Interactive and Intelligent Root Cause Analysis in Manufacturing with Causal Bayesian Networks and Knowledge Graphs

1st Christoph Wehner ⃝
Cognitive Systems Group
University of Bamberg
Bamberg, Germany
christoph.wehner@uni-bamberg.de

2nd Maximilian Kertel ⃝
Technology Development Battery Cell
BMW Group
Munich, Germany
maximilian.kertel@bmw.de

3rd Judith Wewerka ⃝
Digitalization and Innovation
BMW Group
Munich, Germany
judith.wewerka@bmw.de

Abstract—*Root Cause Analysis (RCA)* **in the manufacturing of electric vehicles is the process of identifying fault causes. Traditionally, the** *RCA* **is conducted manually, relying on process expert knowledge. Meanwhile, sensor networks collect significant amounts of data in the manufacturing process. Using this data for** *RCA* **makes it more efficient. However, purely data-driven methods like** *Causal Bayesian Networks* **have problems scaling to large-scale, real-world manufacturing processes due to the vast amount of potential cause-effect relationships (***CER's***). Furthermore, purely data-driven methods have the potential to leave out already known** *CER's* **or to learn spurious** *CER's***. The paper contributes by proposing an interactive and intelligent** *RCA* **tool that combines expert knowledge of an electric vehicle manufacturing process and a data-driven machine learning method. It uses reasoning over a large-scale** *Knowledge Graph* **of the manufacturing process while learning a** *Causal Bayesian Network***. In addition, an** *Interactive User Interface* **enables a process expert to give feedback to the root cause graph by adding and removing information to the** *Knowledge Graph***. The interactive and intelligent** *RCA* **tool reduces the learning time of the** *Causal Bayesian Network* **while decreasing the number of spurious** *CER's***. Thus, the interactive and intelligent** *RCA* **tool closes the feedback loop between expert and machine learning method.**

Index Terms—**Root Cause Analysis, Sensor Networks, Electric Vehicles, Interpretable Machine Learning, Interactive Learning, Bayesian Network, Knowledge Graph**

I. Introduction

Machine learning is heavily used in driver assistant systems of electric vehicles [1]. The electric vehicles benefit from machine learning not only on the road, but also during the manufacturing process as the latter has become more intelligent and data-driven. The manufacturing process is monitored via sensor networks, resulting in significant amounts of data hour by hour. The data includes faults diagnosed by quality control process steps. However, detecting a fault in a complex manufacturing process does not necessarily uncover in which *Process Steps* the fault was induced. Finding the cause-effect relationships *(CER's)* leading to the diagnosed fault is called *Root Cause Analysis (RCA)* [2].

Traditionally, *RCA* is a manual, labor-intense, and thus expensive process [3, 4, 5], involving design of experiments to study manipulations and resulting effects. *RCA* requires high amounts of process expert knowledge. Such that only process experts can conduct a *RCA* in a reasonable timeframe.

Digital support tools promise to make *RCA* more efficient, data-driven, and less dependent on individual process knowledge [2], to reduce costs and improve the performance of the manufacturing process.

Previous work identified *Causal Bayesian Networks* as a machine learning technique to automate *RCA* [6] by learning *CER's* between *Sensor Variables*. The *Causal Bayesian Networks* constructs a root cause graph over the manufacturing process, showing potential root causes [7]. However, learning *Causal Bayesian Networks* for large manufacturing processes becomes prohibitively expensive as the search space for potential cause-effect relations explodes. This problem can be mitigated by including process expert knowledge in the learning process of the *Causal Bayesian Networks* [8]. In addition, expert knowledge can improve the learned root cause graph by identifying spurious *CER's*.

The remaining challenges for scaling *Causal Bayesian Networks* to large-scale manufacturing processes are to model process expert knowledge in detail and to provide an accessible way for the process expert to interact with and improve the root cause graph. The paper investigates those challenges by combining a large-scale *Knowledge Graph* of the manufacturing process with a *Causal Bayesian Network*. The *Knowledge Graph* allows detailed modeling of large-scale manufacturing processes and enables the process expert to improve the root cause graph.

In this work, a support tool for interactive and intelligent *RCA* in the manufacturing process of electric vehicles is proposed. The *RCA* is conducted by finding *CER's* between *Sensor Variables* of the manufacturing process.

In detail, the paper contributes by:
- Detailed modelling of process expert knowledge in a large-scale *Knowledge Graph* for a real-world electric vehicle manufacturing process.
- Automatically considering manufacturing knowledge from the *Knowledge Graph* to drastically prune the search

space while learning the *Causal Bayesian Network*.

- Allowing the process expert to interact with and improve the root cause graph by explaining the *Causal Bayesian Network* where and how to do better.

