



HAL
open science

Diagnosis automation using similarity analysis: application to industrial systems

Ivan Orefice, Wissam Mallouli, Ana R Cavalli, Filip Sebek, Alberto Lizarduy

► **To cite this version:**

Ivan Orefice, Wissam Mallouli, Ana R Cavalli, Filip Sebek, Alberto Lizarduy. Diagnosis automation using similarity analysis: application to industrial systems. 19th International Conference on Software Technologies(ICSOFT), Jul 2024, Dijon, France. pp.331-338, 10.5220/0012719200003753 . hal-04902424

HAL Id: hal-04902424

<https://hal.science/hal-04902424v1>

Submitted on 20 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0
International License

Diagnosis Automation using Similarity Analysis: Application to Industrial Systems

Ivan Orefice¹, Wissam Mallouli¹, Ana R. Cavalli¹, Filip Sebek² and Alberto Lizarduy³

¹Montimage EURL, Paris, France

²R&D ABB Marine and Ports, Västerås, Sweden

³FAGOR Arrasate S.Coop, Spain

name.surname@montimage.com, filip.sebek@se.abb.com, a.lizarduy@fagorarrasate.com

Keywords: Incident diagnosis, Root cause analysis, Automation, Load Position System, Manufacturing Machinery, Security

Abstract: The paper introduces the MMT-RCA framework, an automated incident diagnosis system crucial for maintaining security and reliability in complex systems such as ABB's Load Position Sensor (LPS) and FAGOR's remote manufacturing machinery access. Traditional incident response methods often involve time-consuming and error-prone manual analysis, hindered by limited human expertise. MMT-RCA addresses this challenge by leveraging similarity analysis techniques. It utilizes historical incident data to create a comprehensive repository, capturing characteristics and outcomes of past incidents. By employing sophisticated algorithms, the MMT-RCA framework identifies patterns and correlations among incidents, facilitating the swift identification of similar problems and their root causes. To validate its efficacy, the framework underwent real-world experiments with industrial data from both companies. The results demonstrate the framework's ability to accurately diagnose incidents and identify root causes.

1 INTRODUCTION

Ensuring the security and safety of critical industrial system (Kalam, 2021) is essential to prevent incidents that can result in significant damage, financial losses, and potential harm to human lives. In fact, these systems face various security and safety challenges that can lead to incidents, making incident diagnosis a crucial aspect of their overall protection (Hemsley and Fisher, 2018). One of the key challenges is the complex and evolving nature of threats. Attackers are continuously developing new techniques to exploit vulnerabilities in industrial systems (Kallel et al., 2021), necessitating proactive monitoring, threat intelligence, and robust incident detection capabilities (Salazar et al., 2022). Effective incident diagnosis enables the identification of security breaches and helps determine the root causes of these incidents, allowing organizations to implement appropriate countermeasures and strengthen their security posture.

Another challenge is the detection and diagnosis of safety incidents within critical industrial systems (He et al., 2022). Safety incidents can arise from equipment failures, human errors, procedural violations, or environmental factors. Identifying the root

causes of safety incidents is crucial to prevent their recurrence and improve overall safety measures.

Furthermore, the integration of security and safety incident diagnosis is essential for a comprehensive approach to protect critical industrial systems (Kirkpatrick, 2019). Security and safety incidents are often intertwined and can have cascading effects on each other. For example, a cyber attack targeting an industrial control system can lead to safety hazards if critical safety controls are compromised. Conversely, a safety incident, such as an equipment malfunction, can create vulnerabilities that can be exploited by malicious actors. By integrating security and safety incident diagnosis, organizations can gain a holistic understanding of incidents, their interdependencies, and the underlying causes, enabling them to implement appropriate measures (Lanotte et al., 2023) to address both security and safety concerns effectively.

