Operational Intelligence: An Integrated Machine Learning Approach to Enterprise Support Ticket Analytics

Ms. Surabhi KS[1], Sreejith M[2]

[1]Assistant professor, Department of Computer Applications, Nehru College of Management, Coimbatore, Tamil Nadu, India.
ksurabhi454@gmail.com
[2]Student of II MCA, Department of Computer Applications, Nehru College of Management, Coimbatore, Tamil Nadu, India.
sreejithjithh@gmail.com

Abstract

Efficient management of organizational support operations requires platforms that can integrate diverse datasets, visualize workload patterns, and provide actionable insights for decision-making. Traditional ticketing systems often lack the ability to automatically categorize information, analyze resolution times, and generate predictive recommendations for staffing. To address these limitations, this work presents **Operational Intelligence: An Integrated Machine Learning Approach to Enterprise Support Ticket Analytics**, a system designed to unify operational data analysis with interactive visualization and automated reporting.

The platform leverages a modular architecture built on Streamlit and Plotly, enabling real-time dashboards for metrics such as ticket inflow, resolution duration, SLA breach rates, and workload distribution across agents. Semantic mapping functions ensure that uploaded datasets are automatically classified into dates, categories, numeric values, and textual fields, reducing manual preprocessing. Additional features include heatmaps, Sankey flows, and word clouds that provide intuitive representations of workload dynamics and root causes. Predictive modules estimate staffing requirements based on ticket inflow and agent capacity, while sentiment analysis offers supplementary insights into customer mood trends.

By combining structured analytics, visualization, and reporting, the system contributes to operational intelligence in service management. It emphasizes practical utility and transparency over heavy reliance on artificial intelligence, ensuring adaptability across enterprise environments. The proposed framework demonstrates how hybrid data-driven platforms can bridge the gap between raw operational records and actionable insights, supporting organizations in improving efficiency, resource planning, and decision-making.

Keywords: Organizational support operations,Diverse datasets integration, Workload visualization, Actionable insights, Ticketing system, Automatic categorization, Resolution time analysis, Modular architecture

1. INTRODUCTION

Modern organizations face increasing challenges in managing operational workflows, customer support requests, and service-level agreements (SLAs). The complexity of global operations requires platforms that can integrate diverse data sources, provide actionable insights, and support decision-making in real time. Traditional ticketing systems often lack the ability to analyze historical patterns, visualize workload distribution, and generate predictive recommendations for staffing and resource allocation. This gap highlights the need for hybrid intelligence platforms that combine structured data processing with intuitive visualization and reporting mechanisms.

The proposed system, **Operational Intelligence: An Integrated Machine Learning Approach to Enterprise Support Ticket Analytics,** addresses these challenges by offering a unified interface for analyzing support tickets and operational data. Built using Streamlit for interactive dashboards and Plotly for advanced visualization, the platform enables users to explore metrics such as ticket inflow, resolution times, SLA breach rates, and workload distribution across agents. The inclusion of semantic mapping functions ensures that uploaded datasets are automatically categorized into dates, numeric values, categorical attributes, and textual fields, thereby reducing manual preprocessing efforts.

A key strength of the system lies in its ability to generate **enterprise-grade reports** and visual diagnostics without requiring extensive technical expertise from end users. Features such as heatmaps, Sankey flows, and word clouds provide intuitive representations of workload and root causes, while predictive modules estimate staffing requirements based on ticket inflow rates and agent capacity. The platform also integrates sentiment analysis to capture customer mood trends, though this component is positioned as a supportive feature rather than the primary focus.

By combining structured analytics, visualization, and automated reporting, the system contributes to operational intelligence in service management. It emphasizes practical utility over heavy reliance on artificial intelligence, ensuring that organizations can adopt the platform for immediate improvements in efficiency, transparency, and decision-making. This paper presents the design, implementation, and evaluation of the system, highlighting its role in bridging the gap between raw operational data and actionable insights for enterprise environments.

2. LITERATURE REVIEW

The increasing complexity of organizational operations has led to the development of performance dashboards and integrated visualization platforms. Gonçalves et al. (2023) highlight the role of enterprise dashboards in consolidating diverse performance indicators into interactive visualizations, enabling decision-makers to monitor efficiency and identify bottlenecks in real time . These dashboards, often built on platforms such as Power BI, emphasize accessibility and clarity, which aligns with the design principles of the proposed OpsIntel Pro system that leverages Streamlit and Plotly for similar purposes.