II. RELATED WORK

RCA in manufacturing is dominated by methods defined in the ISO/IEC 31010 [9], including versions of the *Five Why's* [3], the *Failure Mode and Effect Analysis (FMEA)* [4], or the *Fault Tree Analysis* [5]. The methods structure expert knowledge, such that a process expert can use it as a manual for *RCA*. However, the decision-support tools do not take advantage of today's sensor networks and data.

Thus, machine learning methods were introduced in *RCA*, enabling automated and intelligent decision-support tools [10, 2]. In particular, *Causal Bayesian Networks* were identified as beneficial for *RCA* by learning *CER's* between measurements [11, 12, 13, 14].

For large and complex manufacturing processes, the number of possible *Causal Bayesian Network* grows super-exponentially and its derivation becomes challenging [15]. Additionally, one is not only interested in the existence of *CER's*, but also in their strength to prioritize potential root causes. The derivation of the *CER's* is called *structure learning*, while the identification of the strengths is called *parameter learning*. [7] proposes an *FMEA*-based approach, while [6] proposes an ontology of the manufacturing process to determine the structure of the *Causal Bayesian Network*. However, the ontology only considers classes of entities and does not allow pruning the search space for individual measurements. Furthermore, in complex manufacturing processes many *CER's* are unknown to the experts and both methods are unable to uncover them. On the contrary, [15] combines expert knowledge and data to derive the *Causal Bayesian Network*. This idea improves the quality of the root cause graph, as shown by [16]. While *Sensor Variables* can be continuous, most of the aforementioned approaches consider a discretization to decrease the complexity of learning the *Causal Bayesian Network* [13]. However, this practice may lead to such a degree of information loss that the data-driven identification of *CER's* becomes infeasible [17]. Recent approaches [18, 19] for learning complex *Causal Bayesian Networks* based on *Structural Equation Models* (*SEM's*) were applied on manufacturing processes [8, 20] and the benefits of the inclusion of expert knowledge were demonstrated [8].

In this work, we present a large-scale modeling of domain expertise into a *Knowledge Graph* that leverages the sensor network topology for deriving the *Causal Bayesian Network* of a real-world manufacturing process of electric vehicles. Furthermore, we show how the *Knowledge Graphs* contributes to an interactive *RCA* and closes the feedback loop for the *Causal Bayesian Network* learning algorithm of [8].

III. SYSTEM OVERVIEW

The system for interactive and intelligent *RCA* consists of five components (cf. Figure 1), where each is a microservice

on its own.

Fig. 1. The system architecture of the interactive and intelligent root cause analysis tool.

The *Controller* and the *Data Layer* enable the core components. The *Knowledge Graph*, *Causal Bayesian Network*, and *Interactive User Interface* are the core components of the interactive and intelligent root cause analysis tool.

The following describes the components' high-level functionalities and their interactions. Later on, the three core components are described in detail.

The heart of the system is a *Controller* that handles the communication between all other components. It receives requests and triggers processes.

The *Data Layer* accesses data from a vehicle manufacturing-sensor network. Next, the data is preprocessed according to the input format of the *Causal Bayesian Network* and stored in the *Data Layer*. The preprocessing shall be elaborated in Subsection IV-A.

Another primary source of data is the *Knowledge Graph* (cf. Subsection III-A). The *Knowledge Graph* holds formalized expert knowledge about the manufacturing process of electric vehicles, like the temporal-spatial relations of *Stations*, *Process Steps*, and *Sensor Variables*.

The *Causal Bayesian Network* takes the preprocessed sensor data from the *Data Layer* for learning *CER's* between the *Sensor Variables* (cf. Subsection III-B). The information from the *Knowledge Graph* is used while learning to reduce the search space drastically. The resulting *Causal Bayesian Network* shows, which *Sensor Variables* may induce faults in other *Sensor Variables*. Thus, we call the graph learned by the *Causal Bayesian Network* a root cause graph.

The root cause graph is visualized in an *Interactive User Interface* (cf. Subsection III-C). The visualization enables experts to explore root causes in the manufacturing process of electric vehicles. Additionally, it allows for feedback of the process expert on the root cause graph. For example, the process expert may add an edge between two *Sensor Variables* to the *Knowledge Graph* or defines a *Sensor Variable* as a *Root Variable*.

A. Knowledge Graph

A *Knowledge Graph* [21] is a directed labeled graph [22, 23] and, therefore, optimal for storing highly relational data while

preserving its semantics and allowing for deductive reasoning [21].

The *Knowledge Graph* of the *RCA* tool formalizes expert knowledge of the manufacturing process of electric vehicles. This allows the automatic, efficient, and large-scale mining of expert knowledge, which is then used while learning the *Causal Bayesian Network*.

