced an online fault diagnosis method for industrial processes, utilizing a Bayesian network-based probabilistic ensemble learning strategy. This approach is particularly notable for its capacity to handle dynamic and complex industrial environments. Further contributions in this domain include the work proposed by [24] and [25], which conduct fault detection and diagnosis of a gas turbine by relying on ensemble-based approaches and advanced neural network architectures. Finally, [26] provides valuable insights into intelligent fault diagnosis based on deep learning in rotating machinery, a field widely used in pumps, wind turbine generator systems, gas turbine engines, and power plants. These studies collectively underscore the evolving landscape of fault management in power plants, where

data-driven methodologies, particularly those leveraging Bayesian networks and machine learning techniques, are paving the way for more sophisticated and effective fault diagnosis systems.

In the area of data-driven fault management for ORC systems, the study conducted in [27] stands out as a notable exception in a field where research is still emerging. This work exploits the use of a thermodynamic ORC model to generate synthetic datasets related to different types of faults. These simulated faults provide the necessary data to train various supervised learning algorithms, aiming to effectively detect and diagnose faults within the ORC system. Although the proposed approach has demonstrated reasonably strong performance on experimental data, the need for an accurate model of the ORC plant and the knowledge a-priori of the type of faults represent significant limitations to its applicability in real world scenarios.

To address existing limitations, this paper introduces an innovative approach to usupervised fault diagnosis in power generation systems through Bayesian graphical modeling. The core of the proposed methodology involves constructing a graphical model that accurately depicts the regular operational state of the plant, by relying on a dataset of collected measurements and specific annotations from expert engineers denoting the period of a fault occurrence. Based on a Bayesian graphical modeling framework, the methodology allows to quantify the fault-probability for any given node, by modeling such fault-events as exogenous interventions affecting the network structure representing the system at various (unknown) points. Such probabilities are instrumental in pinpointing critical variables that might have influenced the plant's operational changes during the fault.

The presented approach has been validated using real-world data from an operational Organic Rankine Cycle (ORC) plant. While the dataset incorporates annotations marking periods of faults, it is crucial to highlight that the learning phase of the algorithm did not utilize any specific information about the nature of these faults, thereby rendering the process essentially unsupervised. To complement this, a clustering-based approach is introduced to validate the fault period annotations provided by expert engineers, showcasing its potential for automated fault detection. The specific fault-type information was reserved for the validation stage, in which experienced engineers equipped with detailed knowledge of the plant's operations have assessed the output of the fault diagnosis algorithm, thus confirming the practical efficacy and reliability of the proposed methodology in a real-world industrial context. This approach underscores the synergy between data-driven analytical techniques and expert human validation, enhancing the robustness and applicability of fault diagnosis in ORC systems. A schematic description of the proposed pipeline is provided in Figure 1.

Summarizing, this paper aims to address the existing gaps and challenges in the field, offering a more robust, interpretable, and effective methodology for fault diagnosis in ORC systems.

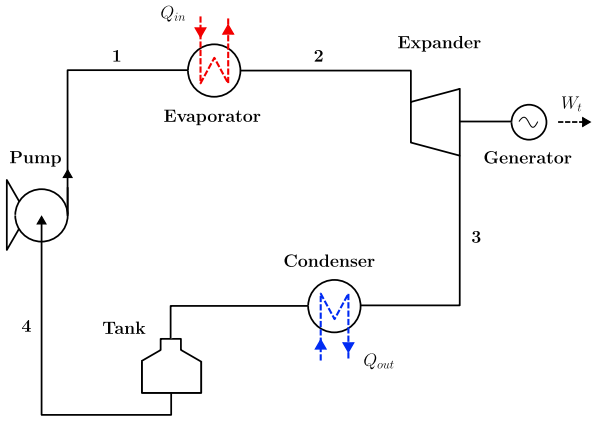

FIGURE 2. Schematic representation of an organic Rankine cycle plant. The numbers indicate the four primary stages of the cycle: evaporation (1 - 2), expansion (2 - 3), condensation (3 - 4) and pumping (4 - 1).

To provide clarity on the structure and flow of this paper, the following outline summarizes each section and its focus. In Section II a detailed description of the ORC plant is provided together with an exploratory data analysis. The proposed Bayesian methodology for fault diagnosis is presented in Section III, with a focus on the application of Bayesian graphical models. The results obtained by validating the proposed technique on real world data are described in Section IV. Finally, Section V concludes the paper.

II. ORC PLANT AND EXPLORATORY DATA ANALYSIS

This section delves into the ORC plant, serving as the foundation for the study. Firstly, a detailed description of the ORC plant is provided in II-A to establish a comprehensive understanding of its operational characteristics and dynamics. This foundational knowledge is essential for contextualizing the data analysis and the findings that follow. Secondly, an exploratory data analysis is conducted in II-B on the dataset collected from the plant. Such initial analysis is crucial not only for understanding the dataset through explanatory visualizations but also for validating the expert engineers' suspicions about the presence of a fault within a specific period, suspected of causing significant performance degradation.

A. ORC PLANT

An ORC plant operates on a principle similar to a traditional Rankine cycle, but it uses an organic fluid with a lower boiling point than water as the working fluid. This characteristic enables ORC systems to efficiently convert low-grade heat sources into electricity, making them particularly well-suited for sustainable energy generation, including waste heat recovery and renewable energy applications. The use of organic fluids with favorable thermodynamic properties allows these systems to exploit heat sources unattainable by conventional steam cycles. Additionally, ORC plants typically require less maintenance and can operate at lower temperatures, enhancing their appeal in various industrial and renewable energy contexts.

The ORC process, exemplified in Figure 2, comprises four primary stages:

- Evaporation (1 - 2): the cycle begins with the organic working fluid in a liquid state. It is heated by an external source, which could be industrial waste heat, geothermal energy, biomass, or solar heat. As the fluid absorbs this heat (generally indicated as Q_{in}), it vaporizes or turns into a high-pressure gas;
- Expansion (2 - 3): the high-pressure vapor then enters a turbine, where it expands and rotates the turbine blades. This mechanical work is converted into electric power (W_t) by a generator coupled to the turbine. As the vapor expands, its temperature and pressure drop;
- Condensation (3 - 4): after leaving the turbine, the low-pressure vapor passes through a condenser. Here, it releases its residual heat (Q_{out}) to the surroundings, often through a cooling water system or air cooling. This process turns the vapor back into a liquid;
- Pumping (4 - 1): finally, the liquid is pumped back to the evaporator, increasing its pressure and preparing it to absorb heat once again, thus completing the cycle. The work required for the pumping process is indicated with W_p.

Finally, the net work output of the cycle, representing the total useful electric power produced, is calculated as:

$$W_{net} = W_t - W_p \tag{1}$$

while the thermal efficiency of the ORC cycle is given by:

$$\eta = \frac{W_{net}}{Q_{in}} \tag{2}$$

thus providing a measure of its effectiveness in converting heat to electricity.

In order to further enhance the cycle efficiency, especially in systems where thermal resource optimization is critical, a variant of the ORC scheme with a regenerator is often considered. Such additional component, which incorporates a split system heat exchanger, is adeptly integrated into the cycle to recover a portion of the thermal energy from the vapor post-turbine and pre-condenser. The recovered heat is then utilized to preheat the working fluid, enhancing the overall thermal efficiency of the system. The regenerator's function can be described by the following steps within the ORC process:

- Regeneration: after the vapor exits the turbine, it passes through the regenerator where it donates a portion of its residual heat to the split system heat exchanger. The working fluid, now preheated by the regenerator, requires significantly less heat from the external source to reach the necessary evaporation conditions;
- Enhanced evaporation (which substitutes the evaporation phase): the preheated fluid enters the evaporator, where it becomes fully vaporized. Due to the preliminary heating from the regenerator, the amount of heat required from the external source is reduced, thereby improving the system's efficiency.