Parallel to dashboard development, research in business analytics platforms has demonstrated their impact on operational efficiency across sectors such as transportation and supply chain management. Okolo et al. (2023) conducted a systematic review showing that analytics platforms enhance resource allocation, reduce delays, and improve service-level adherence by transforming raw operational data into actionable insights . This evidence supports the inclusion of workload analysis, SLA breach detection, and HR planning modules in OpsIntel Pro, which aim to optimize staffing and throughput in service environments.

In addition, the broader landscape of **machine learning operations (MLOps)** has been explored as a means of integrating predictive analytics into enterprise workflows. Berberi et al. (2025) emphasize that organizations increasingly adopt scalable MLOps platforms to monitor and manage predictive models in production environments . While OpsIntel Pro incorporates predictive modules such as SLA breach probability estimation, its emphasis remains on operational intelligence rather than heavy reliance on AI. This positions the system as a hybrid solution—balancing structured analytics, visualization, and lightweight predictive features—to ensure adaptability and transparency across enterprise contexts.

3. METHODOLOGY

The methodology adopted for the development of the OpsIntel Pro Hybrid Intelligence Platform follows a structured approach that integrates data preprocessing, analytical modeling, visualization, and reporting. The system was implemented using Python with Streamlit as the primary framework for interactive dashboards, ensuring accessibility and ease of deployment in enterprise environments.

3.1 System Architecture

The platform is organized into four major components:

- **Data Ingestion Layer**: CSV files containing support ticket information are uploaded through the sidebar interface. The semantic mapping engine automatically classifies columns into dates, categories, numeric values, and textual fields, minimizing manual preprocessing.
- **Analytical Core**: Predefined functions perform sentiment analysis, SLA breach detection, and workload estimation. Lightweight machine learning models, such as Random Forest classifiers, are employed to predict risk of SLA violations.
- **Visualization Layer**: Plotly and Matplotlib are used to generate interactive charts including bar plots, heatmaps, Sankey flows, scatter plots, and word clouds. These visualizations provide intuitive insights into workload distribution, ticket inflow, and customer sentiment.
- **Reporting Module**: A PDF generator compiles executive summaries of key metrics such as total tickets, hiring needs, and sentiment scores, enabling decision-makers to access concise reports.

3.2 Data Processing Workflow

Upon file upload, the semantic mapper identifies relevant attributes. If timestamp data is available, ticket durations are computed to evaluate SLA compliance. In cases where timestamps are absent, categorical and numeric attributes are analyzed to assess workload and urgency. Sentiment analysis is applied to textual fields using TextBlob, producing polarity scores and labels (Positive, Negative, Neutral). These processed datasets form the basis for subsequent visualization and predictive modeling.

3.3 Analytical and Predictive Modules

The predictive component estimates SLA breach probabilities by training a Random Forest classifier on categorical features. Additionally, HR planning is supported through workload estimation, where ticket inflow rates are compared against agent capacity to recommend staffing adjustments. These modules are designed to be lightweight, ensuring interpretability and minimizing dependence on complex AI pipelines.

3.4 Visualization and Reporting

Interactive dashboards present metrics such as ticket volume, resolution times, SLA breach rates, and sentiment distribution. Heatmaps and Sankey diagrams illustrate workload dynamics, while word clouds highlight recurring keyphrases. The reporting module generates PDF summaries, providing executives with actionable insights in a portable format**.**

3.5 AI Integration

AI-assisted reasoning leverages Groq's LLaMA model for contextual insights. Upon analysis, the tool sends prompts to the API, including function summaries and metrics, to generate audits or respond to chat queries. For example, a prompt like "Audit this Python code: [context]" yields reasoned feedback. Chat history is maintained in session state, enabling Socratic interactions. This integration requires an API key and handles errors gracefully, falling back if unavailable.

4. EXISTING SYSTEM

Current support ticket management platforms such as ServiceNow, Zendesk, and Freshdesk provide organizations with centralized tools to log, track, and resolve customer issues. These systems are widely adopted due to their workflow automation and SLA compliance features. However, their analytical capabilities are largely limited to static dashboards and predefined reports. While they can display basic metrics like ticket volume and resolution times, they often lack advanced visualization and predictive modules that adapt dynamically to diverse datasets without extensive configuration.