The expert knowledge of the manufacturing process is modeled as depicted in Figure 2.


Fig. 2. Schema of the manufacturing *Knowledge Graph*.

The *Knowledge Graph* models the manufacturing process of electric vehicles as a sequence of *Lines*. A *Line* has various *Stations*. The *Stations* are also modeled as a sequence. Each *Station* implements multiple *Process Steps*. Again, the *Knowledge Graph* models the sequence in which *Process Steps* are executed. Each *Process Step* measures *Sensor Variables*. The *Knowledge Graph* models explicitly a *"hasNoImpact"* relation between *Sensor Variables* that are known to have no causal effect on each other.

In addition, a *Sensor Variable* can be a member of one subclass, which are *Root*, *Leaf*, and *Irrelevant Variable*.

Formalizing expert knowledge of the manufacturing process in such a way allows automatic reasoning over the manufacturing process. The intelligent and interactive *RCA* tool employs reasoning from the two following dimensions.
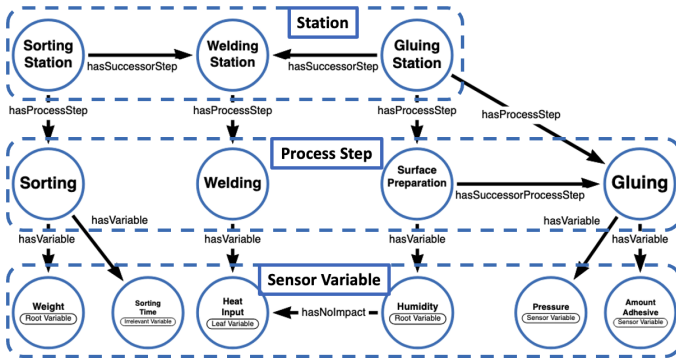

Fig. 3. Example of the manufacturing *Knowledge Graph*.

a) Temporal-spatial relations between Variables: The *Line* is modeled as a sequence of *Stations*, each with a sequence of *Process Steps*. At the bottom of this hierarchy are sensors measuring the *Sensor Variables*. Thus, the *Knowledge Graph* constructs a topology of the sensor network. It enables automatic reasoning to create a partial ordering over the *Sensor Variables*. For every *Sensor Variable*, it is known which *Sensor Variable* is measured before and after and a *Sensor Variable* measured in an earlier *Process Step* may influence a *Sensor Variable* measured in a subsequent *Process Step*, but not vice versa.

Let us consider the *Knowledge Graph* from Figure 3. The topology of the sensor network allows inferring that the *Sensor Variables "Weight"* and *"Sorting Time"* may causally impact the *Sensor Variable "Heat Input"*, but not vice versa and not the *Sensor Variables "Humidity"*, *"Amount Adhesive"*, and *"Pressure"*.

The example shows the reasoning on how a significant amount of possible *CER's* between *Sensor Variables* are excluded a priori and do not have to be considered while training the *Causal Bayesian Network*. This reduces the search space of the *Causal Bayesian Network* learning algorithm drastically, compared to a naive approach of learning the *Causal Bayesian Network* solely on tabular data (cf. Figures 4 and 5).

b) Properties of Variable subclasses: Additionally, the *Knowledge Graph* incorporates expert knowledge over the *Sensor Variables* in the form of their subclass membership. This information is introduced to the *Knowledge Graph* by the feedback of a process expert and is vital to reduce the *Causal Bayesian Network*'s learning algorithm's search space.

Root Variables cannot be impacted by any *Sensor Variables*. However, they may impact other *Sensor Variables*. For example, the *Sensor Variable "Humidity"* in Figure 3 cannot be affected by *Sensor Variables*, as the air humidity is external to the manufacturing process. However, *"Humidity"* may affect other *Sensor Variables* in the manufacturing process.

Leaf Variables are the inverse of *Root Variables*. They may be impacted by *Sensor Variables*, but cannot impact other *Sensor Variables*. The *"Heat Input" Sensor Variable* in Figure 3 is an example of a *Leaf Variable*, as it is the only *Sensor Variable* of the final *Process Step*.

Finally, there are *Irrelevant Variables*. Any manufacturing process measures a minor amount of *Sensor Variables* that do not impact any other *Sensor Variables* in the process. *Irrelevant Variables* are artifacts of unclear requirements, highly specific use cases, or legislation. The *"Sorting Time" Sensor Variable* in Figure 3 is an example, as the sorting time does not impact other *Sensor Variables*. Thus, they shall be excluded from the *RCA*.

