Hybrid Modelling for Anomaly Detection in Industrial Control Systems

Vincent Boerjan, Stefano Schivo and Clara Maathuis
Open University of the Netherlands, Heerlen, The Netherlands

v.boerjan@proton.me
stefano.schivo@ou.nl
clara.maathuis@ou.nl

Abstract: This research addresses the challenge of anomaly detection in Industrial Control Systems (ICS), recognizing the increasing importance of cyber security in these environments due to recent incidents and evolving technical and regulatory frameworks and mechanisms introduced. It does that by proposing a comprehensive hybrid modelling approach to anomaly detection that bridges the gap between theoretical research and practical applications in real-world industrial settings. Specifically, this methodology focuses on generating a custom dataset for anomaly detection, avoiding the limitations associated with artificial datasets. It does that by merging expert-based formal modelling with Machine Learning (ML) modelling in a Model-Driven Engineering approach aiming at assuring the security and reliability of critical control systems from the transportation and logistics domains. This research contributes to these fields by offering a logical, traceable, and adaptable framework for anomaly detection in ICS, addressing the current challenges identified in literature and regulatory requirements.

Keywords: Industrial control systems, Safety, Security, Attack trees, Anomaly detection, Machine learning

1. Introduction

On December 24, 2015, a significant power outage struck large parts of the Ivano-Frankivsk Oblast, Ukraine, leaving 225,000 people without electricity for up to six hours (Lee, Assante, & Conway, 2016). Investigation revealed that the cause of this outage was the infiltration into the grid's computer and Supervisory Control And Data Acquisition (SCADA) system, leading to the subsequent failure of crucial components when unauthorised actors performed irregular actions in the substation's control systems. This incident, and many of its kind, clearly indicate the potential effect of anomalies impacting critical Industrial Control Systems (ICS). Another example of critical infrastructure can be found in the transport and logistics sector. In particular, railways often rely on advanced Automatic Train Protection (ATP) systems to monitor the safety of traffic on the network. One of the primary ATP systems to monitor the safety of traffic on the network is the European Train Control System (ETCS). In essence, ETCS is also based on ICS whose improper functionality could cause safety and availability implications.

The previously described critical systems underscore the importance of a secure and safe design for an ICS. Such systems underline the necessity to correctly detect, report, and act upon anomalies, whether these anomalies are caused by technical events or malicious intent. Nevertheless, from a technical standpoint, the detection of anomalies in ICS remains a challenging and unresolved topic (Ahmed et al., 2020). For instance, Koay et al. (2023) recognise the reliance on labelled datasets and the use of supervised methods for Machine Learning (ML)-based solutions in a large volume of literature. These datasets, however, are not always available for real-world scenarios, precluding the usage of supervised learning methods. At the same time, Koay et al. (2023) identify a distinct gap between theoretical research and practical applications. Specifically, the focus of a large subset of research is the accuracy of the predictions rather than the coverage of real-world attacks or the integration of a model in a production environment. Moreover, Gómez et al. (2019) acknowledge the difficulty in sourcing an ICS dataset in their proposal for a methodology of dataset generation. Lastly, from a legislative and compliance point of view, securing a system is increasingly important. In Europe, the Cyber Resilience Act and the implementation of NIS2 (Network and Information Systems Directive 2022/0383) impose a strict set of regulations related to increased cybersecurity and regulatory compliance (Chiara, 2022; Eckhardt and Kotovskaia, 2023).

Hence, more efforts and studies are necessary to tackle anomalies in ICS systems, especially in contexts that deal with vital services as these face various risks (Rotibi, Saxena and Burnap, 2024). Accordingly, this research proposes a hybrid AI-based system for anomaly detection in ICS in order to tackle existing safety and security issues in critical domains. To do that, Attack Trees expert-based modelling is applied together with ML-based modelling relying on real field data in a transparent, adaptive, and reliable approach that can be consistently applied across different scenarios. In this way, this research aims to contribute to ongoing efforts dedicated to

building intelligent systems that bridge safety and security gaps in critical domains while assuring a responsible and trustworthy stance.

The remainder of this article is structured as follows. Section 2 presents relevant studies conducted in this domain. Section 3 discusses important aspects and considerations taken in the methodological approach used for achieving the aim of this research. Section 4 provides an overview of important results obtained. At the end, concluding remarks and future research perspectives are addressed in Section 5.