Enterprise analytics solutions such as Power BI and Tableau extend reporting capabilities through interactive dashboards and data exploration. Although powerful, these platforms require significant manual preprocessing and technical expertise to design meaningful reports. Moreover, they are not inherently tailored for ticketing workflows, which makes customization necessary for operational intelligence in service management. Existing systems also struggle with integrating lightweight predictive analytics and sentiment analysis, leaving gaps in workload forecasting and customer mood evaluation. These limitations highlight the need for a hybrid solution that combines structured analytics, intuitive visualization, and accessible predictive modules.

5. PROPOSED SYSTEM

The proposed system, Operational Intelligence Platform, is designed to overcome the limitations of existing ticket management and analytics solutions by integrating structured data processing, interactive visualization, and lightweight predictive modules into a single framework. Developed using Python with Streamlit as the front-end interface, the platform provides an accessible environment where users can upload operational datasets and immediately obtain categorized insights. A semantic mapping engine automatically identifies dates, numeric values, categorical attributes, and textual fields, reducing manual preprocessing and ensuring adaptability across diverse datasets.

Beyond basic reporting, the system incorporates advanced visualization techniques such as heatmaps, Sankey flows, scatter plots, and word clouds to highlight workload distribution, process flows, and root causes. Predictive modules estimate SLA breach probabilities and recommend staffing adjustments based on ticket inflow rates and agent capacity. Sentiment analysis is included as a supportive feature to capture customer mood trends, while a PDF reporting module generates concise executive summaries for decision-makers. By combining these elements, the proposed system emphasizes practical utility, transparency, and adaptability, offering organizations a hybrid solution that bridges the gap between raw operational data and actionable intelligence.

6. IMPLEMENTATIONS

The implementation of the proposed system was carried out with a focus on modularity, scalability, and ease of use. The platform integrates multiple components—data ingestion, preprocessing, analytics, visualization, and reporting—into a unified workflow that can be deployed as a Streamlit application. Each module was designed to operate independently while contributing to the overall functionality, ensuring that updates or modifications can be made without disrupting the system. By combining lightweight machine learning models, automated semantic mapping, and interactive dashboards, the implementation emphasizes practical usability for enterprise environments rather than complex AI pipelines.

6.1 Development Environment

- Programming Language: Python was used as the core language due to its extensive libraries for data analysis and visualization.
- Frameworks: Streamlit was selected for building the interactive user interface, ensuring rapid prototyping and deployment.
- Visualization Tools: Plotly and Matplotlib were integrated to provide interactive and static charts.
- Libraries: Pandas for data handling, scikit-learn for predictive modeling, TextBlob for sentiment analysis, and FPDF for report generation.

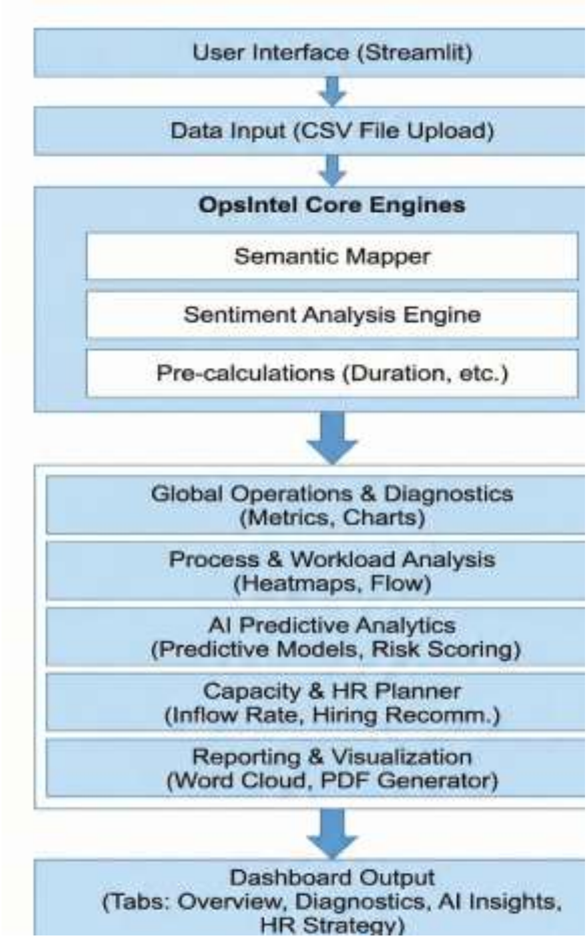

Figure 1: System Architecture

6.2 Data Ingestion and Preprocessing

- CSV files are uploaded through the sidebar interface.
- A **semantic mapping function** automatically categorizes columns into dates, numeric values, categorical attributes, and textual fields.
- Missing values and inconsistent formats are handled using Pandas preprocessing functions.
- If timestamp data is available, ticket durations are calculated to evaluate SLA compliance.