In this paper, we propose an automated diagnosis framework called MMT-RCA a new feature of Montimage Monitoring Tool¹ that aims to leverage similarity analysis techniques to enhance the identifica-

¹<https://www.montimage.com/products>

tion of the root causes of incidents within critical industrial systems. Traditional manual incident analysis processes can be time-consuming and subjective, relying heavily on the expertise and experience of human operators. By incorporating similarity analysis techniques (Black et al., 2019), the framework aims to expedite the incident diagnosis process by identifying patterns and correlations among incidents.

Similarity analysis involves comparing the characteristics and outcomes of incidents to identify similarities and associations that can help determine the underlying root causes. This can be achieved through clustering algorithms that group similar incidents together based on shared or close metric values. By utilizing historical incident data and real-time monitoring data, the framework can continuously learn and improve its diagnosis capabilities.

The integration of similarity analysis techniques in the automated diagnosis framework offers several advantages. Firstly, it can significantly reduce the time required for incident diagnosis, enabling prompt response and mitigation actions. Secondly, it helps minimize the subjectivity and biases associated with manual analysis, providing more objective and data-driven results. Moreover, by identifying the root causes of incidents, organizations can implement targeted remediation strategies, enhance system resilience, and proactively address potential vulnerabilities.

The MMT-RCA framework has been applied to both LPS provided by ABB company² and remote access to manufacturing machinery, a system used inside FAGOR company³, and has yielded interesting results. By leveraging historical incident data and real-time monitoring data from the LPS and the manufacturing system, the framework successfully identified patterns and correlations among incidents, enabling the prompt identification of root causes. The application of the framework to both datasets significantly reduced the time required for incident diagnosis, improving incident response efficiency and overall system uptime. These interesting results demonstrate the effectiveness of the framework in automating the root cause analysis of incidents within these two industrial complex contexts.

The paper is organised as follows. Section 2 presents related work to root cause analysis of incident in industrial systems. Section 3 presents the proposed framework architecture and implementation details. Furthermore, Section 4 presents its application results to LPS system and FAGOR use case. Section 5 concludes the paper and present future work.

²<https://global.abb/group/en>

³<https://fagorarrasate.com>

2 RELATED WORK

There is a significant body of scientific work related to protecting critical industrial systems (Salazar et al., 2022) (Kalam, 2021) (Hemsley and Fisher, 2018) and some of them focused on incidents diagnosis and root cause analysis. The basic concepts and terminologies related to Root Cause Analysis (RCA) are defined in the literature.

The authors of (Kiermeier and Feld, 2018) discuss the challenges of performing root cause analysis in self-organizing industrial systems (SOIS). These systems adapt their behaviour dynamically, making it difficult to establish explicit connections and identify root causes. The authors present a taxonomy of possible root causes in SOIS, with a focus on error sources arising from the online decision-making process. They propose backtracking approaches, distinguishing between automatable and non-automatable procedures. For cases where automatable evaluation is not feasible due to the state space explosion, a visual analytics solution is proposed. The paper also includes a proof of concept for an expert-based assessment, demonstrating the necessary functions for tracing back anomalies in SOIS.

Besides, (Wang et al., 2021) introduces Groot, an event-graph-based approach for root cause analysis in large-scale distributed systems. Groot addresses the challenges posed by microservice architecture, including operational complexities, system scale, and monitoring. It constructs a real-time causality graph based on events, incorporating various metrics, logs, and activities for analysis. Groot can be customized with user-defined events and domain-specific rules, allowing integration of domain knowledge from site reliability engineering (SRE) engineers. The paper demonstrates the usability and effectiveness of Groot in industrial settings, highlighting its practical application and lessons learned. However, the Groot tool presents some limitations associated with its scalability, adaptability to different system architectures, or potential challenges in integrating real-time data from distributed environments.

These scientific works contribute to the understanding and development of root cause analysis methodologies in industrial systems, providing valuable insights and practical guidance for incident investigation, prevention, and overall system improvement. In our paper, we propose a generic root cause analysis solution called MMT-RCA that can extended to different sectors and can easily be integrated in operational environments due to its modularity.