2. Literature Study

Taking into account the multidisciplinary nature of this research, relevant studies from the involved disciplines (i.e., AI, formal methods – attack trees, cyber security, and ICS) are further discussed. First, system modelling and formal methods are incorporated into the research, notably through Attack Trees (Schneier, 1999), which provides a structured and systematic approach. These tools leverage expert knowledge to refine threat modelling, prioritising ICS features, and aligning the framework with business requirements and operational constraints. And secondly, ML techniques are considered, as follows: supervised learning, where labelled data is available, and unsupervised learning, which excels in handling the unlabelled and imbalanced datasets typically encountered in an ICS. Combined, these disciplines are employed to detect anomalies in complex environments and datasets.

The term 'Industrial Control System' encompasses a broad range of systems and components, from individual devices to entire networks, which are utilised to support and/or control industrial processes and operations (Trend Micro, 2024). The performance and effectiveness of these systems are evaluated using metrics such as RAMS (Reliability, Availability, Maintainability, Safety) (DMD-Solutions, 2022), to meet the specific goals of industrial processes.

The primary modelling tool for this research, Attack Trees, have been proposed as a tool for anomaly detection, threat modelling, and security analysis (Kumar et al., 2018). Ray and Poolsapassit (2005) reason that Intrusion Detection Systems fail to detect malicious insider attacks, or miss legitimate action sequences that lead to anomalous activity. They propose (Augmented) Attack Trees to model the probability of insider attacks, minimizing the scale of the Attack Trees thus improving readability and usability. Xingjie et al. (2020) propose a three-stage Attack Tree model from the attacker's perspective, attack recognition, detection and prediction using Long Short-Term Memory, LSTM.

Anomaly detection using machine learning (ML) is a common topic in the literature. For instance, Elmrabit et al. (2020) explore six deep learning methods for this purpose. Goldstein and Uchida (2016), address several methods and their caveats, in particular in relation to high-dimensional data. Lastly, Zong et al. (2018) propose advanced autoencoders to detect anomalies via density estimation. These works demonstrate varied approaches to anomaly detection. Among the core learning paradigms that exist in the ML domain are the supervised and unsupervised learning paradigms. Supervised learning requires labelled data, while unsupervised learning operates without it. For anomaly detection in ICS, unlabelled data is common, making manual labelling impractical due to dataset size and complexity (Koay et al., 2023). Thus, unsupervised, semi-supervised, or hybrid methods are most suitable. Bouman et al. (2024) list 33 unsupervised Anomaly Detection methods , though evaluation still relies on labels. Semi-supervised methods address this by combining small amounts of (manually) labelled data with large quantities of unlabelled data (Zhu, 2005). Hybrid approaches, such as this paper's use unsupervised learning to label data, followed by supervised classification of unseen data.

Hybrid approaches combining real data and expert knowledge are underrepresented in research. While formal models like Attack Trees are welldocumented and applied in Anomaly Detection, a gap exists between their construction and their implementation as a tool to convey expert knowledge in ML-based anomaly detection scenarios. Most studies use labelled dataset for performance evaluation (Koay et al., 2023), focus on data science rather than organisational priorities like traceability or legal compliance, and view anomalies as malevolent actions, neglecting errors and faults that require expertise to identify. As for the selection of the railway sector for this research, Davari et al. (2021) indicate the rail sector presents significant opportunities for research into anomaly detection due to limited academic focus, particularly on hybrid approaches. Addressing this knowledge gap is the main aim of the present paper .

3. Research Methodology

Considering the crucial role of ICS in everyday life, and the ever-changing nature of cyber threats, the need for a reliable, repeatable process is evident. However, while standards for ICS's such as SCADA exist, many ICS's are

purpose-built or otherwise too complex to build a standardised solution. Hence, the aim of this research is to use well-established methods to define a practical method for intrusion detection in industrial control systems. For this purpose, industry standard models, in particular Attack Trees as defined by Schneier (1999), and best-practices in the machine learning field are leveraged. These tools are combined with expert knowledge to expedite typically challenging tasks such as data preprocessing, feature selection, and performance optimisation. A generic design is proposed that relies on these building blocks and includes an exploratory phase, a modelling phase, and a machine-learning phase. In summary, this research aims to address the questions of how Attack Trees can be constructed and leveraged for threat modelling for ICS's, and what the efficacy of a machine learning-based anomaly detection built upon these trees, thus in essence a hybrid approach, would be.