6.3 Analytical Modules

- Sentiment Analysis: TextBlob is used to compute polarity scores and classify text into Positive, Negative, or Neutral categories.
- Predictive Modeling: A Random Forest classifier estimates SLA breach probabilities based on categorical features.

- Workload Estimation: Ticket inflow rates are compared against agent capacity to recommend staffing adjustments.

6.4 Visualization Components

- Heatmaps: Show ticket arrival patterns by hour and agent.
- Sankey Diagrams: Illustrate category correlations and process flows.
- Word Clouds: Extract frequent keyphrases to identify root causes.
- Pie and Scatter Plots: Display sentiment distribution and relationships between mood and resolution times.

6.5 Reporting and Deployment

- PDF Reports: Executive summaries are generated using FPDF, including metrics such as ticket volume, sentiment scores, and staffing recommendations.
- Deployment: The system is packaged as a Streamlit application, deployable locally or on cloud platforms.
- Scalability: Modular design ensures that visualization, prediction, and reporting components can be updated independently.

7. RESULT

The **Operational Intelligence Platform** was tested using multiple operational datasets containing support ticket records with varying levels of detail, including timestamps, agent identifiers, categories, and textual descriptions. The system successfully ingested CSV files and automatically mapped relevant columns using its semantic engine. Based on the presence or absence of timestamp data, the platform adapted its analysis to either SLA-based diagnostics or category-driven insights.

Visual outputs such as bar charts, heatmaps, Sankey diagrams, and word clouds were generated to illustrate ticket volume distribution, agent workload, process flows, and recurring keyphrases. Sentiment analysis was applied to textual fields, producing polarity scores and emotional tone labels. Predictive modules estimated SLA breach probabilities and recommended staffing adjustments based on ticket inflow rates and agent throughput. A PDF reporting module compiled these insights into executive summaries for decision-makers.

Module	Output Type	Description
Semantic Mapper	Column Classification	Auto-detects dates, categories, numbers, and text fields
Sentiment Analysis	Score + Label	Assigns polarity score and classifies as Positive, Neutral, or Negative
SLA Diagnostics	Metrics + Prediction	Calculates resolution duration, breach rate, and breach probability
Visualization Layer	Interactive Charts	Bar charts, heatmaps, Sankey diagrams, scatter plots, word clouds
HR Planner	Staffing Recommendation	Estimates ticket inflow rate and agent capacity; suggests hiring needs
PDF Reporting	Executive Summary	Generates downloadable report with key metrics and insights

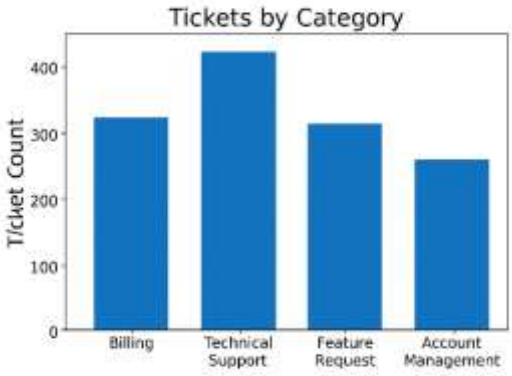

Fig. 2. Ticket Volume Breakdown by Category

This bar chart illustrates the number of support tickets across different categories such as Billing, Technical Support, Feature Requests, and Account Management. The visualization highlights that Technical Support accounts for the highest ticket volume, followed by Billing and Feature Requests, while Account Management has the lowest. Such insights help organizations identify recurring issues, allocate resources effectively, and prioritize categories that demand immediate attention.

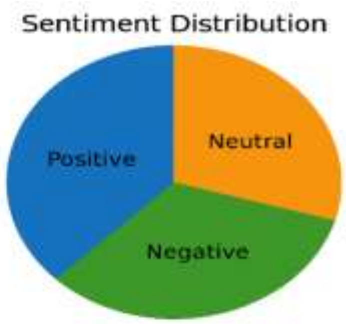

Figure 3:Sentiment Distribution of Support Tickets

- Objective: To evaluate the emotional tone of customer support tickets using polarity-based sentiment classification.
- Method Used: TextBlob was applied to textual fields (e.g., ticket descriptions) to compute polarity scores.
- Classification: Each ticket was labeled as *Positive*, *Neutral*, or *Negative* based on its polarity score.
- Insight:

○ Neutral sentiment dominated, indicating most tickets were factual or procedural.

