10.48047/jocaaa.2024.33.08.108

AI/ML-Driven Service Management for Enhancing Organizational Overall KPIs and Security

"Syed Umair Akhlaq
DIAC Solutions, UK

Abstract

This research addresses IT Service Management (ITSM) challenges to introduce AI to improve Key Performance Indicators (KPIs) such as SLA compliance, Mean Time to Resolve (MTTR), Mean Time to Detect (MTTD), and strengthen cybersecurity. Using the Kaggle IT Incident Log Dataset, we have designed an extensive framework that utilises Random Forest and XGBoost to classify incidents, Linear Regression to predict KPIs and Isolation Forest to detect anomalies. The methodology was highly involved; data was preprocessed, feature was engineered, and grid search was applied using GridSearchCV to optimise the model, thus to guarantee robust estimation. The results indicated that Random Forest realised a classification accuracy of 68% and an F1–score of 0.67, XGBoost obtained 66% accuracy, and an F1–score of 0.64. For the prediction of KPI's, the forecasting model of MTTR gave an RMSE value of 2121.87 hours, which suggests not very effective predictive capability due to the complexity of the dataset. Pr-Auc of anomaly detection performance was at 0.50, where class imbalance constrained the learning. Visualisations such as confusion matrices and an MTTR trend chart show a resolution time reduction of 860 to 825 hours post-AI. These conclusions demonstrate the possibility that AI can improve ITSM efficiency and security, and make way for a more adaptive, predictive service management framework in enterprise environments.

Keywords: *ITSM, Cybersecurity, AI in Service Management, KPI Optimisation, SLA Compliance, Incident Prediction, Machine Learning, Secure IT Operations.*

1. Introduction

The drastic change in technology has disrupted the management of enterprise IT operations, posing simultaneously a multifaceted challenge for organisations. Delivering operational excellence with a focus on solid cybersecurity. Some of the measurements employed in operational efficiency within the IT Service Management (ITSM) include the Service Level Agreement (SLA) compliance, the Mean Time to Resolve (MTTR), and the Mean Time to Detect (MTTD). SLA compliance guarantees that incidents are addressed within the agreed timelines, while MTTR quantifies the average time it takes to resolve incidents, and MTTD measures the speed of detecting incidents. On the other hand, mainstream ITSM tools are rather reactive, using predefined workflow patterns and mandatory manual interventions [1]. This reactive attribute frequently leads to delayed resolutions, SLA goals, and increased exposure to security threats since these tools do not support predicting events or actively detecting anomalies.

The inadequacies of conventional ITSM systems have led to the adoption of Artificial Intelligence to continue solving the deficiencies experienced. AI has grown from simple automation – scripted responses and rule-based escalations – to intelligent decision support systems that can do complex predictive analysis and anomaly detection [2]. Through big data analysis, AI can detect trends in incident data, predict possible disruptions, and identify abnormal behaviour that might indicate a security breach [3]. For example, through AI-based predictive models, high-priority incidents can be predicted, resulting in the proactive allocation of resources. At the same time, anomaly detection algorithms can detect abnormal system logs (for example, unauthorised access attempts), before aggregation to the greater threats [4].

Despite these developments, a vast gap remains in integrating KPI optimisation, with cybersecurity advancements, into the existing AI-ITSM frameworks. While many solutions target better operational metrics (such as decreasing MTTR) or better security (such as anomaly detection), they rarely address both

10.48047/jocaaa.2024.33.08.108

in a unified way. This fragmentation performs a disservice to the interlinked states of efficiency and security in ITSM, where the lags of incident resolutions compound on security and unidentified threats disrupt service delivery. For instance, a long MTTR for a security incident may allow a breach to proliferate, and a shortage of anomaly detection may not pick up early signs of a cyber-attack, causing an SLA violation [5]. This research gap suggests the necessity of a unified framework within which AI is used to optimise KPIS and improve cybersecurity simultaneously, thus the need for a balanced approach for ITSM. In addition, AI amplifies the performance of traditional SIEM systems by bridging the gap between detection and resolution. For example, the system can independently activate workflows like credential revocation, provisioning detailed alerts, or sandboxing the network segments to respond to a threat detection. These self-healing mechanisms limit the dwell time and avoid threat escalation without human intervention.

To fill the gap, this study suggests an AI-driven framework based on the Kaggle IT Incident Log Dataset, which gives real-world incidents for analysis. The framework's goal is to accomplish two primary goals: first, in terms of improving KPIs, incident classification, and prediction of MTTR to make incident handling faster and more efficient, and, second, in terms of improving cybersecurity, through the detection of anomalies, which will allow to identify possible threats timely. By incorporating these components, the framework aims to strengthen ITSM outcomes holistically. The structure of this paper, which is to summarise related work, describe methodology, report findings, outline the implications, and end with a discussion of future directions, makes an original contribution to the development of effective adaptive ITSM in current enterprises.

This study contributes to the field by introducing a new AI-driven framework, which incorporates the process of incident classification, KPIs prediction, and anomaly detection, filling the gap in the need for a unified approach in ensuring operational efficiency and cybersecurity of ITSM. Using real-world data from the Kaggle IT Incident Log Dataset provides practical insights and a scalable solution that leads to adaptive ITSM practices that can respond to the advancement in the IT world. Programmers and IT experts need to find a proper mechanism to utilise sources such as the IT incident log data to make data-driven predictions that can help them forecast the future of IT infrastructure, such as incident rates, ticket duration, etc.

2. Literature Review

2.1 ITIL v4 and Process-Based ITSM

The IT Infrastructure Library (ITIL) v4 provides a holistic framework for IT Service Management (ITSM) by concentrating on value-driven service delivery across processes such as incident management, problem management, and service desk operations [6]. ITIL v4 re-invents the traditional ITSM model by focusing on a holistic approach that routes the IT services towards business goals and incorporates practices such as service level management and continual improvement. [7] KPIs play a central role in this framework, such as SLA compliance (the percentage of incidents resolved within agreed timeframes), Mean Time to Resolve (MTTR, the average time to resolve incidents), and Mean Time to Detect (MTTD, the average time to detect incidents. Such metrics bring a uniform approach toward measuring service performance and operational efficiency. Tools like ServiceNow, BMC Remedy, and Cherwell are implementing ITIL principles, and they provide dashboards and reporting abilities to see those KPI’s in real time [8]. These systems automate workflows, manage ticket assignments and escalations, and increase visibility of service operations. However, their dependence on prescriptive workflows and manual escalations makes them more or less reactive. This reactive nature inhibits their capacity to forecast incidents or adapt to the dynamic, evolving IT environments loaded with frequent updates, new technologies, and emerging threats. Consequently, slow resolution times and being behind schedule relative to early threat detection are common, which calls for greater proactive ITSM solutions.

10.48047/jocaaa.2024.33.08.108

2.2 Traditional KPI Tracking

Conventional ITSM tools use static dashboards to measure KPIs like SLA conformity, MTTR, and MTTD to gain a historical view of service performance [9]. These dashboards consolidate metrics from incident logs, including a percentage of events that meet SLA or average times to resolution across different categories. Their predictive capabilities are low, though, and they necessitate manual analysis by IT personnel to derive trends or preempt issues. Escalation processes, which are frequently initiated through breach of SLA thresholds, require significant dependency on humans in prioritising which susceptibility to delay, especially in high-impact or complex situations [10]. This manual process does not draw on the inherent order and correlation against incident data, for example, repeating issues or seasonal blasts, and thus is inefficient. The frequent outcomes are increased downtime as reactive responses replace proactive ones. Furthermore, possible security vulnerabilities (e.g., anomalous system behaviour or illegal access attempts) are frequently not identified until they result in significant events. Such a reactive approach emphasises the constraints of conventional KPI tracking, especially in settings where timely reaction and anticipatory measures are essential for ensuring quality of service and security.

2.3 AI/ML in Service Automation

Artificial Intelligence (AI) and Machine Learning (ML) have transformed ITSM into an automation and predictive giant, moving the paradigm from reactive to proactive operations. Natural Language Processing (NLP) is a part of this transformation process [11]. This allows systems to categorise tickets according to textual descriptions, for example, incident summaries and user reports, and prioritise them or assign categories automatically. This will minimise manual work and incident response. ML models such as Random Forest and XGBoost go a step further to predict incident urgency and escalation needs, which in turn help organisations distribute resources more effectively on anticipated workload or severity [12]. [13] Anomaly detection using algorithms such as Isolation Forest detects odd patterns of logs that may signify operational inefficiencies, such as system overload, security threats, and data exfiltration attempts. These advancements allow proactive incident management, reducing resolution times and improving service quality through problem minimisation before its effect on users. Nevertheless, implementing such ML techniques differs from one organisation to another depending on data availability, technical knowledge, and infrastructure preparedness. This inconsistency calls for a standardised approach to using them to ensure maximum benefits in ITSM.