- a) Exploratory phase

Prior to the formal modelling or the Machine Learning stage, data collection was performed. This information includes the structure of the organisation, the business priorities, and a mapping of the available data sources. This phase results in three artefacts that can be tailored to the researcher's specific requirements. The information thus collected is employed as indication for the proceeding steps.

In this first phase, the topology of the target ICS is mapped to the input from various stakeholders, experts, and key users of the system. Experts are selected based on their roles, responsibilities, and knowledge in relation to the ICS. Using semi-structured interviews, workshops, and questionnaires, the relevant knowledge for the next steps is gathered. This knowledge includes performance metrics, business goals, and priorities in relation to the ICS as well as the identification of primary data sources. This knowledge gathering phase results in a selection of tangible output documents; both for reference in further steps and cross-validation amongst the information sources.

- b) Attack Tree modelling

Once all the required exploratory information is gathered, the construction of the Attack Trees is initiated. The goal of this step is to tackle the complexity of the ICS. This is a multi-step process that results in the creation of a tool that assists in the ML process for a complex system. All Attack Trees were constructed using Attack-Defense Tree Tool (Kordy et al., 2013) based on insights from subject matter experts and information gathered during the previous stage. In this section, an insight is given into the construction of the trees, starting with the design of the template from which the Individual Attack Trees were built. Next, the design of Primary Attack Tree and assignment of the ICS's parameters to the leaves of the tree is discussed.

- i) Individual Attack Trees

The Individual Attack Tree, IAT, is the most granular part of the formal model. It represents the ICS from the perspective of a single professional with close familiarity to the system, enabling an experience-first hybrid approach to threat modelling. As one IAT is constructed per consulted expert, compatibility amongst the various trees is crucial. The resulting construction method is based on the work by Sonderen (2019), albeit with a deviating structure and the introduction of a global step, to ensure this compatibility requirement, and an Individual step.

1. Prepare Global Template
    (a) Selection of the primary goal
    (b) Selection of sub-goals
    (c) Selection of the weighting criterion
2. Individual Tree Construction
    (a) Creation and detailing of Attacks and Faults
    (b) Assigning weights

In brief, a primary goal, optional sub-goals, and a priority weighting criterion are drafted from the organisation's priorities as established during the initial discovery phase. This results in an IAT-template which is subsequently completed with information, in the form of specific attacks and faults as well as a risk score, gathered from the earlier expert interviews, group interviews, or questionnaires.

- *ii) Primary Attack Tree*

The Primary Attack Tree, PAT, is the result of the merger of all constructed IAT's and as such it represents the ICS as a whole. The construction process consists of three steps:

1. Normalisation of risk scores

2. Pruning of the IAT's
3. Merger of the IAT's

The pruning step is the most delicate part of this process. It entails the elimination of low-impact or low-importance components, resolution of duplicates and incompatible branches, and removal of redundant nodes. It is, however, crucial not to reduce the AT beyond this point to preserve maximum information density.

Upon finalisation of the PAT structure, the tree is decorated with parameters, variables, and characteristics of the ICS, as identified by the experts. This action, in essence, links the theoretical model to the practical observations made earlier.

- c) Machine learning

In the final phase of the generic design, the finalised Attack Trees are leveraged to optimise a ML pipeline. In practice, information in the Primary Attack Tree is utilised to select data sources, priorities for the machine learning algorithm, hyperparameters, and even the evaluation approach.

The machine learning strategy requires careful consideration due to the specificity of the data found in most ICS's. For instance, supervised learning methods rely on training data with labels to make accurate predictions of unseen data, but these labels are often absent with ICS data (Koay et al., 2023). However, unsupervised learning is difficult to evaluate due to the absence of a ground truth, and manually labelling the data brings its own difficulties. For these reasons, a hybrid strategy which leverages the strengths of both approaches is proposed: first, an unsupervised learning approach is implemented to group ICS data into classes, one of which represents anomalies such as system faults or cyber-attacks. Next, the now-labelled data can be used to train a supervised learning algorithm.