The integration of a regenerator into the ORC cycle allows for a more efficient use of thermal energy by recycling it within the system, seamlessly complementing the existing stages of the cycle. While the net work output of the cycle remains constant, the effective heat input is diminished due to the regenerator's preheating effect. This leads to an improved thermal efficiency according to (2), marking a more effective conversion of heat into electricity. The regenerator thus plays a pivotal role in enhancing the performance of the ORC cycle by optimizing the utilization of heat sources and reducing reliance on external heat inputs.

B. EXPLORATORY DATA ANALYSIS

This study utilizes a dataset comprised of real-world measurements, recorded every minute from 48 different sensors strategically placed throughout an operational ORC plant equipped with a regenerator. This extensive dataset spans a nearly five-year period, from December 2016 to April 2021, providing a comprehensive overview of the plant's operational dynamics. The available measurements encompass the generated electric power as well as temperatures, flow rates, and pressures at various points within the plant, thereby indirectly providing data on enthalpies, heat exchanges and plant efficiencies. All such variables are measured on continuous scales. Additionally, annotations from expert engineers suggest the presence of a fault, presumably a fluid leakage from the split, during a specific period (from August 2019 to April 2021).

It is pertinent to note that the dataset employed in this research is subject to a non-disclosure agreement, which precludes the open sharing of the data. Despite this limitation, we have endeavored to describe our analytical process in detail to ensure that the methodology can be understood and applied in other contexts where data access is not restricted. This approach aims to balance the need for confidentiality with the commitment to contribute valuable insights to the field.

The primary objective of the preliminary data analysis, as discussed in the following, is to verify the suspicion of expert engineers regarding the presence of a fault in the ORC plant. To achieve this, a cluster-based approach is employed, focusing on isolating low-performing samples to detect any degradation in power generation. A crucial step in this process is data preprocessing, which ensures the integrity and accuracy of the analysis. As a secondary but valuable aspect, the data preprocessing also yields useful visualizations. These visualizations not only aid in understanding the dataset but also complement the cluster-based approach in identifying and substantiating the suspected fault.

1) DATA PREPROCESSING

Variables included in the analyzed dataset were selected according to their relevance for fault diagnosis and did not contain missing values.

Outliers were rigorously identified using a criterion based on the interquartile range, focusing on data points

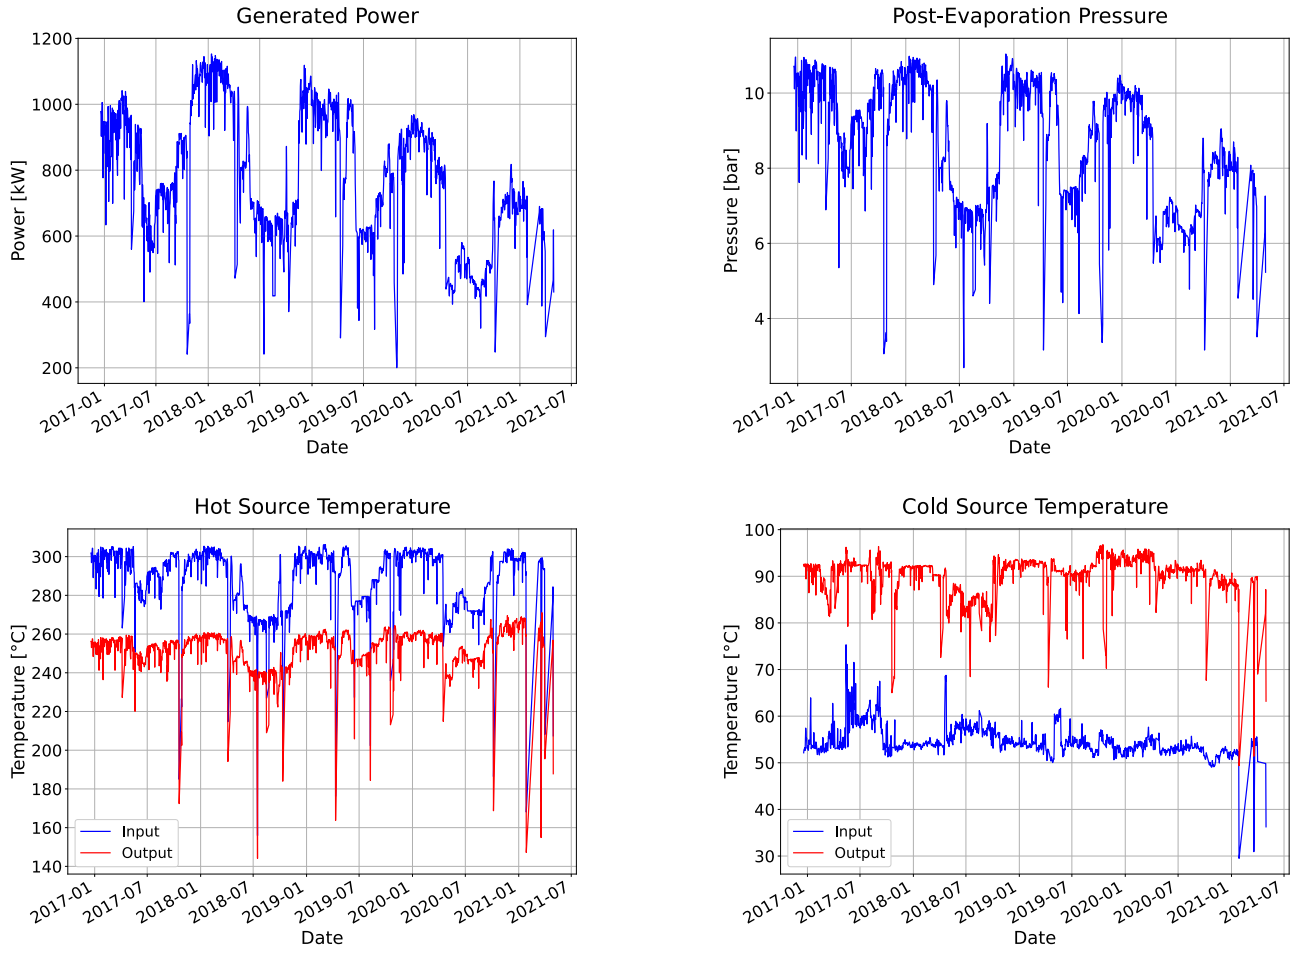

FIGURE 3. Exploratory data analysis of ORC Plant Operation. The set of plots provides a detailed overview of the ORC plant's performance, including daily generated power, pressure post-evaporation, and temperatures of both hot and cold sources, offering insights into the operational dynamics of the plant.

lying outside the 95% confidence interval. This method determined that approximately 5% of the data were outliers, primarily resulting from measurement discrepancies and transient operational disturbances not indicative of the plant's normal operational states. These outliers were carefully reviewed and subsequently excluded to ensure the analysis's integrity.

Inactive periods, marked by zero generated power, were predominantly due to operational downtimes and routine maintenance activities at the plant. Such occurrences are expected over the course of long-term data collection spanning a nearly five-year period. These inactive periods were also excluded from the analysis as they do not yield meaningful insights into the study's objectives. Then, considering the relatively slow dynamics characteristic of an ORC power plant, the minute-based data have been transformed into a daily basis. This is achieved by calculating the daily average values for each variable in the dataset. Such a transformation helps in capturing the broader trends and patterns, smoothing out short-term fluctuations that are less relevant for the plant's overall operational analysis. A graphical representation of a selection of these restructured variables is illustrated in Figure 3, providing a visual insight into the daily operational trends of the plant.