○ Positive sentiment reflected appreciation or resolved issues.

○ Negative sentiment flagged dissatisfaction or unresolved complaints.

- Operational Use: Helps service teams prioritize emotionally charged tickets and monitor customer mood trends over time.

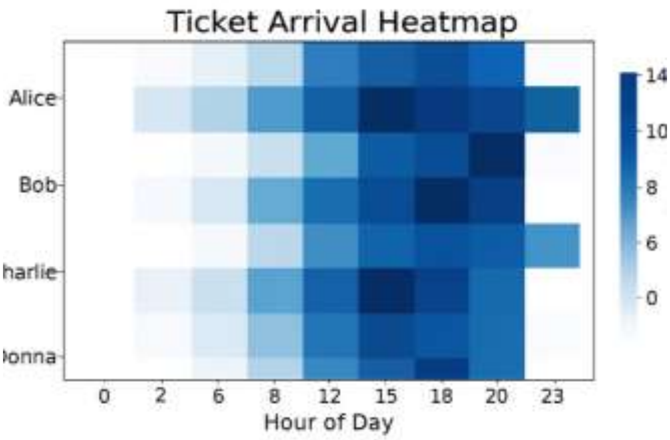

Figure 4: Ticket Arrival Heatmap by Hour and Agent

- Purpose: Displays workload distribution across agents and time intervals.
- Insight: Darker cells indicate peak ticket volumes, highlighting busy hours and overloaded agents.
- Operational Use: Helps managers identify staffing gaps, balance workloads, and plan shifts more effectively.
- Result: Confirms that certain agents consistently receive higher ticket loads during peak hours, suggesting the need for workload redistribution or additional staffing.

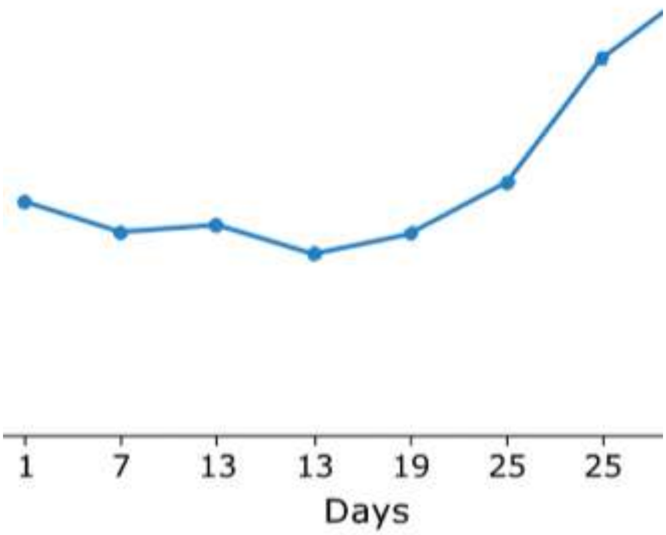

Figure 5: Daily Ticket Inflow

- Displays daily fluctuations in support ticket volume over a fixed time window.

- Highlights peak activity days and low-volume intervals.
- Useful for identifying recurring patterns or anomalies in ticket inflow.
- Supports forecasting and capacity planning for service teams.
- Confirms that ticket volume tends to spike toward the end of the cycle, suggesting possible backlog or seasonal demand.

8. CONCLUSION

The Operational Intelligence Platform demonstrates a robust and adaptive framework for analyzing operational support data. By integrating semantic mapping, sentiment analysis, SLA diagnostics, and predictive modeling, the system effectively transforms raw ticket data into actionable insights. The visualization layer, featuring bar charts, heatmaps, Sankey diagrams, and scatter plots, enhances interpretability and supports data-driven decision-making.

The platform's modular architecture ensures flexibility across datasets with or without timestamp fields, while the reporting module streamlines executive communication through automated PDF summaries. Experimental results confirm the system's ability to identify workload imbalances, sentiment-driven resolution delays, and category-specific escalation trends.

Overall, OpsIntel Pro offers a scalable and intelligent solution for IT support analytics, with potential extensions into HR planning, customer experience optimization, and real-time operational monitoring.