The groundworks of our pipeline are based on the works by Kreuzberger et al. (2023). Discussing all these steps in detail is beyond the scope of this article, so instead the focus remains on two key steps, i.e. unsupervised labelling and supervised learning. For illustrative purposes, the complete pipeline is presented in Figure 1.

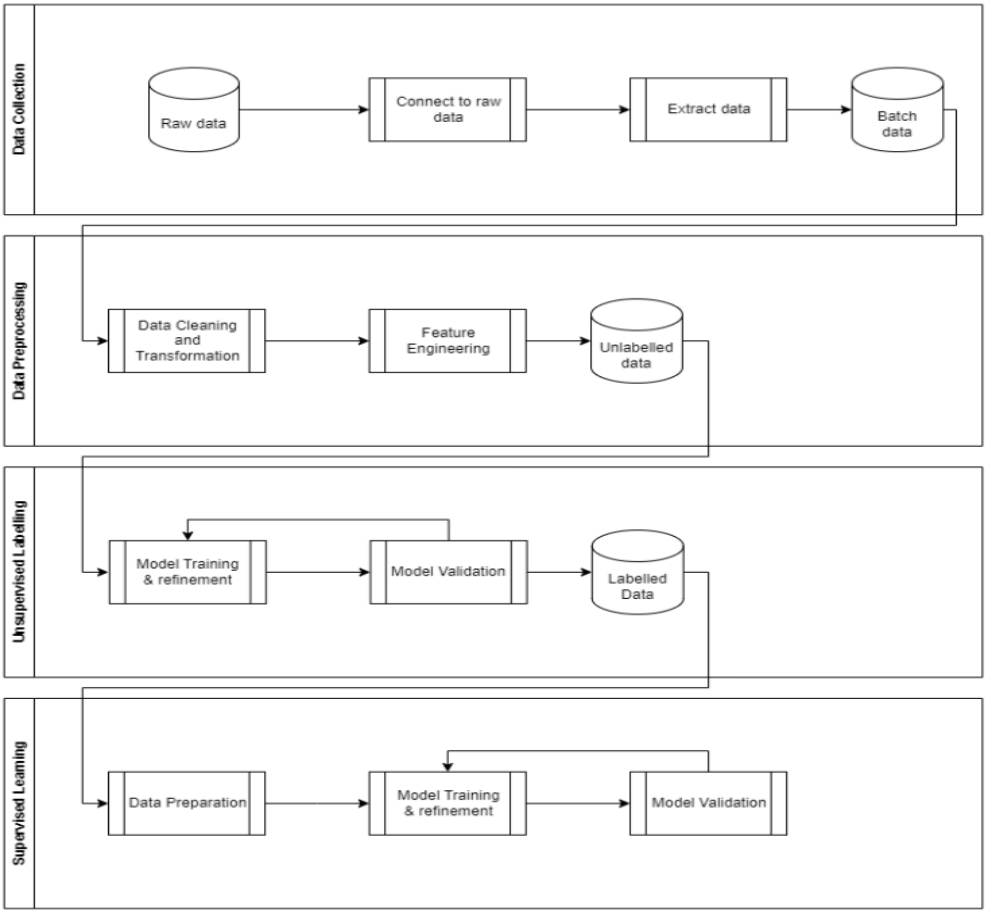

Figure 1: Simplified Machine Learning Pipeline

- *i) Unsupervised Labelling*

For the unsupervised labelling component, various models can be selected. Hence, this is a process of trial and error. In some cases, an initial guess can be made based on the data. For example: Isolation Forest performs well on imbalanced datasets and K-Means is very suitable for clusters with convex shapes (Liu et al. 2008). Model training and tuning is done recursively to optimise the accuracy. As for verification and validation, the efficacy of the labelling can be determined via clustering methods such as the Silhouette Score as proposed by Rousseeuw (1987), via expert review, or a combination thereof. Eventually, this phase results in a labelled dataset, suitable for supervised learning.

- *ii) Supervised Learning*

After labelling the training data, the data of an ICS is typically highly imbalanced with a limited number of anomalies compared to regular data. This can be addressed via oversampling of the minority class or (random) undersampling. As with the unsupervised labelling component, supervised learning can be performed via various methods. For ICS data, Random Forest or XGBoost are proposed as both are ensemble learning methods that offer advantages such as suitability for large datasets with numerous features, resistance to class imbalance, and decreased overfitting. Furthermore, Random Forest is robust in the face of noisy ICS-data, which is especially relevant if sensor data is included (Biau and Scornet, 2016).