2) CLUSTER ANALYSIS FOR FAULT DETECTION

Building on the data preprocessing described earlier, this section focuses on the application of cluster analysis for fault detection in the ORC plant. The primary aim is to validate the hypothesis proposed by expert engineers regarding the occurrence of a fault during the summer of 2019, specifically indicated by a notable decrease in the plant's efficiency. However, this task is challenging due to several factors. Firstly, direct measurement of the heat input, a crucial variable for assessing efficiency, is not available. Nonetheless, indirect information about the heat input is presumed to be encapsulated within other measured variables. Secondly, the natural variability in both the heat input and generated power, influenced by seasonal variations and external environmental conditions, adds complexity to the detection process. Additionally, the plant's efficiency is a dynamic and nonlinear metric, intricately influenced by various state variables of the plant. This necessitates a more nuanced and data-driven approach, such as the cluster

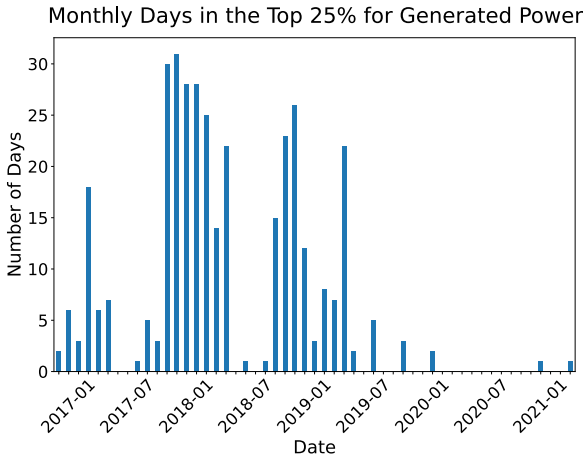

FIGURE 4. Monthly distribution of days falling in the first quartile of generated power. This histogram displays the count of days per month within the first quartile of power generation, highlighting the period of suspected fault occurrence between August 2019 and April 2021 in the ORC plant.

analysis employed here, to discern patterns and anomalies that may indicate the suspected fault.

The analysis specifically aims to determine if there is a decline in generated power under stable environmental conditions. To this end, a clustering approach is utilized to categorize the operational data of the ORC plant. The chosen methodology for achieving this is the K-means clustering algorithm, a well-established technique in data science for partitioning data into distinct groups based on similarity. The process begins by selecting a subset of variables to be used in the clustering, which is able to hold indirect information about heat input and environmental conditions, as indicated by the expert engineers. Specifically, variables related to the temperatures of both hot and cold sources and the mass flow have been considered. This phase is pivotal in distinguishing the different operational states of the plant under varied environmental scenarios.

Upon selection, these variables are normalized to ensure a consistent scale across the dataset. The K-means algorithm is then applied, segmenting the operational data of the ORC plant into clusters. For the accurate identification of the optimal number of clusters in the K-means analysis, both the elbow method and silhouette analysis are employed. These methods were instrumental in determining the most suitable clustering partition, leading to the identification of 7 distinct clusters. The silhouette score for this chosen configuration was found to be 0.41, and while this figure may not be exceptionally high, it is indicative of a reasonable structure within the dataset. This is particularly significant considering the operational data's complexity and the specific context of our case study. Furthermore, the model's inertia, which reflects the total within-cluster sum of square, was calculated to be 2055. This value further supports the efficacy of our clustering approach in distinguishing distinct operational states of the plant under varied environmental conditions.

After segmenting the data according to the established clusters, significant emphasis was placed on analyzing data points within the higher quartile of generated power. This specific focus allows for the isolation of periods of lower performance, which are potential indicators of the plant's malfunctioning. The outcomes of this scrutiny are presented in Figure 4, which shows the monthly distribution of days categorized in this higher performance bracket. A notable pattern emerged, revealing a dense concentration of high-performing data points before July 2019. This period is indicative of the plant's regular functioning. In stark contrast, the timeframe from August 2019 to April 2021 is marked by an almost complete absence of high-performing samples. This significant drop in the upper quartile performance aligns precisely with the period suspected by the engineers for a fault occurrence. The scarcity of high-efficiency data points during this latter period substantiates the hypothesis of a fault, corroborating the engineers' suspicion and indicating a deviation from the plant's normal operational efficiency.

The clustering-based approach delineated here not only serves the specific context of the ORC plant but also holds broader applicability for automatic fault detection across various domains. Particularly, this method is well-suited for scenarios where faults manifest as performance degradation, indicated by a decrease in a specific variable's value. Crucially, this approach thrives when a set of independent variables can be readily identified, typically by expert engineers, as being primarily responsible for the output. These variables essentially capture the operational conditions, making the methodology versatile and adaptable to different systems where similar patterns of fault-induced performance changes are observed.

The exploratory data analysis and the clustering-based detection presented in this paper have effectively validated the expert engineers' suspicions regarding a potential fault in the ORC plant. The forthcoming section introduces a novel methodology for diagnosing the primary causes of the detected fault, grounded in Bayesian structural learning and graphical modeling. This method transcends the limitations of traditional physics-based modeling, which, despite its precision, is often computationally intensive and excessively case-specific. In contrast, graphical modeling provides a more flexible and scalable framework, adeptly capturing the plant's complex dynamics with efficiency and facilitating easy application to other plants. Consequently, the ensuing section is devoted to deploying this innovative approach, with the objective of diagnosing the primary causes of the detected faults and fostering a more comprehensive understanding and adaptable solutions for fault diagnosis in ORC plants.

III. BAYESIAN STRUCTURE LEARNING FOR FAULT DIAGNOSIS

We consider a collection of random quantities $X = (X_j)_{j \in V}$, $V = \{1, \ldots, q\}$, corresponding to observable features of an ORC system, e.g. input/output source temperature, pressure, or generated power. We assume that the multivariate

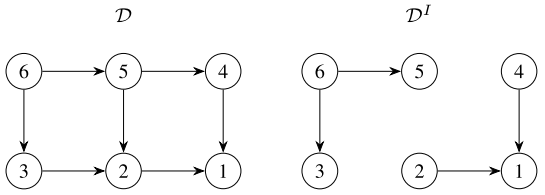

FIGURE 5. An instance of DAG $\mathcal{D}$ and post-intervention DAG $\mathcal{D}^I$ corresponding to the intervention target $I = \{2, 4\}$.

distribution of X can be represented through a directed acyclic graph (DAG for short) $\mathcal{D} = (V, E)$, where E is a set of directed edges of the form $i \to j$. DAG $\mathcal{D}$ encodes a set of structural dependencies between variables which can be representative of the functioning of a physical mechanism. Under $\mathcal{D}$, the joint distribution of X factorizes as

$$p(x \mid \mathcal{D}) = \prod_{j \in V} p(x_j \mid x_{\text{pa}(j)}) \tag{3}$$

where pa(j) are the parents of node j in the DAG, namely all nodes i for which $i \to j$. Equation (3) is known as the observational (or pre-interventional) distribution of X and describes the *regular* functioning of the system as determined by the "law of nature" [28]. In this paper, it is assumed that a fault modifies a certain set of structural dependencies embedded in $\mathcal{D}$, as the result of an external intervention affecting some nodes within the graphical representation of the plant. The rationale behind this assumption is rooted in the typical behavior of faults within such complex systems. Specifically, a fault tends to disrupt connections between different components of the plant, consequently altering the causal effect relationships that define the operational dynamics. Specifically, an intervention on target nodes $I \subseteq V$ destroys the original dependence of each $j \in I$ from its parents and replaces its observational density $p(x_j \mid x_{\text{pa}(j)})$ with $\tilde{p}(x_j)$. In formula:

$$p(x \mid \mathcal{D}, I) = \prod_{j \notin I} p(x_j \mid x_{\text{pa}(j)}) \prod_{j \in I} \tilde{p}(x_j). \tag{4}$$

We refer to Equation (4) as the post-intervention distribution and I as the intervention target. Notice that in (4) all local distributions of nodes $j \notin I$ are the same as in (3), meaning that the system is stable after intervention, save for those components that are instead affected by the intervention. The effect of such an intervention can be represented through a post-intervention DAG, $\mathcal{D}^I$, obtained from $\mathcal{D}$ by removing all edges $i \to j$, for each $j \in I$; see also Figure 5 for an example. D^I reflects the structure of the altered (as the effect of a fault) plant system. In what follows, both $\mathcal{D}$ and I are regarded as unknown and we provide a methodology for joint inference on the mechanism underlying the plant system and the nodes that are involved in the fault, $\mathcal{D}$ and I respectively.