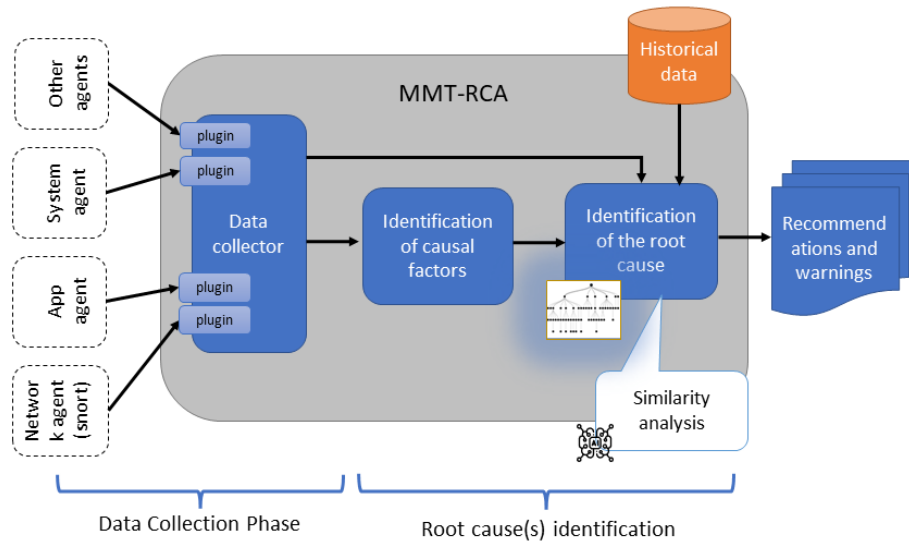


Figure 1: The automated diagnosis framework architecture.

3 THE AUTOMATED DIAGNOSIS FRAMEWORK : MMT-RCA

In this section the complete framework of MMT-RCA will be explained and showed in details.

3.1 The Automated Diagnosis Framework

Figure 1 illustrates the high-level architecture of the automated diagnosis framework. The tool utilizes a **data collector** that gathers information from various sources within the industrial critical system, such as network, application, system and hardware, using dedicated monitoring agents. The data collector employs a plugin architecture to support the extraction of attribute values relevant for identifying the origin of incidents. Machine learning algorithms are applied to select the most significant attributes called **causal factors**, enhancing analysis accuracy and reducing computational resource requirements.

The data collector can be either provided by the system itself or deployed as an agent, captures network traffic and reads and extracts logs in different formats like JSON, CSV or even binary. Relevant metrics are selected, considering performance indicators and specific case studies.

Historical data consisting of labeled events and their associated attribute values is used for learning purposes. By using **similarity** learning, particularly Ranking Similarity Learning, the tool compares new system states with known undesirable states to rec-

ognize root causes. The tool **recommends** relevant countermeasures based on known mitigation strategies.

The RCA tool operates in two phases: knowledge acquisition, also called learning phase, and monitoring phase. The learning phase involves building a historical database of known problems and incidents, while the monitoring phase continuously analyzes real-time system data, queries the historical database, and suggests potential root causes. Root cause identification relies on similarity analysis, treating system states as vectors in a multidimensional space.

Recommendations and visualization of similarity scores, known incidents, and root causes assist system administrators in anticipating system evolution, detecting faults, and taking appropriate mitigation actions.

3.2 Implementation details

As described in the previous section, the RCA is based on two phases: knowledge acquisition and monitoring stage. Both of them require to define a set of metrics or features describing the system status and its related incidents. This set of features are obtained thanks to the use of 4 in-house modules of MMT tool, their combination constitute the MMT-RCA solution:

- MMT-Extract: this module allows parsing the input traces of the MMT-RCA. These inputs can have different formats and sources. That's why MMT-Extract has a plugin architecture to deal with heterogeneous inputs. Today MMT-Extract