Figures 4 and 5 show the impact of reasoning over the temporal-spatial relations between *Sensor Variables* and their subclasses on the search space of the *Causal Bayesian Network* learning algorithm. In Figure 4, every *Sensor Variable*'s fault is potentially caused by any other *Sensor Variable*, as the sensor network's topology modeled in the *Knowledge Graph* is not considered. This amounts to a total of 30 possible *CER's* between the six *Sensor Variables*. However, considering the sensor network topology, the search space of the *Causal Bayesian Network* learning algorithm is reduced to seven potential *CER's* between the *Sensor Variables* (cf. Figure

Fig. 4. The potential *CER's* between the *Sensor Variables* of the manufacturing *Knowledge Graph* (cf. Figure 3) without considering the topology of the sensor network.


Fig. 5. The potential *CER's* between the *Sensor Variables* of the manufacturing *Knowledge Graph* (cf. Figure 3) while considering the topology of the sensor network.


Fig. 6. The assumed true causal graph between the *Sensor Variables*. Its edges are contained in the relations of the potential *CER's* graph of Figure 5.

5). This example shows that automated reasoning over the expert knowledge formalized in the *Knowledge Graph* prunes the search space of the *Causal Bayesian Network* learning algorithm significantly.

B. Causal Bayesian Networks

This section introduces important concepts of *Causal Bayesian Networks*, which represent *CER's* behind data sets, and how it incorporates the *Knowledge Graph* into its training. *CER's* are dependencies, where the manipulation of one *Sensor Variable* changes the value of the other even if all remaining *Sensor Variables* are fixed. Assume the true causal graph is depicted in Figure 6. Then a change of *"Pressure"* influences *"Heat Input"* through the path *"Pressure"* → *"Amount Adhesive"* → *"Heat Input"*. For example, a higher *"Pressure"* might lead to a higher *"Amount Adhesive"*, which in turn results in a larger *"Heat Input"*. Now imagine that after an increase of *"Amount Adhesive"* we manually change all *Sensor Variables* but *"Heat Input"* back to their original level. Then, the manipulation of *"Pressure"* does not impact *"Heat Input"* anymore, as it is only directly influenced by *"Weight"* and *"Amount Adhesive"*, which are as before. Therefore, the influential path *"Pressure"* → *"Amount Adhesive"* → *"Heat Input"* is blocked. We understand these direct impacts as the *CER's*. They are pivotal for root-cause analysis, process control, and process understanding. *CER's* imply an order over the *Sensor Variables*, which we call the *causal order*. In our running example (cf. Figure 6), one possible causal order of the *Sensor Variables* is (*"Sorting Time"*, *"Weight"* *"Humidity"*, *"Pressure"*, *"Amount Adhesive"*, *"Heat Input"*).

This paper's scope is to derive the unknown causal graph. Instead of labor and cost-intensive design of experiments, we would like to leverage our sensor network along the manufacturing process to derive the causal graph in a data-driven manner. However, if we rely exclusively on data, this is challenging due to the following reasons:

1) **Confounding:** Even if we have identified a correlation, say between *"Amount Adhesive"* and *"Heat Input"*, it might be, that there is a third *Sensor Variable*, say *"Machine Operator"*, impacting both and there is no direct impact between the two. We call this third *Sen-*

sor Variable a confounder. Besides pathological cases, confounders lead to extra spurious relationships.

2) **Causal Order:** Even if we can exclude confounding it remains still unclear, whether *"Amount Adhesive"* is the cause and *"Heat Input"* the effect or it is vice versa.

3) **Complexity:** The number of potential graphs grows super-exponentially. Beyond a small number of *Sensor Variables*, it is thus impossible to evaluate all of them. This is especially worrisome if there are many cause-effect pairs.

Expert Knowledge to the Rescue: In the following, we present how expert knowledge contributes to derive the causal graph. As Section III-A mentions, the number of potential causal graphs can be drastically reduced by the sensor network's topology and expert knowledge on the manufacturing process:

1) sensor network topology provides a **partial ordering** of the *Sensor Variables* and

2) expert knowledge helps to exclude certain relationships and thus **avoids spurious relationships**.

Additionally, even a partial ordering drastically reduces the number of potential graphs. The classification of *Sensor Variables* into roots and leafs reduces not only the number of potential edges but also the number of relevant orderings. For example, as *"Weight"* and *"Humidity"* are defined as *Root Variables*, the relevant orderings are those that assign them to any of the first two positions. This reduces the number of permutations by a factor of 30. The effect of *Leaf Variables* is analogous. In the following, we present how we leverage the background information for causal graph identification.