A. MODEL FORMULATION

Consider a dataset comprising n observations from X collected under two regimes, namely a pre- and a post-intervention setting, which correspond to the regular (R) and faulty (F) conditions of the system respectively. Let $\mathbf{X} = (\mathbf{X}^R, \mathbf{X}^F)^\top$ and $\mathbf{X}_A^R$ the sub-matrix of $\mathbf{X}^R$ with columns indexed $A \subseteq V$; similarly for $\mathbf{X}_A^F$. Assuming independence across observations, the likelihood function can be written under a parametric model for X as

$$p(\mathbf{X} \mid \theta, \tilde{\theta}, \mathcal{D}, I) = \prod_{j \in V} p(\mathbf{X}_j^R \mid \mathbf{X}_{\text{pa}(j)}^R, \theta_j) \cdot \prod_{j \notin I} p(\mathbf{X}_j^F \mid \mathbf{X}_{\text{pa}(j)}^F, \theta_j) \prod_{j \in I} p(\mathbf{X}_j^F \mid \tilde{\theta}_j),$$

where $\theta = (\theta_j)_{j \in V}$, $\tilde{\theta} = (\tilde{\theta}_j)_{j \in I}$ are parameters indexing observational and interventional distributions respectively.

Our objective is to develop a methodology for learning the pair $(\mathcal{D}, I)$, namely the underlying DAG model and the intervention targets. To this end, we require a score of the form $p(\mathbf{X} \mid \mathcal{D}, I)$, known as model evidence in the Bayesian framework, which can be evaluated for a set of candidate values of $(\mathcal{D}, I)$. The quantity $p(\mathbf{X} \mid \mathcal{D}, I)$ can be derived in closed-form expression under a Bayesian model formulation based on the assumption of joint normal distribution for X. Finally, a sampling strategy implementing such score within an acceptance-rejection algorithm will return a collection of plausible values for $(\mathcal{D}, I)$ from which statistical inference can be carried out. We provide details in the next section.

B. BAYESIAN INFERENCE

We assume each conditional distribution $p(\cdot)$ in (3) and (4) being normal and specifically

$$\begin{aligned} p(x_j \mid x_{\text{pa}(j)}, \theta_j) &= \phi(x_j \mid \beta_j^\top x_{\text{pa}(j)}, \sigma_j^2), \quad j \in V, \\ \tilde{p}(x_j \mid \psi_j^2) &= \phi(x_j \mid 0, \psi_j^2), \quad j \in I, \end{aligned} \tag{5}$$

where $\phi(x \mid \mu, \sigma^2)$ denotes the probability density function of $\mathcal{N}(\mu, \sigma^2)$. Moreover, $\theta_j = (\beta_j, \sigma_j^2)$ is the parameter (regression coefficients and conditional variance) indexing the observational distribution of variable X_j given $X_{\text{pa}(j)}$, while $\tilde{\theta}_j = \psi_j^2$ the parameter (conditional variance) indexing the post-intervention distribution of X_j. To complete our Bayesian model formulation we need to assign prior distributions to θ and $\tilde{\theta}$, i.e. $p(\theta, \tilde{\theta} \mid \mathcal{D}, I)$. One can assume, under prior parameter independence,

$$p(\theta, \tilde{\theta} \mid \mathcal{D}, I) = \prod_{j \in V} p(\theta_j) \prod_{j \in I} p(\psi_j^2). \tag{6}$$

The authors in [29] show that Normal-Inverse-Gamma priors on $\theta_j = (\beta_j, \sigma_j^2)$ and Inverse-Gamma priors on ψ_j^2 provides a closed-form expression for the model evidence in favor of $(\mathcal{D}, I)$, that is the integrated likelihood

$$p(\mathbf{X} \mid \mathcal{D}, I) = \int p(\mathbf{X} \mid \theta, \tilde{\theta}, \mathcal{D}, I)\, p(\theta, \tilde{\theta} \mid \mathcal{D}, I)\, d(\theta, \tilde{\theta}). \tag{7}$$

The latter, when coupled with a prior on $(\mathcal{D}, I)$, provides a formula to compute *via* Bayes theorem the posterior probability of $(\mathcal{D}, I)$ for any candidate value of $(\mathcal{D}, I)$, that is:

$$p(\mathcal{D}, I \mid \mathbf{X}) \propto p(\mathbf{X} \mid \mathcal{D}, I)\, p(\mathcal{D}, I). \tag{8}$$

Algorithm 1 MCMC Sampler for DAG and Targets

Input: Datasets $\mathbf{X}^{(R)}$, $\mathbf{X}^{(F)}$, number of MCMC iterations S, prior hyperparameters $w \in (0, 1)$, $a \in (0, \infty)$, $b \in (0, \infty)$.
Output: S draws from the posterior of DAG and targets.
Initialize DAG and targets as $\mathcal{D}_0$ (empty DAG) and $I_0 = \emptyset$;
Set $\mathcal{D} = \mathcal{D}_0$ and $I = I_0$.

1: **for** $s = 1, \ldots, S$ **do**
2: Propose a candidate DAG $\widetilde{\mathcal{D}}$ by locally modifying $\mathcal{D}$ and update $\mathcal{D} = \widetilde{\mathcal{D}}$ based on a MH acceptance ratio;
3: Propose a candidate target $\widetilde{I}$ by adding/removing randomly a node in I, and update $I = \widetilde{I}$ based on a MH acceptance ratio;
4: Set $\mathcal{D}_s = \mathcal{D}$ and $I_s = I$.
5: **end for**

We refer the reader to [29] for technical details on the computation of $p(\mathbf{X} \mid \mathcal{D}, I)$. A prior on $\mathcal{D}$, $p(\mathcal{D})$, is assigned through a collection of Bernoulli random variables, Ber(w), $w \in (0, 1)$, indicating the absence/presence of a link between any two nodes. Similarly, prior on intervention target $p(I)$ is assigned through q independent Ber(η) priors each describing the occurrence of a fault at node/variable $j = 1, \ldots, q$. In addition, a Beta(a, b) prior is assigned to η. The resulting *hierarchical* prior was shown to provide results that are less sensitive to subjective prior specification, as for instance obtained under a fixed η value. Finally, one can set $p(\mathcal{D}, I) = p(\mathcal{D})p(I)$. In general, values of w and a, b can be assigned based on prior/expert knowledge relative to the structure of the system and the occurrence of faults at various points of the plant; alternatively, as a default objective choice one can set $w = 0.5$, $a = b = 1$. We refer to Section IV for details regarding our specific choices.

C. ALGORITHM AND COMPUTATIONAL DETAILS

A Bayesian approach to structure learning and target selection requires the computation of the posterior distribution (8) over the space of DAGs and intervention targets $(\mathcal{D}, I)$. However, a full enumeration of all DAG structures is infeasible even when the number of nodes (variables in the system) is moderate. Approximate inference based on Markov Chain Monte Carlo (MCMC) strategies is typically implemented; see for instance [30].

The sampling scheme that we consider is based on a Metropolis Hastings (MH) acceptance-rejection method which iteratively updates DAG $\mathcal{D}$ and intervention targets I based on the value attained by the model evidence $p(\mathbf{X} \mid \mathcal{D}, I)$ and the prior $p(\mathcal{D}, I)$; see also [29, Supplementary material]. A high-level illustration of our sampler is provided in Algorithm 1.