Supervised learning methods offer more opportunities for model finetuning using optimization techniques such as Grid Search or Bayesian optimisation. These techniques can be leveraged to find suitable hyperparameters. It is advisable to align the evaluation metrics based on the priorities of the organisation and thus concentrating on either Precision and AUC-ROC or Recall and AUC-ROC. Precision is desirable if a false positive prediction is costly for the organisation, for example, if it triggers the shutdown of a production line. If a high false negative rate is costly, for instance in safety applications, recall is important. AUC-ROC is proposed for its resistance to data with class imbalance.

4. Results

To evaluate the solution proposed, the approach described earlier is applied to the European Rail Traffic Management System, ERTMS, which is a dynamic train management system responsible for the safe and efficient operation of a modern, digital railway. Thus, ERTMS represents a potential high value / high impact target for security and safety purposes. The expertise and data for this research was provided by a European infrastructure manager.

The results produced by this study fall into two categories. First, a formal model for anomaly detection in ICS's using Attack Trees is proposed based on the method described earlier. This model consists of several Attack Trees that can be used for threat identification and feature engineering. Supporting this framework, a Machine Learning-based approach is suggested to apply the anomaly detection process. This process incorporates the issues of class imbalance, missing labels, and other application-specific hurdles. The resulting hybrid approach allows security engineers to visualise and address the priorities and targets for their ICS and organisation.

During the information gathering phase, 16 staff members were interviewed. It was determined that the experts' and organisation's priorities were best described by the popular RAMS (reliability availability maintainability safety), framework, and that the organisation prioritises the absence of false negatives over false positives, i.e., they rather threat a regular datapoint as an anomaly than miss an actual anomaly. Their main goal regarding the safety and security of the ICS is to prevent any compromise. This information was leveraged as the groundworks for our study, starting with the Attack Trees. Whereas the Confidentiality, Integrity, Availability (CIA) triad is typically addressing security, RAMS can be seen as its safety counterpart.

The IAT's are built using "Compromise ETCS/ERTMS" as primary goal, and the RAMS parameters as subgoals. The template is displayed in Figure 2. A completed IAT was created for ten experts: six provided input on all RAMS-parameters, while four preferred to address their specific fields only. The resulting trees varied extensively in scope and size, with the largest tree having over 60 nodes. Each of the IAT's contains the views and experiences of the expert that helped constructing it, with each terminal leaf representing an anomaly. All anomalies are weighted for severity using a 0 (trivial) to 100 (extremely high risk) scoring.

Figure 2: Template for the Individual Attack Tree

After merging the IAT's into a Primary Attack Tree, the tree is decorated with the ERTMS parameters as found in the official documentation such as ERA Subset 026. This document contains a technical description of the various data packets present in the ETCS data channels such as the maximum allowed train velocity, movement authority length, and track gradient. This data format proposed by Subset 026 matches with the data supplied by the organisation. By assigning these data packets to leaves of the PAT, the relative importance of each packet is clear. This information will later be used during the machine learning phase to guide the data preprocessing and feature selection phases, vastly reducing the workload of these processes. Figure 3 displays a part of the PAT. The tree in this image does not show the linked data packets for security reasons.

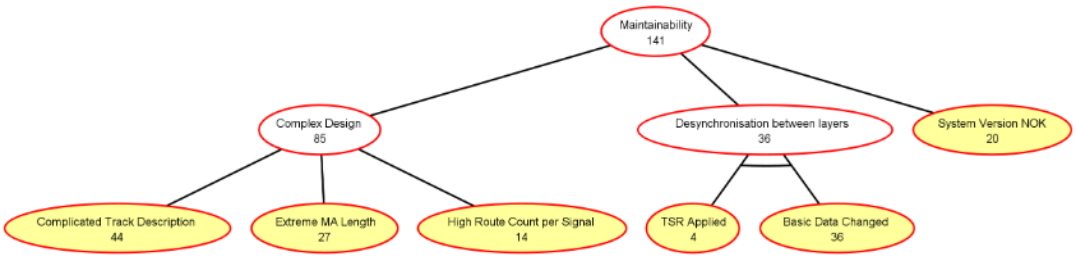

Figure 3: Partial image of the PAT. Numbers indicate the severity of each anomaly

Before proceeding to the actual machine learning process, the data collection process is conducted for training and evaluating the models. Specifically, a scenario with frequent planned stops, consistent schedules, typically error-free routes without outside influence (i.e., outside of the rush hour), is selected. 100 days' worth of train data was exported, covering the period of late 2023 – early 2024. The data is initially supplied in a proprietary format, but converted to CSV for increased flexibility.