Causal Additive Models: The derivation of the *CER's* from observed data is of central interest in many domains [24, 25]. For this task, score-based methods relying on *Structural Equation Models (SEM's)* [26] have recently become increasingly popular due to their ability to incorporate machine learning methods [18, 19, 27, 28]. The underlying assumption is that all *Sensor Variables* can be expressed as a function of inputs and a noise term, that captures unknown influences. In the example of Figure 5, *"Heat Input"* can be described by Equation 1.

$$HeatInput = f(Weight, AmountAdhesive, Noise) \quad (1)$$

We emphasize that this follows an intuitive understanding of a manufacturing process. *Process Steps* transform input material to an output, while they are impacted by machine settings and the environment. We assume that the data follows a *Causal Additive Model*, where Equation (1) is replaced by:

$$HeatInput = f_1(Humidity) + f_2(AmountAdhesive) + Noise$$

(2)

The *Noise* is normally distributed and f_1 and f_2 are non-linear. Under mild assumptions [26] on f_1 and f_2, one can derive the true causal graph from observed data. [19] proposes an approach, that identifies the graph using the following steps:

1) Find the causal ordering of the *Sensor Variables*
2) Identify the *CER's*.

Given N observations of p *Sensor Variables* by $(x_{\ell 1}, \ldots, x_{\ell p}), 1 \leq \ell \leq N$. To derive the causal ordering in step (1), search for the permutation π on $\{1, \ldots, p\}$ such that Equation 3 is minimized and $\widehat{f_{kj}^\pi}$ are learned by maximum likelihood estimation [19].

$$S(\pi) := \sum_{k=1}^{p} \log \left[\sum_{\ell=1}^{N} \left(x_{\ell k} - \sum_{\pi(j) < \pi(k)} \widehat{f_{kj}^\pi}(x_{\ell j}) \right)^2 \right], \quad (3)$$

If the *Knowledge Graph* provides a complete ordering, then we skip the first step, and the graph identification is found using regression techniques [29]. Unfortunately, the ordering is usually partial, as *Sensor Variables* measured at the same station cannot be ordered. We employ the algorithm of [8], which proposes an efficient adaption in case of expert knowledge. It limits step (1) to finding the causal ordering of the *Sensor Variables* within the *Process Steps* while mining the rest of the ordering from the *Knowledge Graph*

C. Interactive User Interface

The *Interactive User Interface* visualizes the learned *Causal Bayesian Network* (cf. Figure 7). Thereby, the *Interactive User Interface* enables no-code usage and adaption of the root cause graph by a process expert.

a) Root Cause Analysis: The workflow supported by the *Interactive User Interface* is as follows (cf. Figure 7):

The process expert searches for the *Sensor Variable* with the diagnosed fault. All possible root cause paths for the faulty *Sensor Variable* are displayed to the process expert. The process expert now sees the information on which other *Sensor Variables* the faulty *Sensor Variable* depends. That way, without further labor-intense analysis, the process expert can move to the physical manufacturing process and check all *Process Steps* related to the identified *Sensor Variables*. The process expert benefits from a highly directed root cause search in the physical manufacturing process. Thus, the interactive and intelligent root cause analysis tool minimizes and sometimes even prevents the downtime of the manufacturing process and maximizes the manufacturing process's output.

b) Expert Feedback: The *RCA* is supported by various means of interaction with the root cause graph (cf. Figure 7).

The process expert can choose the *Product* to be considered in the *RCA*. This results in accurate root cause graphs, as *Sensor Variable* measurements significantly depend on the *Product* measured. In addition, the process expert can select the *RCA* timeframe via the *Interactive User Interface*. Selecting a timeframe enables the process expert to look precisely at the period when the faults were detected. It thus customizes the *RCA* to the individual fault at hand.

However, the root cause graph might include spurious *CER's* and needs to be corrected. In this case, the expert is able to explain the *Causal Bayesian Network* where and how to do better, via adding or removing information to the *Knowledge Graph*. The *Causal Bayesian Network* may learn a *CER* between two *Sensor Variables*, which has to be rejected based on the knowledge of the process expert. The process expert can select such an edge and check "Add to Blacklist", which adds an *"hasNoImpact"* relation between the two *Sensor Variables* to the *Knowledge Graph* (cf. Figure 7 (1)). The expert feedback is considered at the next iterations of learning the *Causal Bayesian Network*, and thus is accounted for in the next *RCA's*.

In addition, spurious *CER's* might be the cause of incomplete information over subclass membership of *Sensor Variables* in the *Knowledge Graph*. Thus, the process expert can add and remove *Root*, *Leaf*, and *Irrelevant Variable* subclass membership from *Sensor Variables* (cf. Figure 7 (2)). This feedback is also considered in the next iteration of the *Causal Bayesian Network*. Assuming that the expert feedback is correct, the root cause graph converges each iteration closer to the true root cause graph.