Output of the algorithm is then a collection of DAGs $\{\mathcal{D}_s\}_{s=1}^{S}$ and targets $\{I_s\}_{s=1}^{S}$ approximately sampled from the posterior (8), where S is the number of finally kept MCMC iterations.

From such MCMC output, we can provide estimates of the underlying DAG structure and targets, or even measures of uncertainty around such quantities. Features of interest are typically the posterior probabilities of edge inclusion, which can be estimated as

$$\widehat{p}(i \to j \mid \mathbf{X}) = \frac{1}{S} \sum_{t=1}^{S} \mathbb{1}\{i \to j \in \mathcal{D}_s\}. \tag{9}$$

Each probability quantifies the strength of a direct dependence between nodes i and j, that we expect to reflect the mechanism underlying the regular functioning of the plant system. In addition, by fixing a threshold for edge inclusion $k^* \in [0, 1]$ to the probabilities above, we can recover a graph estimate by including all those edges $i \to j$ for which $\widehat{p}(i \to j \mid \mathbf{X}) \geq k^*$. See also Section IV for illustrations.

More interestingly to our purposes, from the same output we can provide insights regarding the occurrence of a fault at some point (node) of the system. In particular, for any node $j \in V$ we can recover an estimate of the posterior probability of intervention as

$$\widehat{p}(j \in I \mid \mathbf{X}) = \frac{1}{S} \sum_{s=1}^{S} \mathbb{1}\{j \in I_s\}. \tag{10}$$

Each of these terms quantifies the probability of a departure from the regular functioning of the system, represented by the DAG structure $\mathcal{D}$, relative to node j. Accordingly, higher values of $\widehat{p}(j \in I \mid \mathbf{X})$ suggest stronger evidence of a malfunctioning occurred at the plant-component corresponding to node-variable X_j which can be thus considered as the responsible of the fault.

IV. RESULTS

In this section, the focus is on the application of the methodology described in Section III to the ORC plant's data.

Initially, data preprocessing and feature selection, as detailed in IV-A, are addressed. Results from the implementation of the methodology for structure learning and target identification are presented in IV-B

The resulting graph describes the conditional independencies among nodes and the probabilities of interventions affecting the system. It forms a critical component in the process of fault diagnosis, allowing for the probabilistic identification of nodes potentially impacted by external factors during the period of malfunction. The probabilities of intervention targets are then analyzed and discussed in IV-C. This discussion is instrumental to understand the operational implications of these findings. It offers valuable insights, particularly for engineers and operational personnel, by pinpointing areas within the plant that warrant closer investigation in the event of faults. Finally, details on the implementation of the proposed methodology are described in IV-D.

Thus, the results section provides a comprehensive overview, from the initial stages of data handling and feature

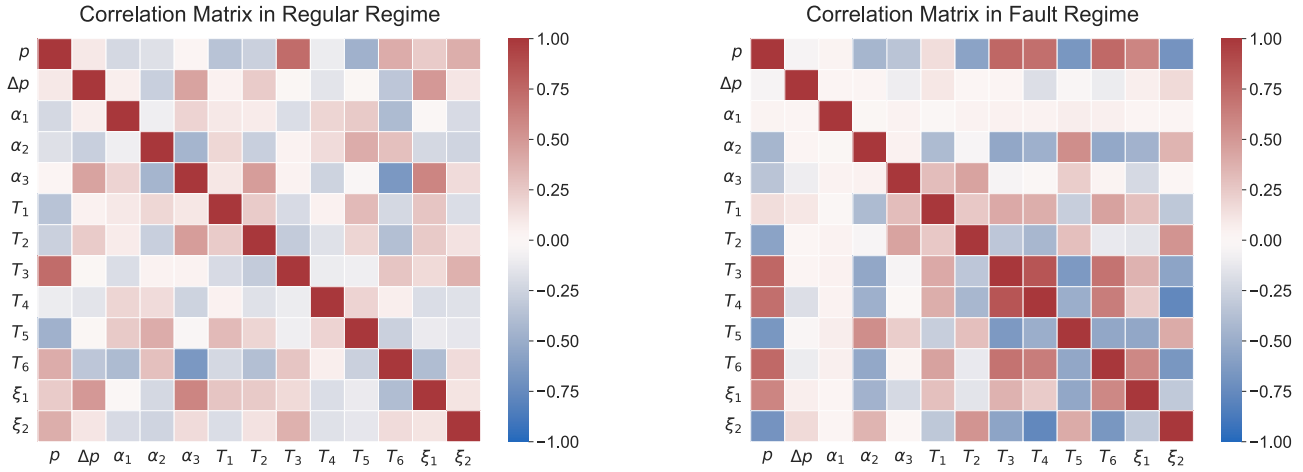

FIGURE 6. Correlation between variables for the distinct operational phases — 'Regular' and 'Fault' — of the ORC plant. It visually represents the inter-variable relationships, aiding in the understanding of variable dynamics under different operational conditions.

selection to the in-depth analysis of the learned graph and its significance in fault diagnosis.

A. DATA PROCESSING AND FEATURES SELECTION

In this subsection, the processes of data preparation and feature selection are elucidated, forming the groundwork for subsequent analytical phases. The data preparation stage was informed by insights from the company, revealing two distinct operational phases of the ORC plant: a smooth operational period preceding August 2019, and a subsequent phase characterized by malfunction, including fluid leakage and overheating, lasting from August 2019 to April 2021. To capture this dichotomy, a binary 'Fault' indicator variable was introduced, designated as 0 for periods of regular operation and 1 for the malfunction periods. Recognizing these two phases as separate experimental conditions, the study aims to discern structural differences in the plant's operation potentially resulting from unknown interventions. This differentiation is crucial for accurate fault diagnosis and understanding the plant's performance under varying conditions.

Given the dataset's complexity, particularly the significant inter-variable correlations, a strategic approach was necessary to address potential challenges in the graph learning phase, such as redundancy. To mitigate this, a variable selection process based on pairwise correlations was implemented. During the regular operational period, a maximum absolute correlation threshold of 0.75 was set. Variables with pair correlations exceeding this threshold were selectively excluded, with preference given to retaining those exhibiting a lower average correlation among pairs. This methodical selection resulted in a more streamlined dataset, where each variable adheres to the maximum absolute correlation threshold criterion in the regular phase. Consequently, the final selection comprised 16 variables, effectively reducing complexity while retaining critical information. The variables

TABLE 1. Specific variables used in the study, detailing their operational relevance within the ORC system.

Variable	Description
p	Organic Steam Pressure at Evaporator Outlet
Δp	Evaporator Steam/Liquid Differential Pressure
α_1	Organic Fluid Control Valve for Split System
α_2	Split System Thermal Oil Valve Opening Percentage
α_3	Main Thermal Oil Valve Opening Percentage
T_1	Input Temperature of Hot Water
T_2	Ambient Temperature
T_3	Organic Fluid Temperature at Condenser Outlet
T_4	Non-condensable Temperature in V103
T_5	Organic Fluid Outlet Temperature from Split System
T_6	Regenerator Outlet Temperature
ξ_1	Turbine Spindle Vibration Transmitter
ξ_2	Pump Vibration Transmitter

selected for analysis can be classified into distinct groups based on their measurement types and functions within the ORC plant:

- Pressure measurements: p and Δp;
- Organic fluid control valve in the split system: α_1;
- Operational position indicators: α_2 and α_3;
- Temperature measurements: T_1, T_2, T_3, T_4, T_5 and T_6;
- Vibration measurements: ξ_1 and ξ_2;

Each group of variables plays a pivotal role in the comprehensive analysis of the plant's operational dynamics and is instrumental in identifying the key factors contributing to its performance and efficiency. Detailed descriptions of these variables are available in Table 1.

The outcome of this selection process is visually depicted in Figure 6. The correlation map illustrates the interrelationships among the chosen features within the dataset, distinctly categorized into the 'Regular' and 'Fault' operational regimes. This graphical representation serves as a foundational reference for understanding the variables' interactions and their relevance in different operational states of the plant.