2.4 AI in Cybersecurity

The ability of AI to improve cybersecurity in ITSM is profound, primarily through leveraging SIEM analytics and insider threat detection [14]. SIEM systems use AI to correlate logs and alerts from multiple sources such as firewalls, servers, applications, and in the process, instrumenting potential threats such as unauthorised access, malware, infection, or a data breach [15]. This correlation functionality enables real-time threat detection and response to the minimum possible period of opportunity available to the attacker. The ML models reinforce SIEM by taking in the user behaviour and identifying unusual patterns of logins, bursty downloads, or access to areas with restricted access that may point to insider threats [16]. These models create baselines of regular activity using historical figures, which flag abnormal activity for investigation. Security management by automated alert prioritisation based on AI ensures a timely response to priority critical security incidents, which conform to frameworks such as the National Institute of Standards and Technology (NIST) Cybersecurity Framework and the Centre of Internet Security (CIS) controls [17]. This integration enhances the ITSM posture by using proactive defence mechanisms. Nevertheless, such cybersecurity improvements tend to develop mainly independently of operational KPI optimisation, creating silos that prevent the holistic improvement of ITSM competency.

2.5 Research Gap

In light of advances in AI and ML, existing AI-ITSM integration frameworks suffer from significant constraints. A lot of attention is paid to areas like operational efficiency, e.g., decreasing MTTR or increasing compliance with SLA, while cybersecurity integration is often ignored. For example, NLP-based

10.48047/jocaaa.2024.33.08.108

ticket classification simplifies ticket classification but does not address problems related to security discrepancies. SIEM systems are excellent at detecting threats but lack integration with KPI, such as MTTR prediction. This approach is disjointed and thus prevents a unified ITSM strategy, considering the interrelation of operational and security goals. Furthermore, static models trained on historical data cannot cope with an ever-changing IT environment, new threats, and incident patterns, and the lack of feedback loops or re-training further widens vulnerabilities. The main research gap is a lack of a combined model improving KPIs (e.g., MTTR, SLA adherence) and cybersecurity (e.g., anomaly detection) in ITSM. Existing solutions lean towards one or the other, disregarding their interlinking nature – delays in resolution can exacerbate security risks, and unaddressed abnormalities interfere with services. This study tackles this with an integrated AI framework through incident classification, KPI prediction, and anomaly detection. It tests it with the Kaggle IT Incident Log Dataset to improve ITSM holistically.

Although with numerous perspectives, despite various applications, today's AI-based security systems have limited or no integration with ITSM workflows. Much of systems depend on alerting without automating mitigative actions. Furthermore, black-box models are opaque, with security teams unable to trust or take action based on insights provided by AI. The lack of feedback loops between incident response and model refinement also contributes to increasing non-adaptability in a changing threat landscape.

3. Methodology

3.1 Proposed Methodology

The proposed methodology provides an integrated AI-driven framework to improve IT Service Management (ITSM) results and strengthen cybersecurity. This framework comprises three central elements: incident classification and KPI prediction, with a feedback loop for continuous improvement and adaptation in a dynamic IT environment.

The incident classification module can classify incidents using Random Forest and XGBoost classifiers. These models utilise features like priority, assignment group, and ticket urgency, calculated as the sum of impact and urgency, to determine the incidents' class into predefined categories. For proper performance, hyperparameter tuning is performed via GridSearchCV to optimise parameters such as some estimators (100 or 200) and the maximum depth (5 or 10) on the Random Forest. XGBoost is similarly configured, though multiclass classification accuracy is emphasised. Such a tuning process improves classification accuracy by creating a trade-off between model complexity and the ability to generalise; the models become more appropriate to the diversity of ITSM incidents and minimise overfitting.

The KPI prediction module relies on Linear Regression to predict Mean Time to Resolve (MTTR), an important KPI in ITSM. Features such as escalation count (the sum of reassignment and reopen counts) and action frequency (system modification count/MTTR) serve as predictors, reflecting dynamics of incident resolution. Linear Regression offers a simple yet effective way to predict resolution times for proactive allocation of resources and SLA management. By predicting the MTTR, this module allows organisations to plan for workload requirements and resource assignment, consistent with overall organisational objectives, to drive efficiencies in operations.

The anomaly detection engine uses Isolation Forest to discover security anomalies in incident logs based on features such as urgency of the ticket and action frequency. Isolation Forest was selected because it is effective in high-dimensional datasets and provides out-of-range signals by isolating data points by employing random partitioning. A contamination factor 0.05 is established, indicating that 5% of the data will be marked as anomalies, which conforms to normal anomaly rates in ITSM contexts. The engine interfaces with a security layer that associates detected anomalies to individual assets and prioritises alerts based on levels of severity and confidence derived from the model output. This prioritisation guarantees a timely review of possible threats, which minimises the risk of escalation and aligns with the optimum practice in cybersecurity.

10.48047/jocaaa.2024.33.08.108

A feedback loop is therefore embedded to make predictions more accurate and the framework more flexible. Misclassified incidents or incorrect MTTR predictions are examined to adjust model weights or their feature importance, leading to better future performance. For instance, if an incident is misclassified, the feedback loop creates new boundaries for the classifier's decision. At the same time, differences in predictions of the MTTR can cause reweighting of the features, such as the escalation count. Comprehensive set of metrics is used to evaluate all models: Degree of accuracy and balance in classification (precision, recall, and F1-score); method of prediction accuracy in regression (Mean Absolute Error (MAE) and Root Mean Square Error (RMSE)); and calculation of performance in imbalanced datasets for anomaly detection (Precision-Recall Area Under Curve (PR-AUC)). This multi-metric evaluation comprehensively reviews the framework's effectiveness for its varied tasks.

The proposed AI framework, apart from incident classification, KPI prediction, and anomaly detection, can be applied to the key ITSM areas, including Change Management, Configuration Management, Knowledge Management, and Availability Management. These domains target IT change control, asset recording, operational knowledge sharing, and service uptime, respectively. Using the same AI techniques, Classification, Regression, and Anomaly Detection, this framework enables proactive planning, risk aversion, and strategic decision making. Its flexibility allows for integration in numerous ITSM functions, making it scalable and operationally efficient in a comprehensive enterprise IT operation while endorsing reliability, efficiency, and intelligence in service management practices.

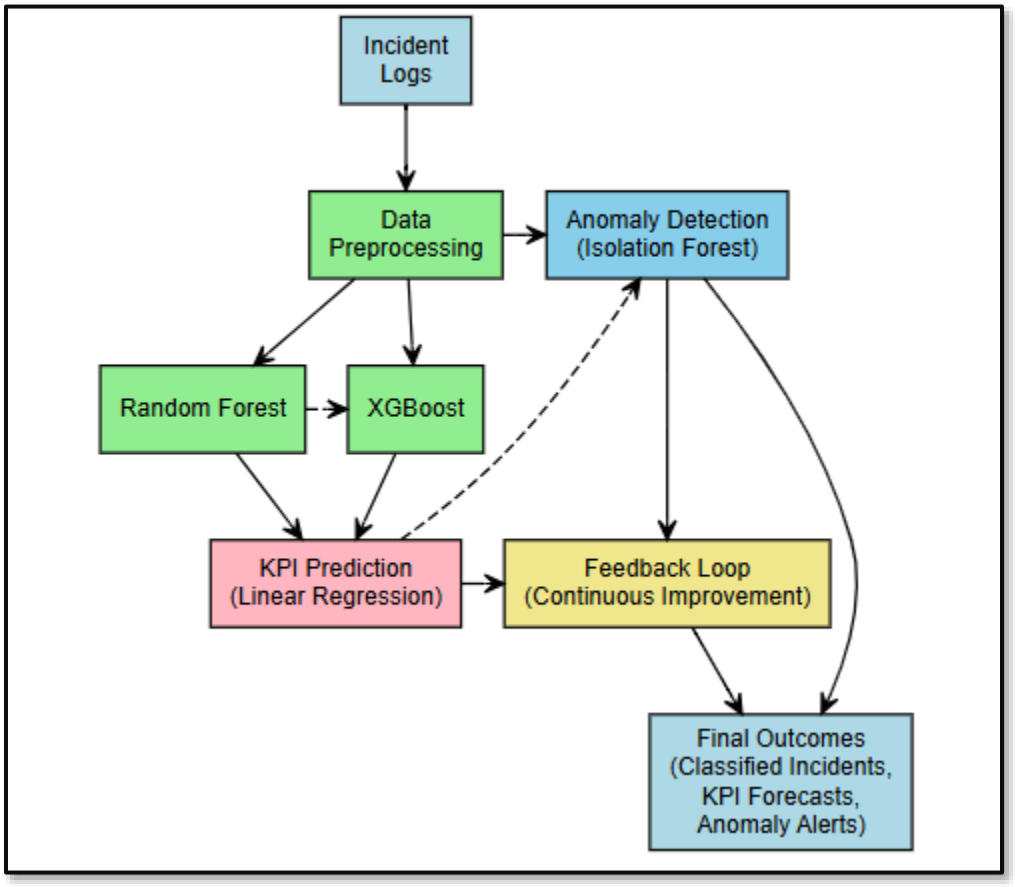

Figure 1: Proposed Methodology

10.48047/jocaaa.2024.33.08.108

3.2 Dataset Description

The Kaggle IT Incident Log Dataset serves as the foundation of this study by presenting a realistic example of ITSM work. This dataset contains fields like Incident ID, Timestamp, Category, Assignment Group, Priority, Activity Type, and status indication, like resolved_at and closed_at. After cleaning, the data set includes 987 records and therefore is appropriate for identifying incident classes, modelling KPIs, and detecting anomalies. Timestamps (e.g., opened_at, resolved_at) allow time-based feature engineering like MTTR calculation – resolved_at – opened_at, which is crucial for KPI prediction. Fields like priority and assignment_group of categorical type will give us context on the level of incident's severity and how resources should be allocated, and fields like closed_code enable us to pinpoint security-related incidents. The dataset's variety, including incident types and operational scenarios, allows the verification of the suggested framework in practice against a live application, making it applicable to enterprise ITSM environments.