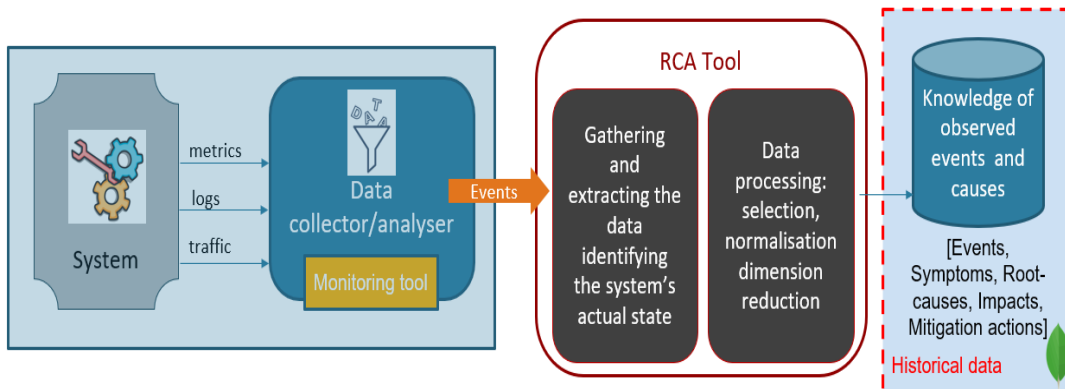


Figure 2: The Knowledge acquisition phase.

has more than 650 plugins and can deal with pre-captured logs/datasets or live streaming logs. The data extracted can be coming from the raw data (a specific attribute in the log) or computed data by correlating several events. Example: computing packet loss rate in a session, or speed from 2 positions of a vehicle etc.

- MMT-Security: this module allows to perform Complex Event Processing (CEP) to detect security incidents and attacks. To do so, it aggregates and applies a logic to the extracted variables.
- MMT-Similarity: Based on the historical data (i.e. list known metrics values denoting each security incident), a distance is computed for newly monitored datasets. This allows to compute the similarity of events and identify the root causes of problem with a specific accuracy. This identification is based on similarity analysis algorithms that can

be chosen in the configuration file.

- MMT-Operator: allows the display of the results of analysis and classify the possible causes and provide recommendations and warnings.

The learning phase is depicted in Figure 2. During this phase, labelled incidents need to be used. They can be either old detected incidents or simulated ones. The monitoring phase illustrated in Figure 3 relies on precaptured logs or can be done at runtime where incidents are detected and an automatic diagnosis is performed. The MMT-RCA software solution is implemented as open source solution in C, Python and Javascript and its modules are available at <https://github.com/Montimage>.

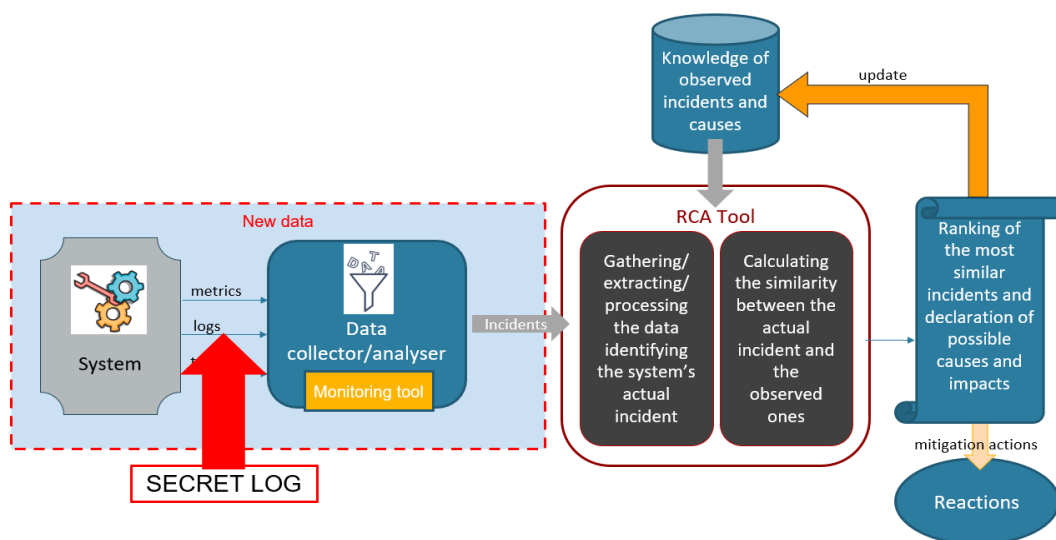


Figure 3: The monitoring and diagnosis phase.