Initial data analysis demonstrated that our data was not yet suitable for most machine learning methods due to it having many time-dependent variables. This was resolved via the creation of lagging and leading features. This process is likely to occur in other ICS's as well, as industrial processes typically have a timing aspect. After this final preparatory step and basic data preprocessing, such as the removal of Not Applicable values, the next step in the hybrid process involves applying the information contained in the Attack Trees.

Initially, the information in the attack trees is used during the feature selection process. As our data initially has hundreds of features, the machine learning methods might encounter issues such as overfitting. Reducing the amount of features, via the weighted scores in the primary attack tree and the binning of features under its leaves, brings down the total feature count to 54 with minimal effort. These 54 features, however, also include features which were not directly present in the original data, but that could be calculated from it. Inclusion of these features should improve the usability of the model. Further feature engineering actions include discarding highly correlated features, reduction of dimensionality via Principal Component Analysis, and discarding features with zero variance.

After the data preparation, four machine learning methods were applied to the data: DBSCAN, K-Means clustering, Isolation Forrest, and one-class SVM. The clustering methods (DBSCAN and K-Means) were evaluated using the silhouette score. All methods were evaluated via mutual comparison and domain knowledge via expert consultation. Of the four methods, only K-Means and Isolation Forest provided acceptable accuracy. The result for one of these methods, Isolation Forest, is demonstrated in Figure 4

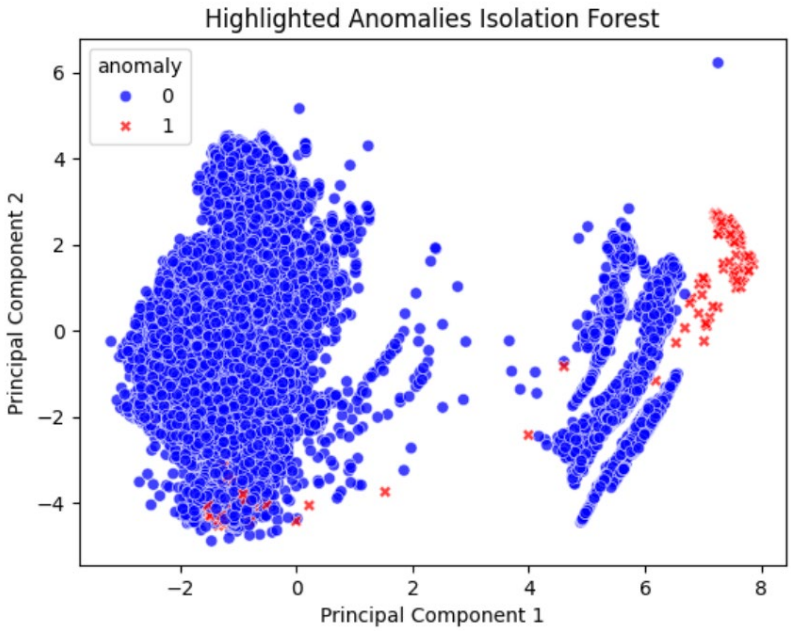

Figure 4: Results of Isolation Forest unsupervised labelling in a simplified 2D representation

Results from both unsupervised labelling methods were used as input for supervised learning. Before proceeding to the actual machine learning, a few more data preparation steps were taken. Specifically, the data was resampled due to a low anomaly rate. Next, two methods were selected: XGBoost and Random Forest. Both are ensemble learning techniques that perform well on data with class imbalance. Instead of guesstimating initial hyperparameters, Bayesian optimisation is used to optimise the performance of both models. After reaching convergence, statistical methods are used to perform a final evaluation, leading to the results demonstrated in Table 1.