3.3 Data Preprocessing

Pre-processing the data is essential for a dataset to be appropriate for modelling. The process starts with how missing values in fields such as u_symptom, problem_id, rfc, vendor, and caused_by are handled by filling with "Unknown" or "None" to ensure consistency and prevent data loss. Opened_at, sys_created_at, sys_updated_at, resolved_at, and closed_at are converted to datetime format using pd.to_datetime and set to coerce errors to NaN to take care of invalid entries. Rows having null values in vital fields opened_at, resolved_at, priority, and assignment_group are discarded as these fields are necessary for calculating MTTR & classification.

Quantitative variables like priority, assignment_group, category, incident_state, and contact_type is encoded using LabelEncoder, which converts them into numerical values to facilitate compatibility of ML models. Feature engineering makes the dataset usable by adding new variables: MTTR is calculated by subtracting the value of opened_at from resolved_at and converting the difference into hours; ticket_urgency is computed as an aggregate value of impact and urgency – which captures overall severity of incidents; escalation_count is computed as an aggregate value of reassignment_count and reopen_count, indicating incident complexity; action_frequency is calculated as the ratio of sys_mod_count to MTTR – which is indicative of intensity of actions per unit of resolution time. By dropping remaining NaN values in MTTR or action_frequency, the dataset integrity is maintained to create a cleansed dataset that is ready for analysis.

3.4 Generic AI Model Framework

A generic AI model framework has been created that simplifies the entire procedure of data-driven IT Service Management (ITSM) improvement, including data preprocessing and model assessment. The first step of the framework consists of data ingestion and cleaning, providing high-quality input with structured extraction and efficient work with null values. This is then followed by feature engineering in which domain-specific attributes such as Mean Time to Resolve (MTTR), escalation count, and ticket urgency are engineered to capture the operational dynamics. Based on the task (classification, regression, or anomaly detection), proper models such as Random Forest, Linear regression, or Isolation Forest are chosen. To maximize the performance of the models, hyperparameter tuning will be carried out, which will be used to balance accuracy and generalizability by using GridSearchCV. Effectiveness of each model is then measured using certain metrics relevant for each model, such as F1-score for classification, RMSE for regression and PR-AUC for imbalanced anomaly detection. Critically, the framework includes a feedback loop in which mislabeled incidents or inexplicable predictions rework model parameters and feature weights iteratively. This modular, adjustable framework not only makes incident management better but has also been developed to enable other ITSM areas like change management and configuration management through structured historical log data for predictive insights and preemptive actions.

10.48047/jocaaa.2024.33.08.108

3.5 Model Implementation

Random Forest and XGBoost models are trained on a feature set of priority, assignment_group and ticket_urgency for incident classification. A train test split of 80-20 is used with an 80% training set and a 20% testing set, so robust validation exists. GridSearchCV optimises Random Forest hyperparameters by trying combinations of n_estimators (100, 200) and max_depth (5, 10), maximising F1-score for precision and recall. XGBoost is tuned to use mlogloss evaluation metric to support multiclass classification, guaranteeing correct predictions for the 17 incident categories in the dataset.

KPI prediction utilises the Linear Regression tool for the MTTR prediction with the help of the ticket_urgency, escalation_count, and action_frequency. This model is a basis for incident resolution time forecasting, assisting organisations in accurately predicting and controlling the incident resolution timelines. The same 80-20 train-test split guarantees uniformity for evaluation in tasks.

Anomaly detection uses Isolation Forest targeting the issue of ticket_urgency and action_frequency for the identification of security-related anomalies. Anomalies have labels according to the closed_code field in which the term "Security" is present, and the model assigns scores (-1 to anomalies and 1 to the normal instances). An indicator of contamination, 0.05 points to the level of proposed prevalence of anomalies, helping to take a balanced approach towards identifying the possible security issues in ITSM workflows.

3.6 Evaluation Metrics

Performance of classification is measured with the help of precision (proportion of correct positive predictions), recall (proportion of actual positives correctly predicted), and F1-score (harmonic mean of precision and recall) to measure the balance of accuracy. Confusion matrices illustrate prediction distributions: correct classification and incorrect classification for categories. For KPI prediction, MAE quantifies the average deviation between predicted and actual MTTR, and RMSE evaluates the differences squared, reflecting high errors. Anomaly detection relies on PR-AUC to measure performance in the context of class imbalance, since the presence of a class imbalance is severe in the dataset having a high built-in skew towards non-anomalies. ROC-AUC is skipped in case of insufficient class diversity, which will result in a meaningful evaluation. All these metrics offered a comprehensive review of the framework's effectiveness in fulfilling the varied needs of incident classification, KPI prediction, and anomaly detection in ITSM.

10.48047/jocaaa.2024.33.08.108

4. Results

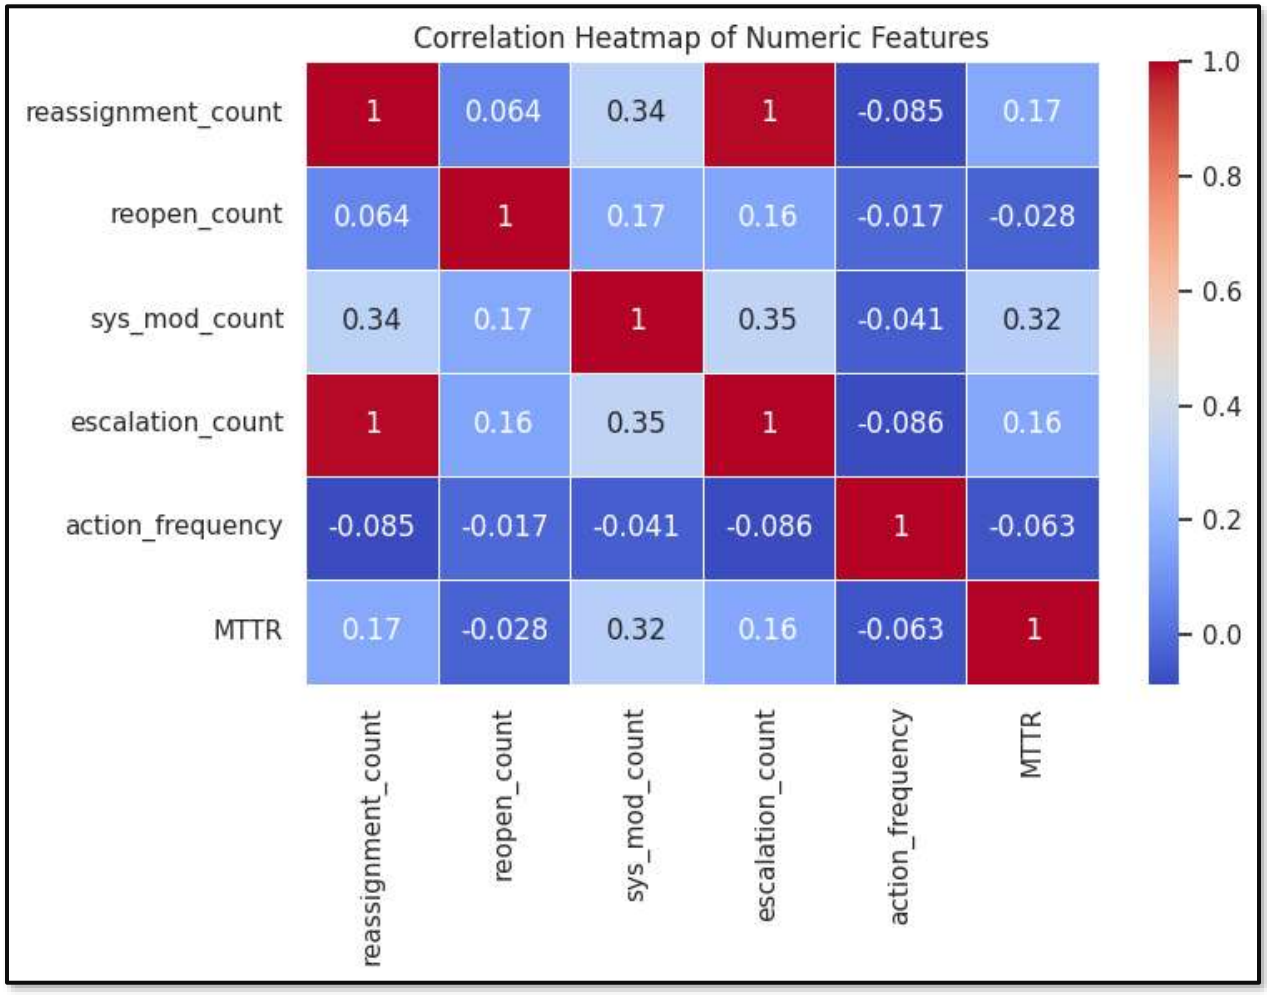

Figure 2: Correlation Heatmap

Figure 2 shows a Correlation Heatmap of Numeric Features depicting relationships between variables, including reassignment_count, reopen_count, sys_mod_count, escalation_count, action_frequency, and MTTR. Positive correlation (e.g., 1.0) is observed for reassignment_count and escalation_count and sys_mod_count and escalation_count (0.35), meaning frequent reassignments and modifications increase escalation. Negative correlations, such as a -0.085 between reassignment_count and action_frequency, indicate lower action frequency with increased reassignments. MTTR demonstrates moderate positive correlations (for example, r = 0.32 with sys_mod_count) because it depends on the modification frequency.