4 APPLICATION TO IT SYSTEMS

4.1 ABB Case Study

ABB Load Position Sensor (LPS) is a camera-based tracking system that determines a hanging crane load position with help of attached LED markers on the load. As can be seen in Figure 4, the system consists of a camera on top of a crane that keeps tracks of LED markers on a load that is in motion. In addition, there is single Programmable Logic Controller (PLC) that analyses the camera images. The sensor acts as feedback to the automatic control of the crane.

However, it has happened that noise detected by the camera is tagged as markers by the PLC. This situation might lead to incorrect movements of the crane and compromise the safety of the system.




Figure 4: The Load Position Sensor (Salazar et al., 2022)

4.1.1 Potential Causes of Incidents

In this experiment, we focus only on 3 types of challenges that could cause an incident:

- **Problem 1:** Internal bugs in camera firmware that could cause faulty output;
- **Problem 2:** Markers are split by a hanging cable between the camera and the marker;
- **Problem 3:** Random noises from the environment in the collected data.

Three datasets of each problem have been gathered by ABB. Secret logs (without labelling) has been generated for testing purposes.

4.1.2 Data Extraction and Knowledge Acquisition

A first attempt of using raw features in the datasets has been ineffective as the information inside the log was not enough to describe each of the problems. Therefore, new features were derived by using MMT-Extract applying computations and aggregation of many rows of the log. These features are listed in Table 1. A core step in the workflow of the tool is the attributes selection. In order to do so, Recursive Feature Elimination (RFE) (Choi et al., 2011) and clusterization (Ezugwu et al., 2022) were applied to the feature set and was computed a correlation matrix, whereby it was selected a subset of metrics that are more relevant to the incidents and enable higher performance: metrics ids (1), (2), (3), (4), (5), (6), (18), (20).

4.1.3 Learning Phase

During the learning phase, the tool takes in input the subset of selected features for each labelled dataset. From here on, normalisation is implemented: it permits to have all the feature in the range $[0,1]$ but also to ease the comparison among the states of the incidents. After the normalisation, parameters of Gaussian distribution per problem are computed (mean, standard deviation, min and max).

4.1.4 Evaluation results

MMT-RCA recognised correctly most of the examples (with a similarity score above 80%) while the tool had little trouble recognising the log without any faults, but these examples still had a acceptable similarity score (above 70%). The overall accuracy of the tool was 95.1% while the amount of examples that had not been correctly recognised was only 14 on a total of 292 examples. Finally, also the results obtained with the secret log used as input of MMT-RCA are satisfying; the log was composed of a scarce number of examples of normal behaviour and a large amount of examples related to the first problem. RCA classified successfully the examples to the first problem while the ones not containing an issue were associated to problem 0 (normal behaviour). It is important to underline that to obtain these results only one example of each incident log was enough to learn and build the knowledge database. This makes the MMT-RCA unique and interesting for the industrial context.