Table 1: Statistical Analysis

Metric	K-Means		Isolation Forest	
	Random Forest	XGBOOST	Random Forest	XGBOOST
Accuracy	0.99	0.99	0.99	0.99
Precision	0.88	1.0	0.74	0.94
Recall	1.0	0.84	1.0	0.88
F1 Score	0.93	0.91	0.85	0.91
AUC ROC	0.99	0.92	0.99	0.94
Cross-Validation (Recall)	1.0	0.94	1.0	0.96
Cross-Validation (F1)	0.85	0.95	0.84	0.97

This table incidentally indicates why careful model evaluation is crucial: Due to the class imbalance, interpreting these values is not straightforward. Accuracy, for example, is not informative as a model that labels every instance as not anomalous would score good accuracy. Instead, precision, recall, F1-Score and AUC ROC are slightly less sensitive to class imbalance. Likewise, the typical precision-recall trade-off is quite visible as most models score better at one of these metrics than the other. Note that the selected cross-validation method, stratified k-folds, is particularly suitable for imbalanced classes.

Summing up, the selection of the best approach depends on the priorities of the organisation. For our case, as indicated during the interviews, this is the rate of false negatives. As a false negative can ultimately lead to overspeed, a signal passed at danger, or even a collision or derailment, it is much preferable to experience an untimely train stop than to allow false negatives in the system. Considering this priority, it is clear that K-Means in conjunction with Random Forest is the most suitable approach for our data set. This approach offers the highest recall score while having a higher F1 and AUC ROC score than the otherwise similar performing Isolation Forest and K-Means combination. The confusion matrix, which reflects upon the true positive, false positive, true negative, and false negative anomaly predictions, for this approach is given in Table 2. It is clear that no

false negatives were predicted, as is the goal of the organisation. However, some false positives, i.e. anomalies that were predicted but were not present, remain.

Table 2: Confusion Matrix for the hybrid method of K-Means and Random Forest

Actual Negative	8047	11
Actual Positive	0	82
	Predicted Negative	Predicted Positive

5. Conclusions

Recent incidents and evolving technical and regulatory considerations and requirements show the importance of assuring safety and security of ICS. As this remains an important challenge in various critical infrastructure services domains, it is important to build reliable adaptive intelligent solutions that would allow the prevention, detection, and response to incidents based on issues such as anomalies to assure a proper functioning of these systems and their corresponding vital services. Accordingly, in this research a comprehensive hybrid AI modelling approach is proposed for anomaly detection in ICS bridging the gap between theoretical advances and practical implementation in real-world industrial settings. The methodology centres on generating a custom dataset that avoids limitations inherent in artificial, pre-labelled datasets, thereby offering a more accurate reflection of real ICS scenarios. By merging expert-based formal modelling with ML modelling within the Model-Driven Engineering framework, this research assures the safety and reliability in critical transportation and logistics control systems providing an adaptive, transparent, and adaptable system.

As the system proposed was only tested on one ICS type (i.e., ETCS) within a single organization, even if the results obtained reflect the system's reliability, the insights may not seamlessly transfer to other ICS platforms (e.g., SCADA). Nonetheless, because the approach is architecture-agnostic and relies on packet-based messaging, it is designed to be broadly applicable. Additionally, the single-case implementation also prevents a clear measurement of workload impact, even though it notably reduced feature engineering demands. Moreover, the reliance on unsupervised learning techniques made evaluating anomaly detection more complex—yet this accurately reflects the real-world constraints in many industrial settings. In future research, studies could focus on data embedded in the attack trees proposed to classify anomalies by tracing paths from roots to goals. At the same time, synthetic anomaly generation using Generative AI techniques like Generative Adversarial Networks (GANs) or semi-supervised learning could bolster supervised methods, increasing confidence in the detection process. Furthermore, deep learning strategies, such as active learning, Large Language Models (LLMs), and transfer learning, may further address the scarcity of anomaly examples by generating data, focusing on high-value samples, or leveraging knowledge from adjacent fields. Another future perspective could be considered by building the trees as knowledge graphs with LLM's.

References

Ahmed, N.K. *et al.* (2010) 'An empirical comparison of machine learning models for time series forecasting', *Econometric reviews*, 29(5–6), pp. 594–621.

Biau, G. and Scornet, E. (2016) 'A random forest guided tour', 25, pp. 197–227.