IV. EVALUATION

The evaluation is conducted on a real-world electric vehicle manufacturing process. One *Line* is taken as an example to show the *RCA* with the interactive and intelligent tool.

A. Data

The data is two-folded. There is the *Knowledge Graph* and the *Data Layer* as a source of data.

The *Knowledge Graph* is implemented in *Neo4J* [30]. This allows querying all knowledge via the Cypher API. Cypher is a deductive reasoning query language [30]. The real-world *Knowledge Graph* holds a total of 100,015 nodes and 417,944 relationships. This amount of expert knowledge constitutes a large-scale knowledge graph. For the *Line* of this evaluation, there are 53 *Stations* with 96 *Process Steps* and 1683 *Sensor Variables* and a total of 2143 relations modeled.

To learn the *Causal Bayesian Network* we prepare the data by assigning the *Sensor Variables* to individual output products. Then, we preprocess the data by removing columns with only one value. Further, we remove one part of a column pair, if they have a correlation above 0.95 to avoid collinearity. Afterwards, we remove columns and rows, where more than 50% of the values are missing. As the rate of missing values
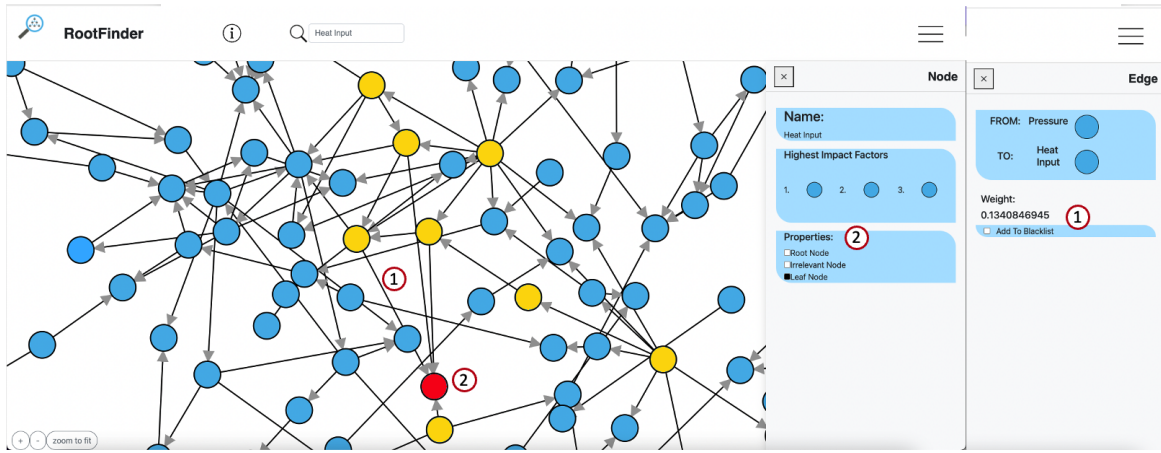
Fig. 7. The user interface of the interactive and intelligent root cause analysis tool. The fault is marked red. *Sensor Variables* part of the learned root cause path are marked yellow. The numbers in the red circle highlight the interactions.


Fig. 8. Decrease of learning time and *CER's* with an increase of data from the *Knowledge Graph*.

is low, we simply impute the column's mean for missing values. The presented approach is generic and can be applied dynamically to different products and time windows. For this section, we consider data of one day and for one specific product type.

B. Experiments

The contribution of the paper is to give a real-world example of how to interactively improve the performance of *Causal Bayesian Networks* for *RCA* with large-scale modelled expert knowledge in the form of a *Knowledge Graph*. The contribution is shown by three experiments.

1) Decrease in learning time of the Causal Bayesian Network by including the Knowledge Graph: As argued above, one major motivation for including expert knowledge is a decrease in the learning time of the *Causal Bayesian Network* algorithm. Figure 8 shows the average decrease in learning time with an increase in expert knowledge. The experiment was conducted five times. For each run-through, randomly 25%, 50%, 75%, and 100% of data from the *Knowledge Graph* were selected for training the *Causal Bayesian Network*. A decline by 79.0% in training time can be observed, with an increase in data from the Knowledge Graph. Thus,

expert knowledge optimizes and enables the training of *Causal Bayesian Network* in large-scale manufacturing processes.