Table 1: List of attributes extracted with MMT-Extract for ABB use case

Ref	Attribute	Description
(1)	NSE_{th}^i	Number of entries markers were frozen in the same position (x,y)
(2)	NST_{th}^i	Number of times markers were frozen in the same position (x,y)
(3)	MDAM	Minimum distance between all the entries for same marker
(4)	ADAM	Average distance between all the entries for same marker
(5)	ADT_TNM	Average ΔT the trolley was not moving with speed >0.5
(6)	TTNM	Number of times the the trolley was not moving with speed >0.5
(7)	ADT_HNM	Average ΔT the hoist was not moving with speed >0.5
(8)	THNM	Number of times the the hoist was not moving with speed >0.5
(9)	AvgTrolleyPos	Average position of the trolley in the log (x,y)
(10)	AvgHoistPos	Average position of the hoist in the log (x,y)
(11)	$AvgMarker_{th}^i$	Average coordinates x and y per each marker
(12)	AvgMhSpeed	Average speed of Main Hoist
(13)	AvgGaSpeed	Average speed of Gantry
(14)	AvgTrSpeed	Average speed of Trolley
(15)	AvgMhAcceleration	Average acceleration of Main Hoist
(16)	AvgGaAcceleration	Average acceleration of Gantry
(17)	AvgTrAcceleration	Average acceleration of Trolley
(18)	TotalDistanceMarker	Total euclidean distance per marker in a window of time
(19)	AvgNumberOfMarkers	Average number of markers considered valid by PLC
(20)	NotAligned	Number of times the markers were not aligned (angle <170)

4.2 Fagor use case

To use their manufacturing equipment, Fagor provides a Remote Desktop Protocol (RDP) connection to remote users. Despite the robust RDP framework, challenges may arise, leading to potential disruptions in reaching the destination host. Monitoring is crucial to identify and address these challenges promptly.

Additionally, a VPN tunnel is used to transmit network traffic between the Industrial PC and the Tele-service host. In order to manage the traffic and guarantee the security and privacy of the operations, the point-to-point VPN has a remote firewall that may block traffic between the Industrial PC and the endpoint. Figure 5 illustrates the described architecture and three different incident scenarios has been simulated and pcap files provided for the learning phase.

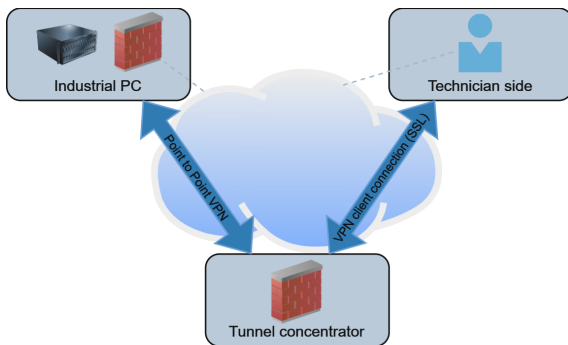


Figure 5: FAGOR system topology

- a trace describing the situation in which the Industrial PC running, but RDP protocol was not activated (**Problem 1**);
- a problem known as **Problem 2** where communication with the industrial PC was blocked by the firewall;
- a pcap in which the firewall could not be contacted (**Problem 3**).

4.2.1 Data Extraction and Knowledge Acquisition

Several network features have been computed for this analysis and are listed in Table 2. The feature selection process was still required despite the small number of metrics that were recovered. Features (1) and (2) were finally discarded.

4.2.2 Learning phase

Similar to how it was done for the ABB use case, RCA was given the datasets as input so that it could learn the states of every problem. After normalising the values, the states were stored in a MongoDB collection, while the Gaussian distribution parameters were kept in a separate collection. At the end of this step there were four states in the historical data, one for each scenario.

Table 2: List of attributes extracted and sent to RCA for FAGOR use case

Ref	Attribute	Description
(1)	RDP	Percentage of RDP packets in the entire trace
(2)	TPKT	Percentage of TPKT packets in the entire trace
(3)	ICMP	Percentage of ICMP packets in the entire trace
(4)	ARP	Percentage of ARP packets in the entire trace
(5)	MDNS	Percentage of MDNS packets in the entire trace
(6)	LENGTH	Length of the trace (number of packets)
(7)	NO-RDP-TPKT	Absence of both RDP and TPKT packets in the entire trace
(8)	NO-RDP-TPKT-ICMP	Absence of both RDP and TPKT and ICMP packets in the entire trace

4.2.3 Evaluation Results

The majority of issues were acknowledged by the tool as problems (with a similarity score in this case above 75%, which was considered as a good threshold). Only a few instances were mistakenly classified as belonging to a separate problem, demonstrating the tool’s reliability: in fact, the tool’s overall accuracy was 91.42%, with only 35 out of a total of 408 samples failing to acquire the correct identification. Finally, the results from injecting unlabeled trace data into MMT-RCA are encouraging; the tool was able to classify each trace to the appropriate incident with a similarity score of $\approx 75\%$, with the exception of the normal incident, which has a lower similarity rate but is still correctly identified.