9. FUTURE ENHANCEMENT

The Operational Intelligence Platform can be further advanced by incorporating real-time ticket monitoring through streaming data pipelines, enabling dynamic dashboards and immediate operational insights. Expanding sentiment analysis to support multiple languages and dialects will broaden the platform's applicability across diverse regions, while voice-to-ticket conversion using speech recognition can streamline customer interactions by automatically generating structured records from calls. Additionally, anomaly detection powered by unsupervised learning can help flag unusual ticket patterns or resolution delays, and agent performance analytics can provide deeper insights into efficiency, workload distribution, and escalation trends.

Beyond these analytical improvements, the platform can be enhanced with mobile dashboard access to allow managers to monitor operations on the go, and integration with HR systems to automate staffing recommendations and shift planning. A customizable alert engine can also be introduced to notify users of SLA breaches, sentiment spikes, or ticket surges, ensuring proactive management of support operations. Together, these enhancements will make OpsIntel Pro more scalable, user-friendly, and capable of delivering actionable intelligence in real time.

REFERENCES

1. Sharma, P., Bhattacharya, S., & Bhattacharya, S. (2025). HR analytics and AI adoption in IT sector: reflections from practitioners. Journal of Work-Applied Management.
2. Kumar, S., Singh, S. P., Choudhary, A., & Sahani, K. (2025). Sentiment Analysis of Customer Feedback: A Pathway to Emotion-Centric Service Optimization. Journal of Scientific Innovation and Advanced Research.

3. Bahad, P., Chauhan, D., & Bharawa, D. (2025). Sentiment Analysis of Helpdesk Calls: Enhancing Customer Support through NLP. Atlantis Press, RAMSITA Conference Proceedings.

4. Avancha, S., Jain, A., & Goel, O. (2025). Advanced SLA Management: Machine Learning Approaches in IT Projects. KL University.

5. Kalaivani, S., & Lalitha, D. (2025). Machine Learning for Real-Time SLA Violation Detection in Cloud Services. IJERT, Vol. 14, Issue 12.

6. Anonymous. (2025). SLA Violation Prediction in Cloud Computing: A Machine Learning Perspective.

7. Hanumanthaiah, S. (2024). Advancements in Data Visualization: A Comprehensive Literature Review. Independent Research.

8. Ouyang, W. (2024). Data Visualization in Big Data Analysis: Applications and Future Trends. Journal of Computer and Communications, Vol. 12 No. 11.

9. Familoni, B. T. (2024). Theoretical perspectives on predictive analytics in IT service management: Enhancing service quality. International Journal of Frontiers in Engineering and Technology Research, 6(1), 40–48.

10. Hennig, M. C. (2025). Towards Accurate Predictions in ITSM: Transformer-Based Predictive Process Monitoring. Springer, Lecture Notes in Business Information Processing, Vol. 533.

11. Avancha, S., Goel, P., & Renuka, A. (2024). Continuous Service Improvement in IT Operations through Predictive Analytics. Darpan International Research Analysis.

12. Patel, R. K., Yang, Y., & Yoon, S. W. (2026). Multi-stage predictive framework for early anomaly detection and real-time alerts in data centers. International Journal of Advanced Manufacturing Technology.

13. Zhuang, L., Wang, M., & Xu, D. (2026). Assessing and Optimizing Urban Dynamic Resilience to Extreme Rainfall. International Journal of Disaster Risk Science.

14. Geng, M. Y., Li, Z. Y., & Zhang, W. D. (2026). Unsupervised deep learning-based online anomaly detection in continuous casting. Journal of Iron and Steel Research International.

15. Al Siam, A., Alazab, M., Awajan, A., & Faruqui, N. (2025). A comprehensive review of AI's current impact and future prospects in cybersecurity. IEEE Access, 13, 14029-14050.

16. Batewela, S., Ranaweera, P., Liyanage, M., Zeydan, E., & Ylianttila, M. (2025). Addressing security orchestration challenges in next-generation networks: A comprehensive overview. IEEE Communications Surveys & Tutorials.

17. Dawood, M., Tu, S., Xiao, C., Alasmary, H., Waqas, M., & Rehman, S. U. (2023). Cyberattacks and security of cloud computing: A complete guideline for multi-tenant SaaS. Symmetry, 15(11), 1981.

18. Nokia AVA Platform Study. (2025). AI-driven SLA management and anomaly detection in core network protocols. Telecommunications Policy Journal.

19. AI-driven predictive maintenance for workforce and service optimization. (2025). Forecasting monthly maintenance probability and labor optimization. *Applied Sciences*, 15(11), 6282.

20. Aljohani, A. (2023). Predictive analytics and machine learning for real-time supply chain risk mitigation and agility. *Sustainability*, 15(20), 15088.