2) Decline of spurious CER's in the Causal Bayesian Network by including the Knowledge Graph: As argued above, *Causal Bayesian Networks* for manufacturing processes tend to include spurious *CER's*. Figure 8 shows the average decline in spurious CER's with an increase in expert knowledge. The experiment was conducted as described in IV-B1. A steep decline in the amount of learned *CER's* is observable. The result indicates that data from the *Knowledge Graph* prunes 60.6% of spurious *CER's*. This leads to a more accurate root cause graph and thus increases the performance of the root cause analysis tool. The following experiment shall underline the latter claim.

3) Expert feedback improves the Causal Bayesian Network: To evaluate the usefulness of the interactive component, we extend the evaluation proposed by [8]. A learned *Causal Bayesian Network* is compared to a partially known root cause graph, using an *adapted Structural Hamming Distance (aSHD)*, which describes the deviation of the learned *Causal Bayesian Network* from the partially known root cause graph [8]. A lower number indicates a better result. The learning algorithm is applied on 100 random samples of eight *Sensor Variables*, using the *temporal-spatial ordering*. For each learned *Causal Bayesian Network*, the *aSHD* is calculated. The mean and standard deviation of the *aSHD* are 0.47 and 0.74, respectively. If an expert now adds one *hasNoImpact* relation to the *Knowledge Graph*, the mean and standard deviation of the *aSHD* go down to 0.44 and 0.70. This corresponds to an improvement of the *aSHD* by almost 6% and the learning algorithm is stabilized, as the standard deviation is lower. The results illustrate the benefits of including expert knowledge and exemplify one iteration step of the interactive and intelligent root cause analysis tool.

V. CONCLUSION

This work proposes an interactive and intelligent root cause analysis tool for the manufacturing process of electric vehi-

cles. It shows how to model detailed expert knowledge of a real-world manufacturing process in a large-scale *Knowledge Graph*, which is used for automatic reasoning over potential root causes. It is described how the reasoning is used to prune the search space of a *Causal Bayesian Network* learning algorithm. In addition, this work shows how to include expert feedback in the learning of the *Causal Bayesian Network* via the *Interactive User Interface* and the *Knowledge Graph* to identify and exclude spurious *CER's*. The evaluation of the interactive and intelligent root cause analysis tool on a real-world manufacturing process of electric vehicles proves the tool's feasibility.

Machine learning makes *RCA* more efficient. However, it will only succeed if we rely on data and expert knowledge. Let us put the human back in the loop.

REFERENCES

[1] C. Wehner, F. Powlesland, B. Altakrouri, and U. Schmid, "Explainable online lane change predictions on a digital twin with a layer normalized lstm and layer-wise relevance propagation," in *Advances and Trends in Artificial Intelligence. Theory and Practices in Artificial Intelligence*, H. Fujita, P. Fournier-Viger, M. Ali, and Y. Wang, Eds. Cham: Springer International Publishing, 2022, pp. 621–632.

[2] E. e Oliveira, V. L. Miguéis, and J. L. Borges, "Automatic root cause analysis in manufacturing: an overview and conceptualization," *Journal of Intelligent Manufacturing*, 2022. [Online]. Available: https://doi.org/10.1007/s10845-022-01914-3

[3] O. Serrat, *The Five Whys Technique*. Singapore: Springer Singapore, 2017, pp. 307–310.

[4] K. M. Tay and C. P. Lim, "On the use of fuzzy inference techniques in assessment models: part ii: industrial applications," *Fuzzy Optimization and Decision Making*, vol. 7, pp. 283–302, 2008. [Online]. Available: https://doi.org/10.1007/s10700-008-9037-y

[5] E. Ruijters and M. Stoelinga, "Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools," *Computer Science Review*, vol. 15-16, pp. 29–62, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1574013715000027

[6] S. Pradhan, R. Singh, K. Kachru, and S. Narasimhamurthy, "A bayesian network based approach for root-cause-analysis in manufacturing process," in *2007 International Conference on Computational Intelligence and Security (CIS 2007)*, 2007, pp. 10–14.

[7] M. Kirchhof, K. Haas, T. Kornas, S. Thiede, M. Hirz, and C. Hermann, "Root cause analysis in lithium-ion battery production with fmea-based large-scale bayesian network," 06 2020.

[8] M. Kertel, S. Harmeling, and M. Pauly, "Learning causal graphs in manufacturing domains using structural equation models," 2022. [Online]. Available: https://arxiv.org/abs/2210.14573

[9] ISO/IEC, "Risk management — Risk assessment techniques," International Organization for Standardization, Geneva, CH, Standard, Jun. 2019.