4.3 Performance evaluation

To assess the performance of MMT-RCA, multiple experiments were carried out regarding scalability and execution time. The amount of root cause scenarios was increased up to 12 which represent a high number of potential causes in a realistic case study. Each different scenario was created taking into account a likely possible scenario by applying a perturbation to a subset of the features. For each case, the tool has been applied to the examples and were collected quantitative metrics such as the accuracy and execution time.

The experiments conducted can be divided into two main groups:

- the first group is related to those experiments which did not include data augmentation and had a number of examples used during the learning phase equal to $2 \cdot |\text{root} - \text{causes}|$ according to the number of problems used in that case;
- in the experiments of the second group instead was used data augmentation, so that the number of examples for every case was ≈ 400 .

This division allowed to understand how MMT-RCA performed with a small or a large amount of examples used during the learning phase. For all the experiments of both groups it has been applied feature selection, so as to get the highest possible value of accuracy: the number of features ranges from 7 to 9.

Firstly, the experiments done to evaluate MMT-RCA showed it can reach high levels of accuracy, whether many examples are present with which to learn the states of incidents or not. As anybody would expect, the higher is the number of examples during the learning phase, the more precise and definite is the state, therefore the more accurate is the tool; on the other hand, MMT-RCA still reaches high values of accuracy when it encounters only few examples, that is the formidable outcome of the instrument: just few examples are needed (in our experiments each problem was described by only two instances) to create a state which is in any case reliable.

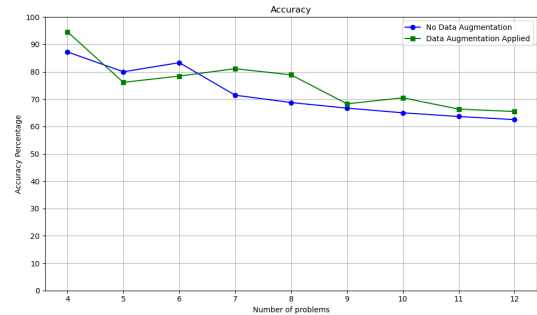


Figure 6: Accuracy of MMT-RCA

Secondly, the experiments showed a real good performance of the tool in terms of speed, since increasing the number of problems did not lead to an exponential increase of execution time for both groups; instead, the execution time grows linearly, which is a outstanding result: increasing the number of problems did not correspond to an interminable wait.

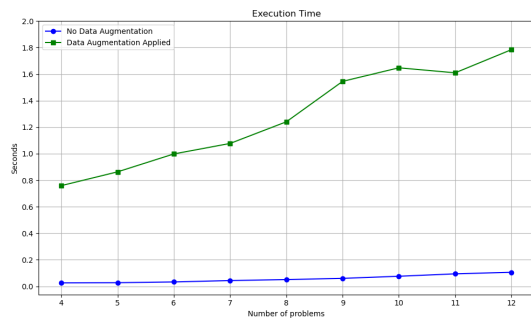


Figure 7: Execution Time of MMT-RCA

5 CONCLUSIONS

In conclusion, the proposed approach of leveraging historical incident data and employing similarity analysis algorithms in the MMT-RCA framework has shown promising results in incident management and root cause identification. By building a comprehensive incident repository and analyzing patterns and correlations among issues, the framework effectively identifies similar incidents and their underlying causes. The validation experiments conducted using real-world incident data from industrial settings provided by both ABB and FAGOR have demonstrated the framework's accuracy in diagnosing incidents and significantly reducing the time needed for manual analysis. The automation of the diagnosis process not only improves incident response time but also enables proactive maintenance, leading to increased system uptime and enhanced operational efficiency. Overall, the implementation of this approach holds great potential for enhancing incident management practices and optimizing the performance of industrial systems.