10.48047/jocaaa.2024.33.08.108

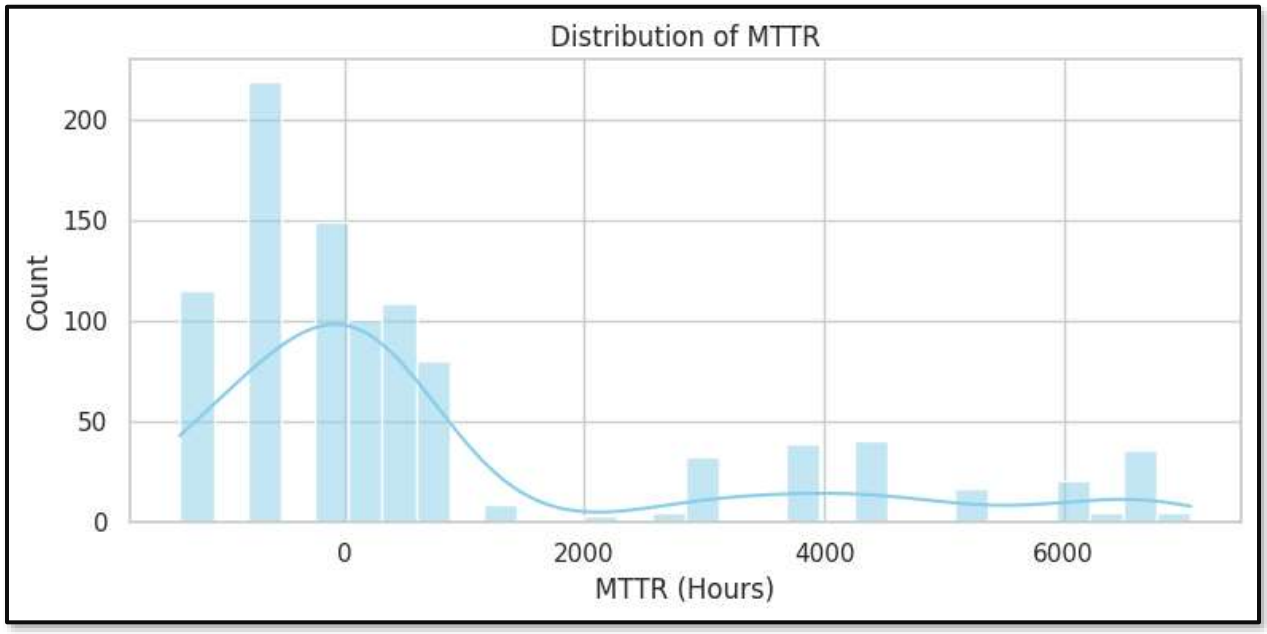

Figure 3: Distribution of MTTR

Figure 3 illustrates the distribution of MTTR and the frequency of Mean Time to Resolve (MTTR) values in hours. The histogram shows a skewed distribution with a marked peak at 0 hours, corresponding to many incidents closed almost on the spot (approximately 200 counts). There is a secondary peak between 500 and 1500 hours, with counts varying between 100 and 150, meaning a large proportion of agitation handling has moderate incident resolution time. The distribution falls off toward greater values beyond hours of 2000, with less frequent occurrences taking broader spans of 4000-6000 hours (counts fall to 20-50), and beyond having a long tail to greater values. Once the overlaid density is produced, this trend is smoothed, and this skewness is confirmed; MTTR values concentrate around lower ranges. This variation of the distribution highlights the differences in time to resolve incidents, where most are resolved almost immediately; a minor percentage require much longer time frames, which may be attributed to the complexity of the dataset.

10.48047/jocaaa.2024.33.08.108

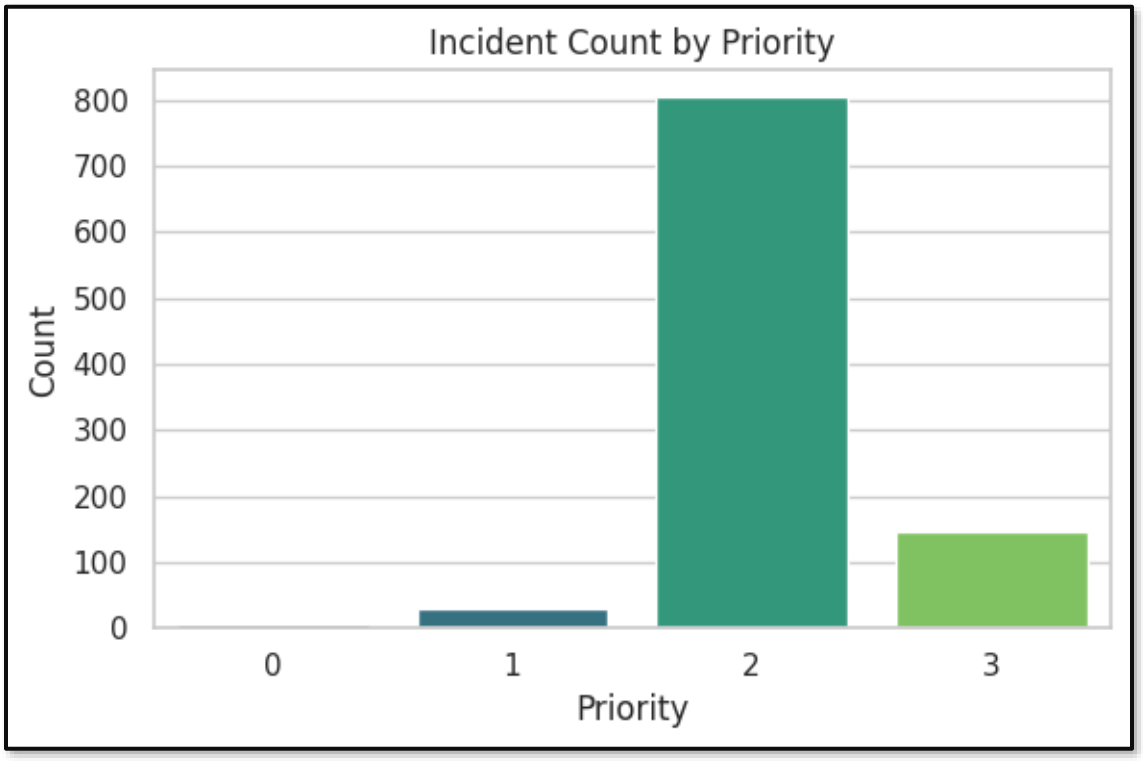

Figure 4: Incident Count by Priority

Figure 4 shows the Incident Count by the Priority level, summarised as a distribution of incidents across three levels of priorities (0, 1, 2, and 3). Based on the bar chart, the concentration at priority 2 is most prominent, with more than 700, which denotes a high intensity of, more or less, critical incidents. Priority 3 follows with around 150 incidents, with the smaller yet notable quantity of higher priority cases. On the other hand, priority 1 has a negligible count (close to 0), and priority 0 displays no incidents. This indicates an imbalance where most incidents fall into the mid to high priority band. Such distribution highlights the moderate urgency of the issues studied in the dataset while concentrating more on issues with a low or even critical priority, which might impact the model quality in the area of classification tasks. In addition, it can make certain underrepresented priority levels worth further distinction.