TABLE 2. Hyperparameters for the MCMC algorithm discussed in III-C.

Parameter	Value
Number of iterations (S)	1,000,000
Prior probability for target inclusion (a)	0.0625
Prior probability for target inclusion (b)	0.1
Prior probability of edge inclusion (w)	0.1

The final step in data preparation involves standardizing each feature, normalizing them to a mean of 0 and a standard deviation of 1. This normalization is imperative for maintaining a consistent scale across all variables, thereby ensuring comparability and minimizing biases that may arise from differing magnitudes of data values.

B. RESULTS OF THE STRUCTURAL LEARNING PROCESS

The structural learning process, as outlined in III-C, was applied to the dataset object of the presented case study. Key to this process was the configuration of the Markov Chain Monte Carlo (MCMC) algorithm, which was executed in Python. The parameters used for the MCMC algorithm, which are crucial for its performance and outcomes, were carefully selected based on the dataset characteristics and the objectives of the analysis. The specific parameters set for the MCMC algorithm are summarized in Table 2.

In particular a substantial number of iterations, set at 1,000,000, was selected based on empirical evidence suggesting its adequacy in facilitating the convergence of the algorithm. The hyperparameters for the Beta prior were specifically chosen to reflect the probability of target inclusion in the model. The parameter a was set as the inverse of the number of features in the dataset (16 in this case), providing a balanced approach to feature inclusion. Similarly, b was set to 0.1, offering a reasonable level of specificity for the model. Furthermore, the prior probability of edge inclusion, denoted by w, was set to 0.1. This value was selected to maintain a conservative yet flexible approach to determining relationships between nodes in the graph. It reflects a moderate level of stringency in establishing connections, ensuring that the model is neither too sparse nor overly complex.

The outcomes of the MCMC analysis, as delineated below, are instrumental in unveiling the intrinsic structure and interconnections within the dataset. These results shed light on the conditional dependencies and potential points of intervention (or faults) in the ORC plant system. As previously mentioned, the DAG derived from our model encapsulates the regular operational regime of the system. Conversely, intervention targets are identified as nodes impacted by unknown interventions, indicative of system malfunctions.

A primary focus of the MCMC results lies in the posterior probabilities of edge inclusion. The edge inclusion matrix presents the likelihood of a directed link between each pair of nodes in the learned graph during the regular regime, as shown in Figure 7.

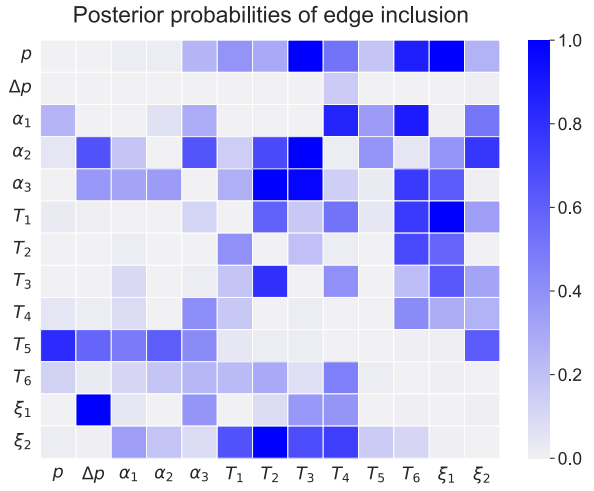

FIGURE 7. Heatmap of the posterior probabilities for each potential directed link in the ORC plant's learned graph, providing insights into the strength and likelihood of connections between different nodes in the system.

TABLE 3. Posterior probabilities for each variable within the 'Fault' data subset, representing the likelihood of interventions during the malfunctioning period as identified by the MCMC analysis. In blue variables with intervention probability higher than 0.5.

Variable	Probability of Intervention
p	0.01
Δp	0.99
α_1	0.99
α_2	0.99
α_3	0.89
T_1	0.34
T_2	0
T_3	0
T_4	0.03
T_5	0.81
T_6	0.05
ξ_1	0
ξ_2	0.46

Additionally, the MCMC output yields posterior probabilities of intervention targets. These probabilities correspond to the 'Fault' data subset – the period labeled as defective by the company engineers. Essentially, they represent the likelihood of interventions on specific nodes during the malfunctioning period of the plant system. Such probabilities are provided in Table 3 for each of the considered variables.

By applying a threshold $k^* = 0.5$ for edge inclusion, a point estimate of the posterior over DAGs is obtained. This is reported in Figure 8, where each node is colored according to the corresponding estimated posterior probability of intervention. Similarly, in order to identify which nodes are likely to be affected by faults one could consider a threshold for the posterior probability of intervention, such as 0.5. As an alternative, an Expected False Discovery Rate (EFDR) [31]) could be adopted for the purpose of optimal selection of such threshold, both for graph and target estimations.

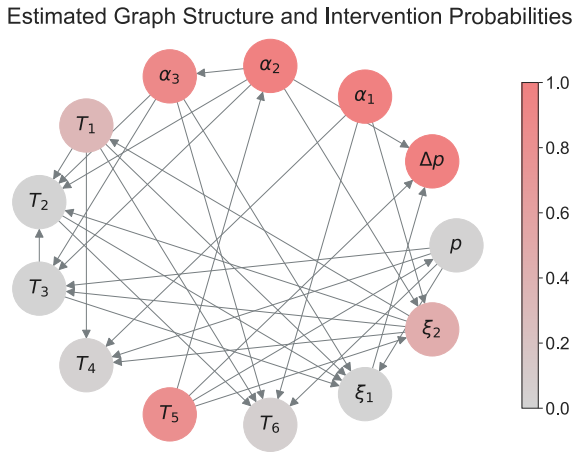

FIGURE 8. Estimated graph structure of the ORC plant system, derived by applying a threshold of $k^* = 0.5$ to the posterior probabilities of edge inclusion, which means that only edges with probabilities equal to or exceeding 0.5 are included in the graph. Each node is colored according to its estimated posterior probability of fault.

C. DISCUSSION OF POSSIBLE FAULTS

In this subsection, we delve into the interpretation and discussion of the results obtained from the MCMC analysis, particularly focusing on the posterior probabilities of interventions in various nodes of the ORC plant system. The interventions, as suggested by these probabilities, may be indicative of potential faults in the corresponding parts of the plant.

Expert engineers have pinpointed the split system heat exchanger in the ORC circuit, which recovers heat from a biomass boiler, as a critical component for optimizing the plant's thermal efficiency. In this context, the split system functions as a secondary heat exchanger, designed to recuperate residual heat from the boiler's exhaust gases. This recovered heat is then transferred to the ORC through the split system exchanger. The term split system originates from the fluid flow configuration in the ORC, where a portion of the working organic fluid is directed towards the regenerator, while another segment is diverted towards the split system exchanger. This setup allows for effective transfer of the recovered heat to the organic fluid, thereby enhancing the overall efficiency of the ORC-boiler system.

Among the variables selected in the previous section, the ones related to the split system are α_1, α_2 and T_5. The high probabilities in Table 3 associated with these variables strongly suggest a malfunction in the split system, aligning with the engineers' suspicion of a fault related to this component. Specifically, α_1, α_2, with a probability of 0.99, and T_5, showing a probability of 0.81,indicated a significant change in the operational behavior of the split system.

Note that the variables Δp and α_3, although not directly linked to the split system, exhibit considerable probabilities of intervention (0.99 and 0.89, respectively), suggesting their roles in the plant's response to the malfunction. In contrast, the other variables object of the analysis show low (i.e. less than 0.5) to negligible probabilities, indicating their lesser involvement or impact during the fault period.

In summary, the analysis of intervention probabilities provides critical insights into the plant's operational dynamics during the fault period, specifically pinpointing the split system as the focal area of concern. This information is invaluable for engineers and plant operators, guiding them towards targeted investigations and remedial actions for the faults identified in the ORC plant system.