ACKNOWLEDGMENT

This work is partially supported by the European Union's Horizon Europe research and innovation program under grant agreements No 957212 (VERIDEVOPS) and No 101070455 (DYNABIC). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union.

REFERENCES

Black, P., Gondal, I., Vamplew, P., and Lakhota, A. (2019). Evolved similarity techniques in malware analysis. In *18th IEEE International Conference On Trust, Security And Privacy In Computing And Communications /*

13th IEEE International Conference On Big Data Science And Engineering, TrustCom/BigDataSE 2019, Rotorua, New Zealand, August 5-8, 2019, pages 404–410. IEEE.

- Choi, H., Yeo, D., Kwon, S., and Kim, Y. (2011). Gene selection and prediction for cancer classification using support vector machines with a reject option. *Comput. Stat. Data Anal.*, 55(5):1897–1908.
- Ezugwu, A. E., Ikotun, A. M., Oyelade, O. N., Abualigah, L. M., Agushaka, J. O., Eke, C. I., and Akinyelu, A. A. (2022). A comprehensive survey of clustering algorithms: State-of-the-art machine learning applications, taxonomy, challenges, and future research prospects. *Eng. Appl. Artif. Intell.*, 110:104743.
- He, Z., Chen, P., Luo, Y., Yan, Q., Chen, H., Yu, G., and Li, F. (2022). Graph based incident extraction and diagnosis in large-scale online systems. In *37th IEEE/ACM International Conference on Automated Software Engineering, ASE 2022, Rochester, MI, USA, October 10-14, 2022*, pages 48:1–48:13. ACM.
- Hemsley, K. and Fisher, R. E. (2018). A history of cyber incidents and threats involving industrial control systems. In Staggs, J. and Sheno, S., editors, *Critical Infrastructure Protection XII - 12th IFIP WG 11.10 International Conference, ICCIP 2018, Arlington, VA, USA, March 12-14, 2018, Revised Selected Papers*, volume 542 of *IFIP Advances in Information and Communication Technology*, pages 215–242. Springer.
- Kalam, A. A. E. (2021). Securing SCADA and critical industrial systems: From needs to security mechanisms. *Int. J. Crit. Infrastructure Prot.*, 32:100394.
- Kallel, S., Cuppens, F., Boulahia-Cuppens, N., Kacem, A. H., and Othmane, L. B. (2021). Special issue on risk and security of smart systems. *J. Inf. Secur. Appl.*, 61:102925.
- Kiermeier, M. and Feld, S. (2018). Visual analytics for root cause analysis in self-organizing industrial systems. In *16th IEEE International Conference on Industrial Informatics, INDIN 2018, Porto, Portugal, July 18-20, 2018*, pages 315–320. IEEE.
- Kirkpatrick, K. (2019). Protecting industrial control systems. *Commun. ACM*, 62(10):14–16.
- Lanotte, R., Merro, M., and Munteanu, A. (2023). Industrial control systems security via runtime enforcement. *ACM Trans. Priv. Secur.*, 26(1):4:1–4:41.
- Salazar, Z., Cavalli, A. R., Mallouli, W., Sebek, F., Zaïdi, F., and Rakoczy, M. E. (2022). Monitoring approaches for security and safety analysis: Application to a load position system. In *15th IEEE International Conference on Software Testing, Verification and Validation Workshops ICST Workshops 2022, Valencia, Spain, April 4-13, 2022*, pages 40–48. IEEE.
- Wang, H., Wu, Z., Jiang, H., Huang, Y., Wang, J., Köprü, S., and Xie, T. (2021). Groot: An event-graph-based approach for root cause analysis in industrial settings. In *36th IEEE/ACM International Conference on Automated Software Engineering, ASE 2021, Melbourne, Australia, November 15-19, 2021*, pages 419–429. IEEE.