10.48047/jocaaa.2024.33.08.108

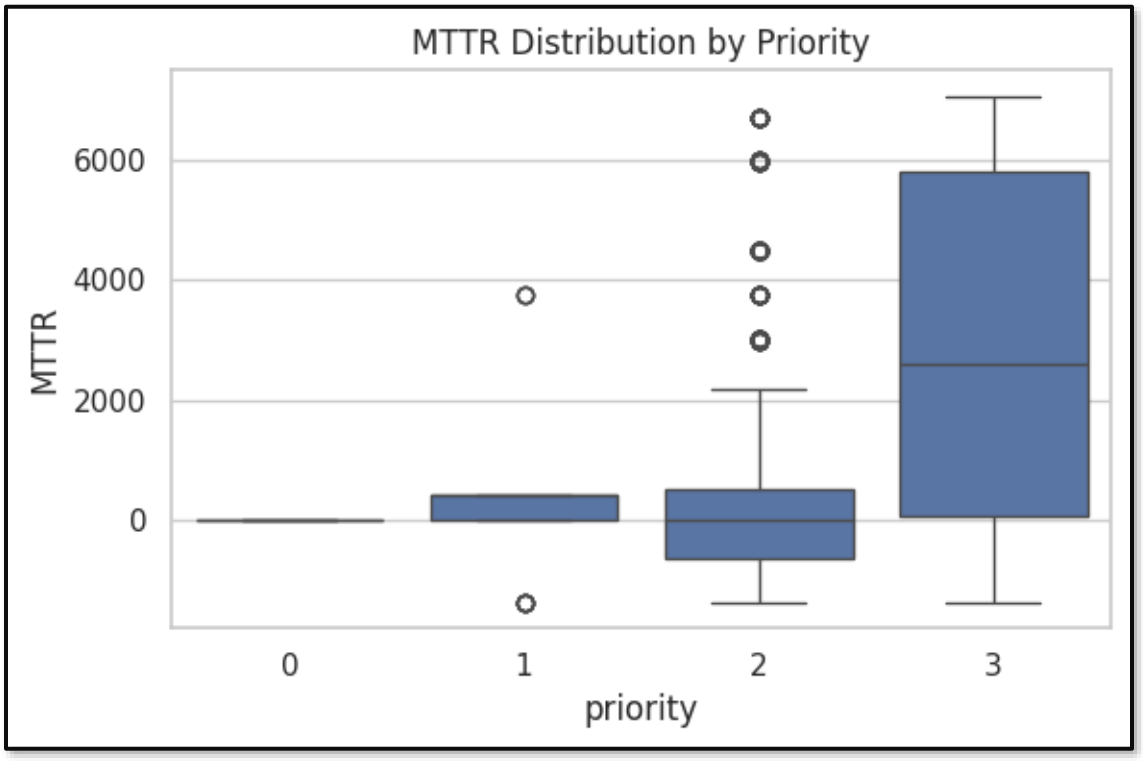

Figure 5: Box Plot (MTTR by Priority)

Figure 5 shows the MTTR Distribution by Priority with a box plot of MTTR across priority 0, 1, 2, and 3. With most values near 0 hours, priority 0 has a very low variation of MTTR, meaning immediate resolution. Median MTTR of around 500 hours resides for Priority 1, yet a narrow IQR and the ability to reach as high as 4000 hours indicate some time-based escalation. Priority 2 has a median MTTR of approximately 1000 hours with a larger IQR range, and several outliers up to 6000 having a higher level of variability and longer resolution times. Priority 3 (with the highest median MTTR (approximately 4500 hours) and a wide IQR) reflects the highest median, fluctuating resolution times, which go up to 6000 hours. This trend shows that high priority incidents (2 and 3) take much more time to manage, probably because of their complexities, whereas the low priority incidents (0 and 1) are resolved fairly faster, which corresponds with levels of their urgency and expectations.

10.48047/jocaaa.2024.33.08.108

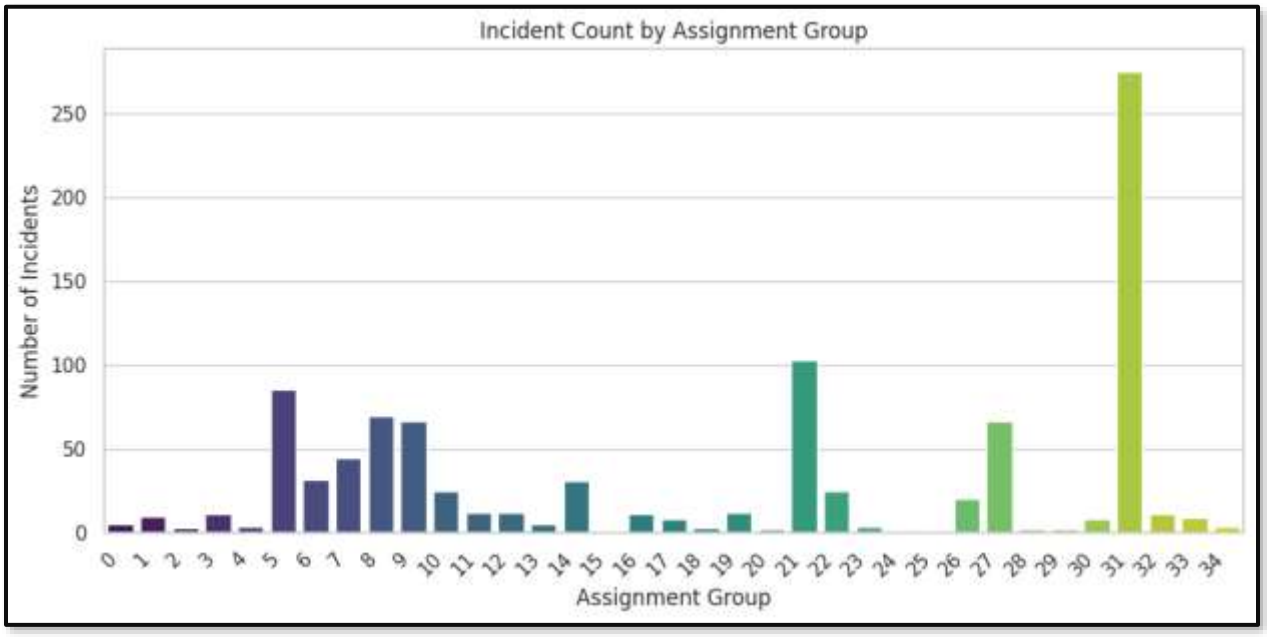

Figure 6: Bar plot for assignment group

Figure 6 depicts the Incident Count by Assignment Group, which displays the assignment groups (labeled 0 through 34) where the incidents have been distributed. The bar chart shows a highly skewed distribution with a large peak in assignment group 34, greater than 250 incidents, which therefore suggests that this group handles most of the cases. Assignment group 21 is next with about 100 incidents, and thus a medium volume, while assignment groups 0, 1, 3, 5, 6, and 9 all have approximately 50 – 70 incidents, and therefore, medium activity. The other major groups (e.g., 2, 4, 7, 8, 10-20, 22-33) have frequencies of less than 50 incidents, with many near zero, thus low or zero incident allocation. This imbalance shows that there are a few assignment groups, such as 34 and 21, that are heavily weighted, either by specialization or by higher incident complexity. The uneven distribution may affect the allocation of resources and performance of the model, thus, there will need to be implemented to ensure a balance of workloads across groups.

10.48047/jocaaa.2024.33.08.108

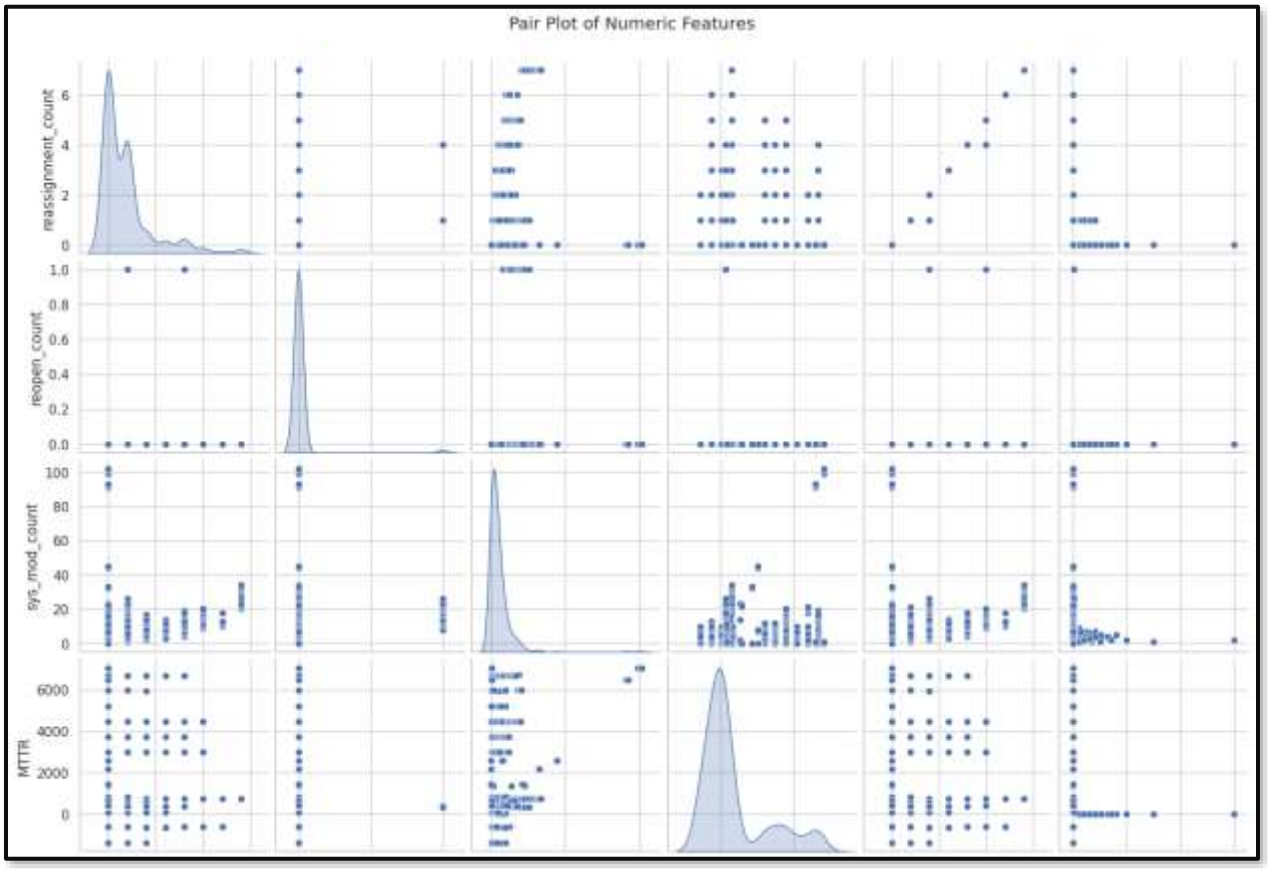

Figure 7: Pair Plot

Figure 7 shows the Anomaly Score Distribution, which is a histogram of anomaly scores produced by the Isolation Forest algorithm. The distribution is bimodal, with a high peak at 1.0 (lack of anomalies), which exceeds 900 counts; thus, most data points (approximately 937 in 987) are normal. The smaller peak at -1.0 [anomalies] has approximately 50 counts, similar to the 0.05 contamination factor, meaning 50 records will be singled out as potential anomalies. The overlaid kernel density curve smooths this trend, validating the concentration of scores at the extremities, while density is sparse between -1.0 and 1.0. This severe difference indicates how well the model can separate the anomalies from normal cases; however, the small number of anomalies emphasizes the class imbalance of the dataset. This distribution outlines the difficulty in detecting rare occurrences in security event that particularly calls for enriched datasets to enhance the anomaly detection capability.