[10] Q. Ma, H. Li, and A. Thorstenson, "A big data-driven root cause analysis system: Application of machine learning in quality problem solving," *Computers and Industrial Engineering*, vol. 160, p. 107580, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360835221004848

[11] A. Alaeddini and I. Dogan, "Using bayesian networks for root cause analysis in statistical process control," *Expert Systems with Applications*, vol. 38, no. 9, pp. 11230–11243, 2011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417411003952

[12] A. Lokrantz, E. Gustavsson, and M. Jirstrand, "Root cause analysis of failures and quality deviations in manufacturing using machine learning," *Procedia CIRP*, vol. 72, pp. 1057–1062, 2018, 51st CIRP Conference on Manufacturing Systems. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2212827118303895

[13] K. Marazopoulou, R. Ghosh, P. Lade, and D. Jensen, "Causal discovery for manufacturing domains," 05 2016.

[14] J. Li and J. Shi, "Knowledge discovery from observational data for process control using causal bayesian networks," *IIE Transactions*, vol. 39, no. 6, pp. 681–690, 2007. [Online]. Available: https://doi.org/10.1080/07408170600899532

[15] L. Abele, M. Anic, T. Gutmann, J. Folmer, M. Kleinsteuber, and B. Vogel-Heuser, "Combining knowledge modeling and machine learning for alarm root cause analysis," *IFAC Proceedings Volumes*, vol. 46, no. 9, pp. 1843–1848, 2013, 7th IFAC Conference on Manufacturing Modelling, Management, and Control. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667016345633

[16] G. Weidl, A. Madsen, and S. Israelson, "Applications of object-oriented bayesian networks for condition monitoring, root cause analysis and decision support on operation of complex continuous processes," *Computers and Chemical Engineering*, vol. 29, no. 9, pp. 1996–2009, 2005. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S009813540500133X

[17] D. Zhang, L. Wang, Q. Hong, and K. Zhang, "Research on fault diagnosis of steam turbine based on bayesian network," *Journal of Physics: Conference Series*, vol. 1754, p. 012136, 02 2021.

[18] X. Zheng, C. Dan, B. Aragam, P. Ravikumar, and E. Xing, "Learning sparse nonparametric dags," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 3414–3425.

[19] P. Bühlmann, J. Peters, and J. Ernest, "CAM: Causal additive models, high-dimensional order search and penalized regression," *The Annals of Statistics*, vol. 42, no. 6, pp. 2526 – 2556, 2014. [Online]. Available: https://doi.org/10.1214/14-AOS1260

[20] G. Menegozzo, D. Dall'Alba, and P. Fiorini, "Cipcad-bench: Continuous industrial process datasets for benchmarking causal discovery methods," in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2022, pp. 2124–2131.

[21] A. Hogan, E. Blomqvist, M. Cochez, C. D'amato, G. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A.-C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, and A. Zimmermann, "Knowledge graphs," *ACM Computing Surveys*, vol. 54, pp. 1–37, 2021.

[22] F. Harary, R. Z. R. Z. Norman, and D. Cartwright, *Structural models: an introduction to the theory of directed graphs*. Wiley, 1965.

[23] J. Gallian, "A dynamic survey of graph labeling," *Electron. J. Combin., Dynamic Surveys*, vol. 19, 11 2000.

[24] A. Wille, P. Zimmermann, E. Vranová, A. Fürholz, O. Laule, S. Bleuler, L. Hennig, A. Prelić, P. von Rohr, L. Thiele *et al.*, "Sparse graphical gaussian modeling of the isoprenoid gene network in arabidopsis thaliana," *Genome biology*, vol. 5, no. 11, pp. 1–13, 2004.

[25] G. N. Saxe, S. Ma, L. J. Morales, I. R. Galatzer-Levy, C. Aliferis, and C. R. Marmar, "Computational causal discovery for post-traumatic stress in police officers," *Translational psychiatry*, vol. 10, no. 1, pp. 1–12, 2020.

[26] J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference: Foundations and Learning Algorithms*. Cambridge, MA, USA: MIT Press, 2017.

[27] S. Shimizu, T. Inazumi, Y. Sogawa, A. Hyvärinen, Y. Kawahara, T. Washio, P. O. Hoyer, and K. Bollen, "Directlingam: A direct method for learning a linear non-gaussian structural equation model," *J. Mach. Learn. Res.*, vol. 12, no. 0, p. 1225–1248, jul 2011.

[28] Y. Yu, J. Chen, T. Gao, and M. Yu, "DAG-GNN: DAG structure learning with graph neural networks," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 7154–7163. [Online]. Available: https://proceedings.mlr.press/v97/yu19a.html

[29] A. Shojaie and G. Michailidis, "Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs," *Biometrika*, vol. 97, no. 3, pp. 519–538, 07 2010. [Online]. Available: https://doi.org/10.1093/biomet/asq038

[30] M. Lal, *Neo4j Graph Data Modeling*. Packt Publishing, 2015.