D. IMPLEMENTATION DETAILS

To ensure the transparency and reproducibility of our research, we provide detailed information about the computational tools, libraries employed in our simulations, and data analysis, as well as the specifics of the computing environment used for these tasks. Our study utilized Python version 3.10.13, selected for its widespread acceptance within the scientific community and the comprehensive support it offers through a diverse ecosystem of libraries tailored for machine learning and data processing applications. The key Python packages that were instrumental in our simulations include:

- **NumPy (1.24.3)**: essential for its powerful numerical computing capabilities, enabling efficient handling and processing of large datasets;
- **Scikit-learn (1.3.0)**: employed for implementing the K-means clustering algorithm among other machine learning utilities;
- **Pandas (2.1.0)**: crucial for the efficient management, manipulation, and analysis of our datasets;
- **SciPy (1.11.2)**: used for optimization and statistical tasks, including operations such as `cholesky`, `cho_solve`, and `gammaln`, integral to our analyses;

while those employed for visualization are:

- **Matplotlib (3.7.3)** and **Seaborn (0.12.2)**: utilized for the generation of various plots and figures, facilitating the visual interpretation of our findings;
- **NetworkX (3.2.1)**: applied for plotting and analyzing the structure of the graphs associated with our Bayesian models.

Furthermore, to maintain consistency within our computational framework and to fully leverage the capabilities offered by Python's ecosystem, we adapted and implemented the MCMC algorithm, as presented in [29]. This algorithm played a crucial role in our methodology for learning the structure of the causal networks and assessing the probabilities of interventions within the distributed energy systems under examination.

This study's simulations were executed on a personal computing system equipped with Windows 11, powered by an Intel i7-1260P CPU, and supplemented with 32 GB of memory. This hardware configuration provided the necessary computational power and efficiency to perform the extensive simulations and data analyses required by our research.

Code implementing the proposed methodology is publicly available in our GitHub Repository, which in view of the mentioned non-disclosure agreement includes a synthetic dataset for illustration.

V. CONCLUSION

In this paper, a novel Bayesian graphical modeling approach for fault diagnosis in Organic Rankine Cycle (ORC) power generation systems was introduced and explored. Utilizing real-world operational data from an ORC plant, the methodology centers around constructing a graphical model that depicts the plant's regular functioning and estimating the probabilities of interventions at each node during the fault period. This approach successfully identified key variables that potentially contributed to operational changes during the fault, highlighting its effectiveness in fault isolation.

A significant aspect of this study is the reduced dependency on human annotations for both fault detection and diagnosis, which are here conducted through unsupervised techniques. The research demonstrated that automated clustering-based approaches could independently identify fault periods. These identified periods serve as critical inputs in a Bayesian learning framework for causal networks, thereby illuminating potential faulty nodes within the plant's systems.

Furthermore, the practical validation of the findings, conducted with experienced plant engineers, ensured that the proposed model was well-aligned with real-world plant operations and conditions. While a comparison with supervised learning benchmarks would enrich algorithm validation, it's important to note that supervised methods, though powerful in data-rich environments with extensive labeling, face significant challenges in the context of fault management in distributed energy systems like ORC plants. The difficulty lies in acquiring a detailed, labeled fault dataset broad enough to encompass all possible fault states, compounded by the high demands of expert knowledge and time for data labeling. Additionally, the inherent nature of fault data as a minority class leads to highly unbalanced datasets, posing further challenges to supervised learning approaches.

Our unsupervised learning model, designed to operate independently of labeled data, emerges as particularly beneficial under these constraints. It seeks to detect patterns and anomalies within operational data indicative of faults, relying on the dataset's inherent structure rather than extensive labeled examples. Crucially, the Bayesian nature of our approach allows the integration of external knowledge, such as expert engineers' insights into specific faults, as priors in the model's parameters. This capability to meld data-driven analysis with expert intuition significantly enhances the model's versatility and applicability, offering a bridge between purely algorithmic methods and those enriched by human expertise.

In summary, the approach presented in this paper contributes substantially to the field of data-driven fault diagnosis. By integrating Bayesian graphical modeling with actual plant data, it enhances the accuracy, interpretability, and efficiency of diagnosing faults in complex energy systems like ORC plants. This advancement is particularly relevant given the evolving landscape of fault management in power generation systems.

Looking ahead, potential areas for further research include refining the model to handle real-time data streams, thereby enhancing its predictive capabilities and responsiveness. Further exploration into the integration of this methodology with other machine learning and deep learning techniques could also yield a more comprehensive and robust framework for fault management. Additionally, the development of a unified framework that leverages graphical modeling for both fault detection and diagnosis represents a promising avenue for future research. This unified approach aims to streamline the fault management process, ensuring a cohesive and efficient strategy that harnesses the strengths of graphical modeling in addressing the complex challenges inherent in ORC systems.

Our model assumes that multivariate data are i.i.d. within each (fault and regular) regime. However, data collected from dynamical systems typically exhibit time dependencies between variables other than contemporaneous ones. One possible extension of our framework for fault detection could be based on a graphical Vector Auto Regressive (VAR) model; see for instance [32] for a Bayesian methodology. The latter would consider a regression structure among variables measured at different lags, as well as a dependence structure, represented by a directed network, modelling contemporaneous dependencies.

ACKNOWLEDGMENT

The authors express their gratitude for the support provided by three anonymous reviewers, an Associate Editor and the Editor during the revision of the manuscript.

REFERENCES

[1] A. Mahmoudi, M. Fazli, and M. R. Morad, "A recent review of waste heat recovery by organic Rankine cycle," *Appl. Thermal Eng.*, vol. 143, pp. 660–675, Oct. 2018.

[2] S. Quoilin, M. V. D. Broek, S. Declaye, P. Dewallef, and V. Lemort, "Techno-economic survey of organic Rankine cycle (ORC) systems," *Renew. Sustain. Energy Rev.*, vol. 22, pp. 168–186, Jun. 2013.

[3] M. Irl, C. Schifflechner, C. Wieland, and H. Spliethoff, "Advanced monitoring of geothermal organic Rankine cycles," *Renew. Energy*, vol. 217, Nov. 2023, Art. no. 119124.

[4] L. H. Chiang, E. L. Russell, and R. D. Braatz, *Fault Detection and Diagnosis in Industrial Systems*. Springer, 2000.

[5] R. Isermann, *Fault-Diagnosis Systems: An Introduction From Fault Detection to Fault Tolerance*. Springer, 2005.

[6] W. Sun, A. R. C. Paiva, P. Xu, A. Sundaram, and R. D. Braatz, "Fault detection and identification using Bayesian recurrent neural networks," *Comput. Chem. Eng.*, vol. 141, Oct. 2020, Art. no. 106991.

[7] D. Miljkovic, "Fault detection methods: A literature survey," in *Proc. 34th Int. Conv. MIPRO*, May 2011, pp. 750–755.

[8] A. Mesbah, S. Streif, R. Findeisen, and R. D. Braatz, "Active fault diagnosis for nonlinear systems with probabilistic uncertainties," *IFAC Proc. Volumes*, vol. 47, no. 3, pp. 7079–7084, 2014.

[9] J. K. Scott, R. Findeisen, R. D. Braatz, and D. M. Raimondo, "Design of active inputs for set-based fault diagnosis," in *Proc. Amer. Control Conf.*, Jun. 2013, pp. 3561–3566.