10.48047/jocaaa.2024.33.08.108

Random Forest Classification Report

	precision	recall	f1-score	support
0	0.50	1.00	0.67	1
2	0.50	0.54	0.52	13
3	1.00	0.20	0.33	5
4	0.75	0.60	0.67	5
5	0.33	0.14	0.20	7
6	0.75	0.50	0.60	12
7	1.00	0.67	0.80	3
8	0.64	0.69	0.67	13
9	1.00	1.00	1.00	1
10	0.75	0.33	0.46	9
11	1.00	0.50	0.67	4
12	0.91	0.88	0.90	34
13	0.56	0.94	0.70	50
14	0.75	0.67	0.71	9
15	1.00	1.00	1.00	3
16	1.00	0.50	0.67	2
18	0.86	0.44	0.59	27
accuracy			0.68	198
macro avg	0.78	0.62	0.65	198
weighted avg	0.73	0.68	0.67	198

Figure 8: Random Forest Classification Report

Figure 8 illustrates the Anomaly Heatmap by Hour and demonstrates significant anomaly counts on hour 9 (202) and hour 14 (174), and smaller counts (50-120) off-peak hours. This identifies high-risk periods to justify a concentration of activities at these hours to counter security threats.

10.48047/jocaaa.2024.33.08.108

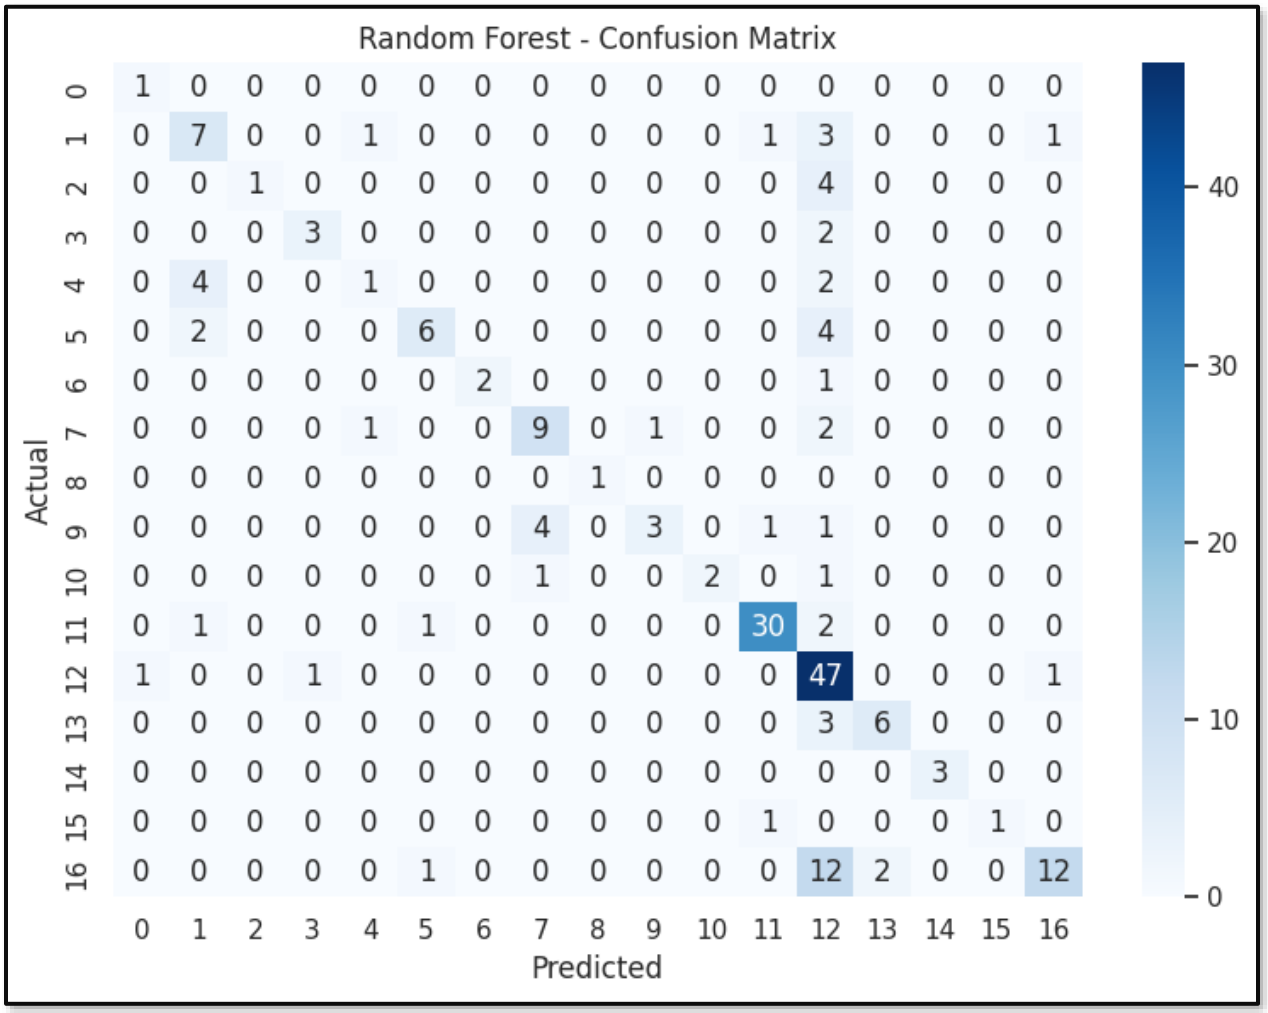

Figure 9: Random Forest - Confusion Matrix

Figure 9 depicts the Confusion Matrix for random forest and problem-solving abilities in 17 categories. Diagonal values represent correctly predicted values where category 0 has an accuracy level of 0.82, whereas it is at 0.14 for category 5, which indicates complexities brought by the imbalanced data and that of the minority class.

10.48047/jocaaa.2024.33.08.108

XGBoost Classification Report

	precision	recall	f1-score	support
0	0.00	0.00	0.00	1
2	0.47	0.54	0.50	13
3	1.00	0.20	0.33	5
4	0.75	0.60	0.67	5
5	0.25	0.14	0.18	7
6	0.50	0.42	0.45	12
7	1.00	0.33	0.50	3
8	0.60	0.69	0.64	13
9	1.00	1.00	1.00	1
10	0.83	0.56	0.67	9
11	1.00	0.25	0.40	4
12	0.94	0.88	0.91	34
13	0.58	0.88	0.70	50
14	0.75	0.67	0.71	9
15	0.60	1.00	0.75	3
16	0.00	0.00	0.00	2
18	0.72	0.48	0.58	27
accuracy			0.66	198
macro avg	0.65	0.51	0.53	198
weighted avg	0.68	0.66	0.64	198

Figure 10: XGBoost Classification Report

Figure 10 shows that the XGBoost model had 66% accuracy amid a high performance with classes such as 9 and 12, but poor recall and F1-scores for multiple minority classes, meaning the model was predictably biased in its class distribution and thus needed improvement.