Bouman, R., Bukhsh, Z. and Heskes, T. (2024) 'Unsupervised anomaly detection algorithms on real-world data: how many do we need?', *Journal of Machine Learning Research*, 25(105), pp. 1–34.

Chiara, P.G. (2022) 'The Cyber Resilience Act: the EU Commission's proposal for a horizontal regulation on cybersecurity for products with digital elements: An introduction', *International Cybersecurity Law Review*, 3(2), pp. 255–272.

Davari, N. *et al.* (2021) 'A survey on data-driven predictive maintenance for the railway industry', 21(17), p. 5739.

DMD-Solutions (2022) *Reliability, availability, maintainability and safety* [Online]. Available at: https://dmd.solutions/reliability-availability-maintainability-safety/ [Accessed: 2024-11-12].

Eckhardt, P. and Kotovskaia, A. (2023) 'The EU's cybersecurity framework: the interplay between the Cyber Resilience Act and the NIS 2 Directive', *International Cybersecurity Law Review*, 4(2), pp. 147–164.

Elmrabit, N. *et al.* (2020) 'Evaluation of machine learning algorithms for anomaly detection', in *2020 international conference on cyber security and protection of digital services (cyber security)*. IEEE, pp. 1–8.

Goldstein, M. and Uchida, S. (2016) 'A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data', *PloS one*, 11(4), p. e0152173.

Gómez, Á.L.P. *et al.* (2019) 'On the generation of anomaly detection datasets in industrial control systems', *IEEE Access*, 7, pp. 177460–177473.

Koay, A.M., Ko, R.K.L., Hettema, H. and Radke, K. (2023) 'Machine learning in industrial control system (ICS) security: current landscape, opportunities and challenges', *Journal of Intelligent Information Systems*, 60(2), pp. 377–405.

Kordy, B. *et al.* (2013) 'ADTool: security analysis with attack--defense trees', in *International conference on quantitative evaluation of systems*. Springer, pp. 173–176.

Kreuzberger, D., Kühl, N. and Hirschl, S. (2023) 'Machine learning operations (mlops): Overview, definition, and architecture', *IEEE access*, 11, pp. 31866–31879.

Kumar, R. *et al.* (2018) 'Effective analysis of attack trees: A model-driven approach', in *Fundamental Approaches to Software Engineering: 21st International Conference, FASE 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Proceedings 21*. pp. 56–73.

Lee R.M, Assante M.J. and T., C. (2016) 'Analysis of the cyber attack on the Ukrainian power grid', *Electricity information sharing and analysis center (E-ISAC)*, 388(1–29), p. 3.

Liu, F.T., Ting, K.M. and Zhou, Z.-H. (2008) 'Isolation forest', in *2008 Eighth IEEE International Conference on Data Mining*. IEEE, pp. 413–422.

Ray, I. and Poolsapassit, N. (2005) 'Using attack trees to identify malicious attacks from authorized insiders', in *Computer Security--ESORICS 2005: 10th European Symposium on Research in Computer Security, Milan, Italy, September 12-14, 2005. Proceedings 10*. Springer, pp. 231–246.

Rotibi, A., Saxena, N., & Burnap, P. (2024). Winning the battle with cyber risk identification tools in industrial control systems: A review. IET Cyber-Physical Systems: Theory & Applications, 9(4), 350-365.

Rousseeuw, P.J. (1987) 'Silhouettes: a graphical aid to the interpretation and validation of cluster analysis', *Journal of computational and applied mathematics*, 20, pp. 53–65.

Schneier, B. (1999) 'Attack trees', *Dr. Dobb's journal*, 24(12), pp. 21–29.

Sonderen, T. (2019) *A manual for attack trees*. University of Twente.

Trend Micro (2024) *Industrial Control System - Definition* [Online]. Available at: https://www.trendmicro.com/vinfo/be/security/definition/industrial-control-system [Accessed: 2024-11-12].

Xingjie, F. *et al.* (2020) 'Industrial control system intrusion detection model based on LSTM & attack tree', in *2020 17th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*. IEEE, pp. 255–260.

Zhu, X.J. (2005) 'Semi-supervised learning literature survey'.

Zong, B. *et al.* (2018) 'Deep autoencoding gaussian mixture model for unsupervised anomaly detection', in *International conference on learning representations*.