[10] M. Thirumarimurugan, N. Bagyalakshmi, and P. Paarkavi, "Comparison of fault detection and isolation methods: A review," in *Proc. 10th Int. Conf. Intell. Syst. Control (ISCO)*, Jan. 2016, pp. 1–6.
[11] D. M. Raimondo, G. R. Marseglia, R. D. Braatz, and J. K. Scott, "Fault-tolerant model predictive control with active fault isolation," in *Proc. Conf. Control Fault-Tolerant Syst. (SysTol)*, Oct. 2013, pp. 444–449.
[12] Y. Xu, Y. Sun, J. Wan, X. Liu, and Z. Song, "Industrial big data for fault diagnosis: Taxonomy, review, and applications," *IEEE Access*, vol. 5, pp. 17368–17380, 2017.
[13] H. Zhao, S. Sun, and B. Jin, "Sequential fault diagnosis based on LSTM neural network," *IEEE Access*, vol. 6, pp. 12929–12939, 2018.
[14] A. D. Smith, P. Dłotko, and V. M. Zavala, "Topological data analysis: Concepts, computation, and applications in chemical engineering," *Comput. Chem. Eng.*, vol. 146, Mar. 2021, Art. no. 107202.
[15] S. Jiang, S. Qin, J. L. Pulsipher, and V. M. Zavala, "Convolutional neural networks: Basic concepts and applications in manufacturing," in *Artificial Intelligence in Manufacturing*. Amsterdam, The Netherlands: Elsevier, 2024, pp. 63–102.
[16] J. Cen, Z. Yang, X. Liu, J. Xiong, and H. Chen, "A review of data-driven machinery fault diagnosis using machine learning algorithms," *J. Vib. Eng. Technol.*, vol. 10, no. 7, pp. 2481–2507, Oct. 2022.
[17] S. R. Saufi, Z. A. B. Ahmad, M. S. Leong, and M. H. Lim, "Challenges and opportunities of deep learning models for machinery fault detection and diagnosis: A review," *IEEE Access*, vol. 7, pp. 122644–122662, 2019.
[18] S. Zhang, S. Zhang, B. Wang, and T. G. Habetler, "Deep learning algorithms for bearing fault diagnostics—A comprehensive review," *IEEE Access*, vol. 8, pp. 29857–29881, 2020.
[19] B. Cai, L. Huang, and M. Xie, "Bayesian networks in fault diagnosis," *IEEE Trans. Ind. Informat.*, vol. 13, no. 5, pp. 2227–2240, Oct. 2017.
[20] H. Abbasi Nozari, M. Aliyari Shoorehdeli, S. Simani, and H. Dehghan Banadaki, "Model-based robust fault detection and isolation of an industrial gas turbine prototype using soft computing techniques," *Neurocomputing*, vol. 91, pp. 29–47, Aug. 2012.
[21] S. Khalid, J. Song, I. Raouf, and H. S. Kim, "Advances in fault detection and diagnosis for thermal power plants: A review of intelligent techniques," *Mathematics*, vol. 11, no. 8, p. 1767, Apr. 2023.
[22] G. Wu, J. Tong, L. Zhang, Y. Zhao, and Z. Duan, "Framework for fault diagnosis with multi-source sensor nodes in nuclear power plants based on a Bayesian network," *Ann. Nucl. Energy*, vol. 122, pp. 297–308, Dec. 2018.
[23] W. Yu and C. Zhao, "Online fault diagnosis for industrial processes with Bayesian network-based probabilistic ensemble learning strategy," *IEEE Trans. Autom. Sci. Eng.*, vol. 16, no. 4, pp. 1922–1932, Oct. 2019.
[24] M. Mousavi, M. Moradi, A. Chaibakhsh, M. Kordestani, and M. Saif, "Ensemble-based fault detection and isolation of an industrial gas turbine," in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Oct. 2020, pp. 2351–2358.
[25] S. Amirkhani, A. Tootchi, and A. Chaibakhsh, "Fault detection and isolation of gas turbine using series–parallel NARX model," *ISA Trans.*, vol. 120, pp. 205–221, Jan. 2022.
[26] S. Tang, S. Yuan, and Y. Zhu, "Deep learning-based intelligent fault diagnosis methods toward rotating machinery," *IEEE Access*, vol. 8, pp. 9335–9346, 2020.
[27] J. Wang, Q. Zuo, G. Liao, F. Luo, P. Zhao, W. Wu, Z. He, and Y. Dai, "Machine learning–based fault detection and diagnosis of organic Rankine cycle system for waste-heat recovery," *J. Energy Eng.*, vol. 147, no. 4, Aug. 2021, Art. no. 04021016.
[28] J. Pearl, *Causality: Models, Reasoning, Inference*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
[29] F. Castelletti and S. Peluso, "Network structure learning under uncertain interventions," *J. Amer. Stat. Assoc.*, vol. 118, no. 543, pp. 2117–2128, Jul. 2023.
[30] F. Castelletti, "Bayesian model selection of Gaussian directed acyclic graph structures," *Int. Stat. Rev.*, vol. 88, no. 3, pp. 752–775, Dec. 2020.
[31] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: A practical and powerful approach to multiple testing," *J. Roy. Stat. Soc. B, Stat. Methodol.*, vol. 57, no. 1, pp. 289–300, Jan. 1995.
[32] D. Ahelegbey, M. Billio, and R. Casarin, "Sparse graphical vector autoregression: A Bayesian approach," *Ann. Econ. Statist.*, nos. 123–124, pp. 333–361, 2016.

FEDERICO CASTELLETTI received the Ph.D. degree in statistics from Università degli Studi di Milano-Bicocca (UNIMIB), in 2018. After a few years as a Postdoctoral Researcher with Università Cattolica del Sacro Cuore, Milan, in 2021, he earned a position as a Senior Assistant Professor with the Department of Statistical Sciences, Università Cattolica del Sacro Cuore. He is currently a Lecturer in a variety of courses at undergraduate, graduate, and Ph.D. levels, including statistical learning, Bayesian statistics, and graphical modeling. He is a member of the Board with the Ph.D. Program in statistics, UNIMIB. His research interests include statistical methods based on graphical models for the analysis of multivariate data, and causal inference from observational and experimental data.

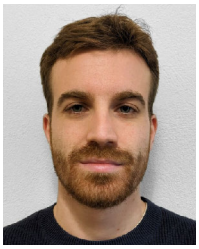

FABRIZIO NIRO received the bachelor's degree in economics and the master's degree in data analytics for business from the Faculty of Economics, Università Cattolica del Sacro Cuore, in 2021 and 2023, respectively. He is currently a Data Analytics Consultant with Reply. His research interests include Bayesian statistics, graphical modeling, and computer science.

MARCO DENTI received the master's degree in industrial engineering from Università degli Studi di Brescia, in 2006. He joined Turboden, in 2010, and he has since ascended through various roles within the company, culminating in his current position as the Chief Operating Officer (COO). Throughout his tenure with Turboden, he has cultivated a deep passion for artificial intelligence and machine learning, leveraging these interests to develop algorithms for organic rankine cycle (ORC) performance monitoring.

DANIELE TESSERA (Member, IEEE) received the Ph.D. degree in computer engineering from the University of Pavia, Italy. He is currently an Associate Professor of computer science with the Department of Mathematics and Physics, Università Cattolica del Sacro Cuore, Italy. He is the coauthor of more than 40 papers in international journals and conference proceedings. His research interests include performance debugging and benchmarking, workload characterization, cloud computing, and artificial intelligence applications.

ANDREA POZZI (Member, IEEE) received the bachelor's degree in industrial engineering, the master's degree in electrical engineering, and the Ph.D. degree in electronics, informatics, and electrical engineering from the University of Pavia, in 2015, 2017, and 2021, respectively. He has held visiting scholar position with TU Braunschweig, in 2016, and UC Berkeley, in 2019. He joined the Faculty of Mathematical, Physical and Natural Sciences, Università Cattolica del Sacro Cuore, Brescia, Italy, as an Assistant Professor of machine learning, in January 2022. After, he is currently a Postdoctoral Researcher with the University of Pavia. His research interests include reinforcement learning, imitation learning, machine learning, approximate dynamic programming, and advanced control theory.