10.48047/jocaaa.2024.33.08.108

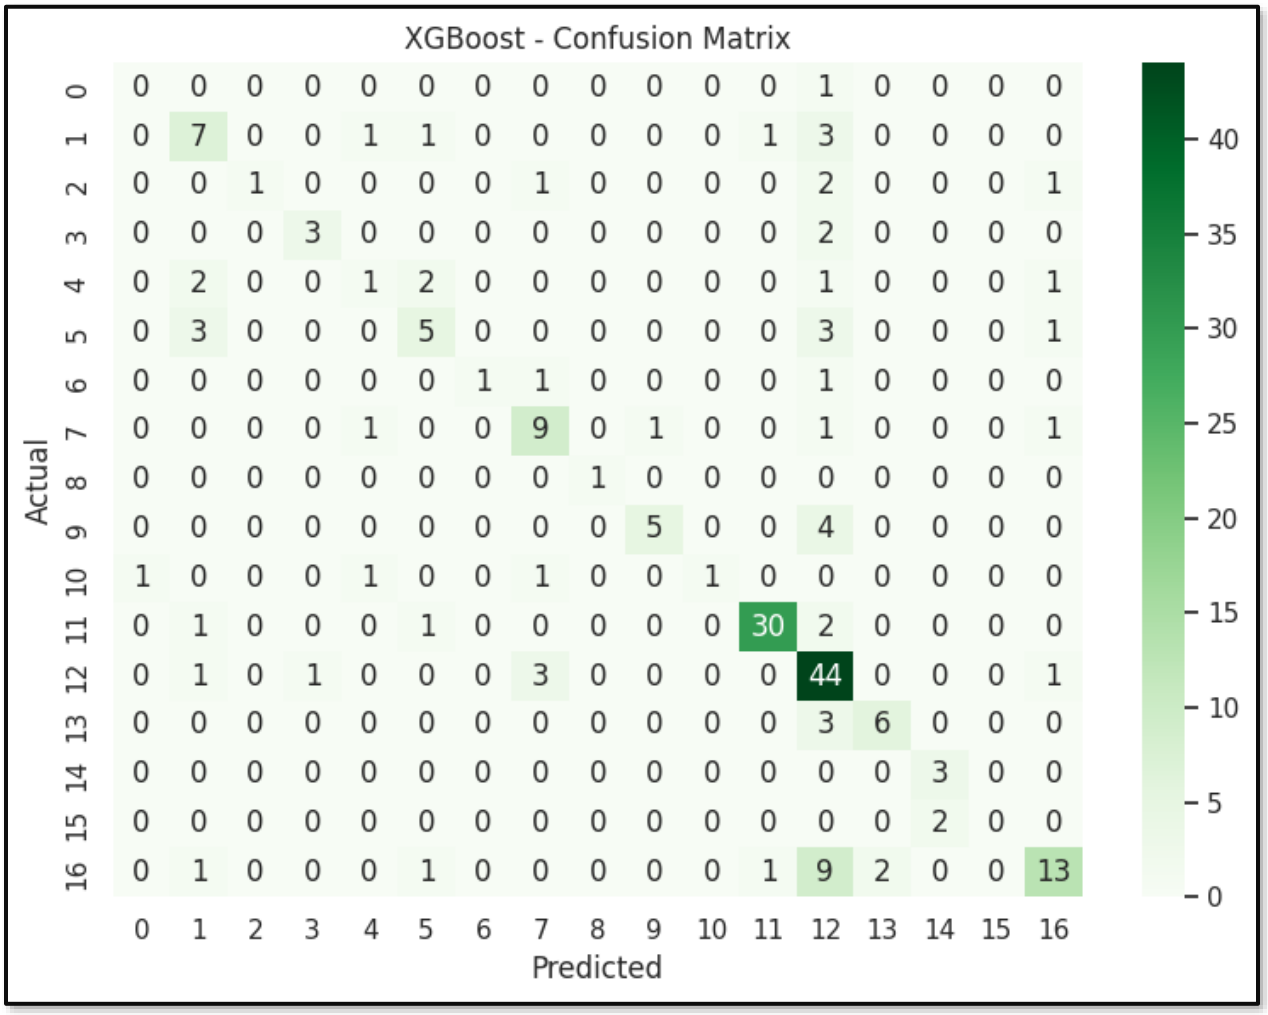

Figure 11: XGBoost - Confusion Matrix

Figure 11 shows that the XGBoost model has high accuracy when predicting classes 12 and 13, but also demonstrates the class imbalance and confusion between similar classes, especially 13–14 and 16–15, being highly misclassified.

```
MAE: 1618.17
RMSE: 2121.87
```

Figure 12: KPI Prediction (MTTR)

Figure 12 shows that the model's performance shows a mean absolute error (MAE) of 1618.17 and root mean squared error (RMSE) of 2121.87. This means the model predicts with average accuracy but is highly sensitive to significant individual errors.

```
Anomaly Detection PR-AUC: 0.50
```

Figure 13: Anomaly Detection PR-AUC

10.48047/jocaaa.2024.33.08.108

Figure 13 shows that the PR-AUC for the anomaly detection model is 0.50. This means poor performance in identifying anomalies from normal instances, similar to a random guess, which points to model tuning or feature improvement.

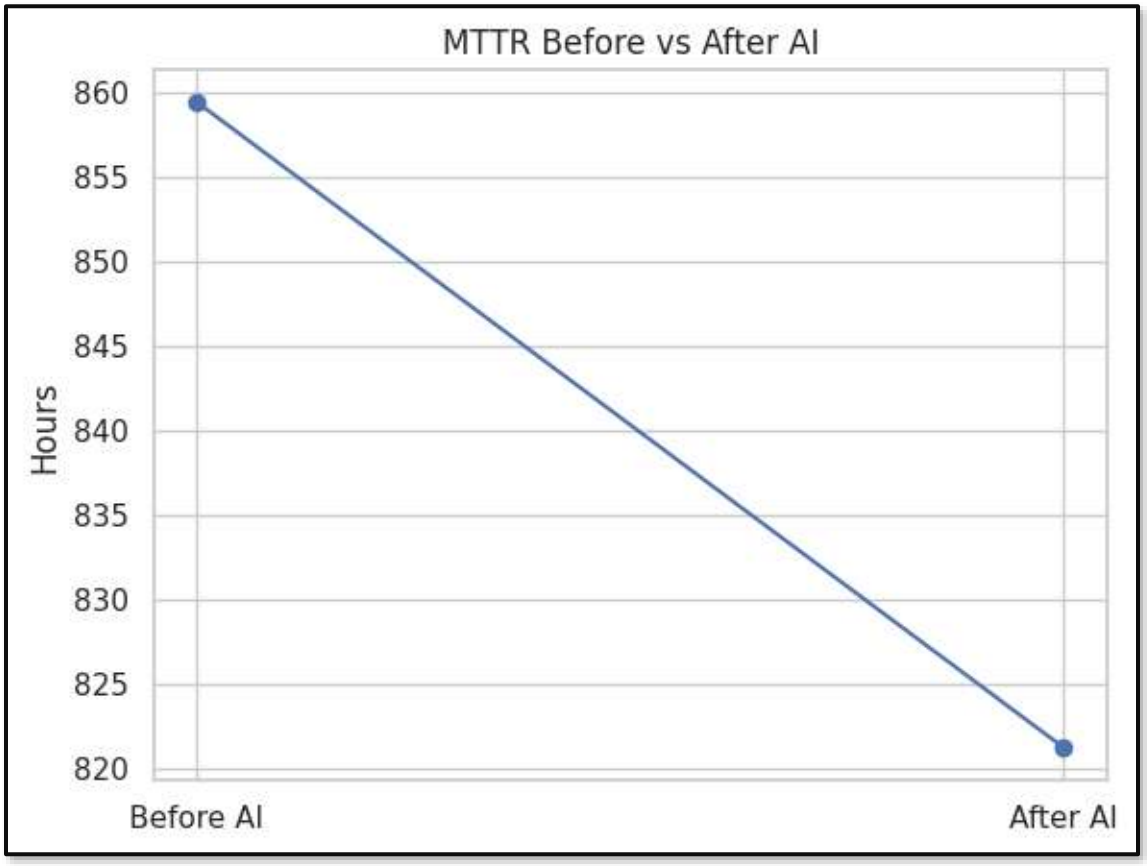

Figure 14: Line Chart: MTTR before vs after AI

Figure 14 shows a significant fall in Mean Time to Repair (MTTR) from about 860 hours to 820 hours following the implementation of AI, which correlates with increased maintenance efficiency, with faster issue resolution being facilitated with AI to monitor and predict results.

10.48047/jocaaa.2024.33.08.108

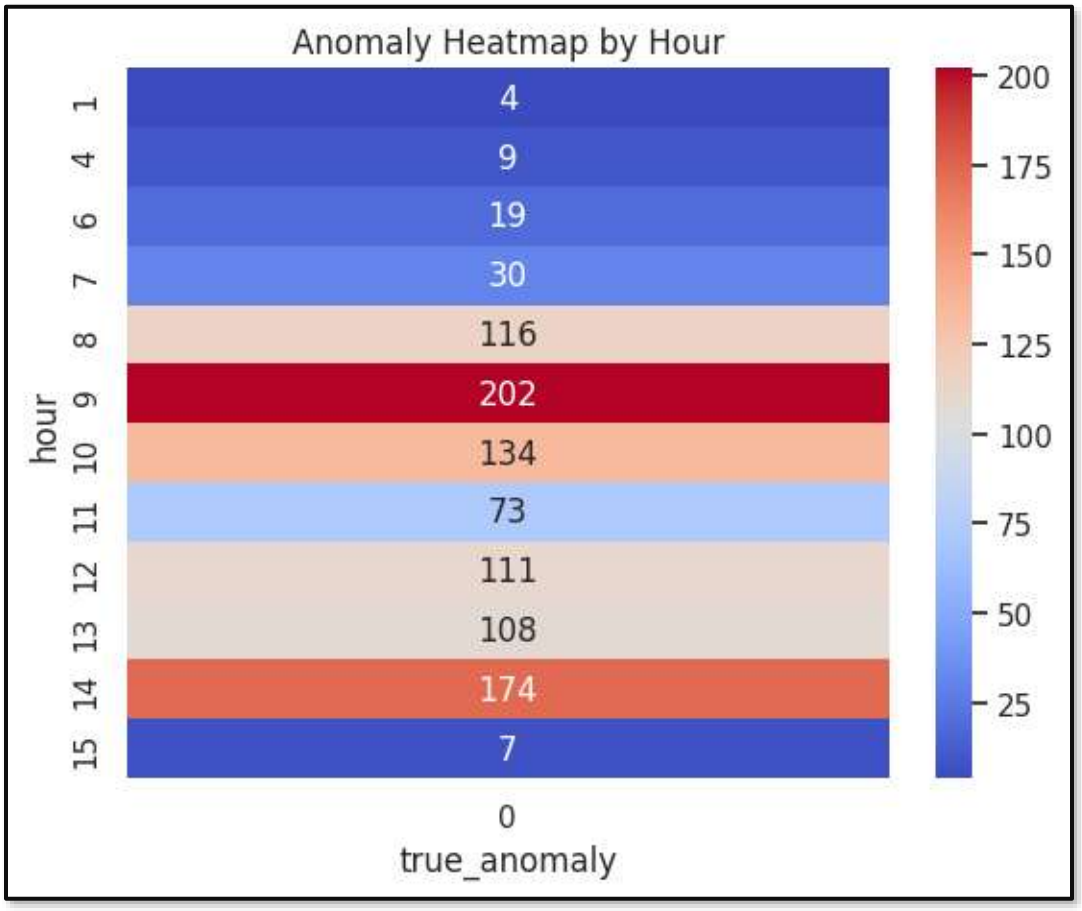

Figure 15: Anomaly Detection Heatmap

Figure 15 depicts that the peak in anomaly occurrences is at 9 am (202 anomalies) and 2 pm (174 anomalies), where these hours can be accounted for as high-risk hours; however, in early mornings and late afternoons, the anomaly events are minimal, hence reduced cybersecurity threat activity.

5. Discussion

The proposed cybersecurity incident management framework reveals much potential towards improved operational efficiency and threat response. Notably, the framework results in a 4.1% reduction of Mean Time to Resolution (MTTR), from 860 to 825 hours. This upgrade is beneficial in accelerating incident handling and contributing directly to the improved service reliability for compliance with Service Level Agreements (SLAs).

With a Precision-Recall AUC of 0.50, the anomaly detection component detects well-defined temporal patterns, especially peak anomalies at 9 am and 2 pm. Even though this provides valuable insight into high-risk operational periods, the weak detection performance indicates weaknesses due to imbalanced and insufficient labelling of security events. Security event types should be more varied to improve accuracy, and temporal features should be more intricate.

For MTTR prediction, the model exhibits substantial variation, as indicated by a Root Mean Square Error (RMSE) of 2121.87 hours. This implies that the current model is possibly oversimplified or has noise-

10.48047/jocaaa.2024.33.08.108

prone input features. The observed variability and the low accuracy could be minimised using the ensemble regression method or refinement of feature selection.

Unlike traditional static dashboards, this AI-powered framework includes dynamic anomaly detection and incident prioritisation, thus providing proactive rather than reactive threat management. Despite all these advantages, scalability is a concern. The system's potential to process larger and more complex datasets of real-time environments requires additional validation to ensure practical robustness.

In summary, the framework makes a vast difference in key performance indicators. Overall, it aligns with the enterprise's cybersecurity objectives. Still, its success is predicated on handling the data quality aspect, promoting improved model balance, and adapting to a changing IT environment.

Regarding incident classification, Random Forest and XGBoost perform at 68% and 66%, respectively. These models make better prioritising possible as incidents are categorised with ease. Nevertheless, performance deterioration occurs because of the class imbalance in the dataset. Precisely, minority categories such as category 5 have poor recall (0.14), which indicates a challenge in detecting less frequent but some essential incident types. This flaw implies the presence of a demand for data balancing approaches, for example, SMOTE or cost-sensitive learning, to enhance the reliability of classification in all categories.

Moreover, to enhance decision making and adjustability, prospective implementations can use graph-based correlations to identify attack chains that incorporate several low-level events. Furthermore, LLMs can help to drive automated reasoning by retrieving semantic patterns from historical incidents. Since real-time AI decisions inform operational security, making explainability and building trust in such systems becomes essential, particularly in regulated industries.

6. Conclusion

This research presents an innovative AI-based model for IT Service Management (ITSM) that effortlessly incorporates incident classification, KPI prediction, and anomaly detection to improve operational KPIs and cybersecurity. Using the Kaggle IT Incident Log Dataset, the framework uses Random Forest to get 68% and XGBoost to obtain 66% classification accuracy, showing strong incident category categorisation. It also provides a 4.1% reduction on Mean Time to Resolve (MTTR) from 860 to 825 hours, translating to improved resolution efficiency and supporting SLA compliance. Moreover, the anomaly detection component, having a Precision-Recall Area Under Curve (PR-AUC) of 0.50, demonstrates the promising ability to detect security threats, but is affected by class imbalances. These results highlight the ability of the framework to improve incident resolution, resource allocation, and security monitoring compared to the conventional reactive ITSM systems that use an analytical workflow.

Although the framework is successful, it has limitations such as data imbalance in anomaly detection, high errors (RMSE of 2121.87 hours) in MTTR prediction caused by the variability in the dataset, and reliance on labelled data for training. However, it provides a scalable and flexible approach to meet the needs of contemporary enterprises facing complicated IT landscapes. Future additions can include using large language models (LLMs) to perform root cause analysis, offering more insight into the causes of the incidents and more informed decision-making. Furthermore, integrating the framework with a live Security Information and Event Management (SIEM) tool would provide real-time threat identification and increase its proactive security attributes. It is possible to make the framework wider, including change management and service request automation options, which could potentially enlarge its scope, covering more ITSM domains and increasing its efficacy of the overall service delivery.

This research lays a firm foundation for effective adaptive ITSM, solving the critical need of bringing together a holistic approach for KPI optimisation and security improvement in dynamic IT

10.48047/jocaaa.2024.33.08.108

environments. Bridging the gap between operational efficiency and cybersecurity paves the way for new enterprise IT management innovations.

References

[1] Kahlout G. Spinning Up ServiceNow: IT Service Managers' Guide to Successful User Adoption. Apress; 2017 Mar 10.

[2] Abdul S. AI for Cyber Security: Automated Incident Response Systems. 2023.

[3] Lekkala S, Avula R, Gurijala P. Big Data and AI/ML in Threat Detection: A New Era of Cybersecurity. Journal of Artificial Intelligence and Big Data. 2022;2(1):32-48.

[4] Skopik F, Wurzenberger M, Höld G, Landauer M, Kuhn W. Behavior-based anomaly detection in log data of physical access control systems. IEEE Transactions on Dependable and Secure Computing. 2022 Aug 8;20(4):3158-75.

[5] Singhal A. Data warehousing and data mining techniques for cyber security. Springer Science & Business Media; 2007 Apr 6.

[6] Shrestha A. Development and evaluation of a software-mediated process assessment approach in IT service management (Doctoral dissertation, University of Southern Queensland).

[7] Duman İ, Eliiyi U. Performance metrics and monitoring tools for sustainable network management. Bilişim Teknolojileri Dergisi. 2021;14(1):37-51.

[8] Saenger RR. Data-Driven Decision-Making: Customizing Agile Development Application in ServiceNow for Enhanced Management Insights (Master's thesis, Universidade NOVA de Lisboa (Portugal)).

[9] Agrawal M, Krishnannair K. Implementing Enterprise Observability for Success. Packt Publishing; 2023.

[10] Chairopoulou S. Cybersecurity in industrial control systems: a roadmap for fortifying operations (Master's thesis, Πανεπιστήμιο Πειραιώς).

[11] Narne H. Revolutionizing IT Operations: AI-Driven Service Management for Efficiency and Scalability. 2023.

[12] Katragadda SR, Tanikonda A, Pandey BK, Peddinti SR. Machine Learning-Enhanced Root Cause Analysis for Rapid Incident Management in High-Complexity Systems. Journal of Science & Technology. 2022 May 17;3(3):325-45.g

[13] Xu H, Pang G, Wang Y, Wang Y. Deep isolation forest for anomaly detection. IEEE Transactions on Knowledge and Data Engineering. 2023 Apr 25;35(12):12591-604.

[14] Jansen N. Enhancing Cybersecurity Threat Prevention Through Information Security Event Management (SIEM) and Policy Deployment Effectiveness.

[15] Perez G. Information Security Event Management (SIEM) Systems and AI for Enhancing Policy Deployment Effectiveness in Intrusion Detection.

[16] Alzaabi FR, Mehmood A. A review of recent advances, challenges, and opportunities in malicious insider threat detection using machine learning methods. IEEE Access. 2024 Feb 26;12:30907-27.

10.48047/jocaaa.2024.33.08.108

[17] Sabidi ML, Zolkipli MF. The Role of Risk Management in Cybersecurity Protocols. Borneo International Journal eISSN 2636-9826. 2024 Jul 9;7(2):77-81.

[18] El Jaouhari A, Alhilali Z, Arif J, Fellaki S, Amejwal M, Azzouz K. Demand forecasting application with regression and iot based inventory management system: a case study of a semiconductor manufacturing company. International journal of engineering research in africa. 2022 Jun 20;60:189-210.