
| RESEARCH ARTICLE

Transforming Incident Management: Leveraging Artificial Intelligence for Enhanced Detection, Classification, and Resolution

Naveen Kumar Chandu

Jawaharlal Nehru Technological University, India

Corresponding Author: Naveen Kumar Chandu, **E-mail:** nchanduinbox@gmail.com

| ABSTRACT

This research explores how artificial intelligence can fundamentally change the way enterprises handle incident management. As companies face increasingly complex IT environments, traditional manual approaches to spotting, categorizing, and fixing problems have struggled to keep pace with operational demands. By bringing together AI capabilities like machine learning for spotting unusual patterns, natural language processing for understanding automated reports, and predictive analytics for assessing problem severity, we can dramatically improve how efficiently incidents are managed. This investigation examines how AI-based approaches enable systems to automatically generate incident reports, make smart severity assessments, calculate changing impacts in real-time, and create thorough documentation through meeting transcription and summarization. While highlighting the benefits of increased speed, better accuracy, and improved scalability, this work also addresses real-world implementation hurdles including data quality concerns, potential algorithmic biases, and the complexities of integrating with existing systems. Companies that successfully implement AI-powered incident management solutions stand to gain stronger operational resilience, faster problem resolution times, and happier customers, giving them a competitive edge in our increasingly digital business landscape.

| KEYWORDS

Artificial intelligence, incident management, anomaly detection, severity classification, automated documentation

| ARTICLE INFORMATION

ACCEPTED: 20 May 2025

PUBLISHED: 10 June 2025

DOI: 10.32996/jcsts.2025.7.5.108

1. Introduction to AI-Driven Incident Management

In today's rapidly evolving digital landscape, companies face extraordinary challenges in maintaining smooth operations across increasingly complex technology infrastructures. Managing incidents has emerged as a critical capability for enterprises, where system failures, performance slowdowns, and service interruptions can directly threaten business operations and customer relationships [1]. Traditional incident management approaches often buckle under the weight of modern challenges - the sheer number, speed, and variety of incidents occurring across today's distributed systems, cloud platforms, and interconnected applications simply overwhelm manual processes.

The growing complexity of enterprise technology has created an environment where human-led detection, categorization, and resolution of incidents struggle to meet operational demands. IT teams and system administrators regularly experience alert overload, struggle with inconsistent severity ratings, and face frustrating delays when depending entirely on human judgment to manage incidents [2]. These difficulties multiply in networked environments where system interdependencies create domino effects that prove challenging to track and document using traditional methods.

Artificial Intelligence offers a game-changing opportunity to transform how we practice incident management through intelligent automation and smart processing. By harnessing the power of machine learning algorithms, natural language understanding, and predictive analytics, companies can build incident management capabilities that autonomously detect

Copyright: © 2025 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (<https://creativecommons.org/licenses/by/4.0/>). Published by AI-Kindi Centre for Research and Development, London, United Kingdom.

problems, accurately classify severity levels, assess impacts as they evolve, and create comprehensive documentation [1]. These AI-powered approaches hold the promise of slashing detection and resolution times while bringing much-needed consistency and quality to incident handling procedures throughout the organization.

This paper examines how AI can be applied across every stage of the incident management lifecycle, from initial detection and report generation through severity assessment and impact documentation. We explore how machine learning models identify unusual patterns in system behavior, classify incidents based on learned patterns from past events, and continuously update impact assessments as situations develop [2]. Additionally, we investigate how AI-powered transcription and summarization tools can capture crucial insights from incident response meetings, preserving organizational knowledge and enabling continuous improvement. Our analysis covers both the theoretical foundations and practical implementation considerations, offering a thorough examination of AI's transformative role in modern enterprise incident management.

1.1 Evolution of Incident Management Practices

The practice of managing incidents has traveled a long road since its early days in IT service management frameworks. Early approaches relied heavily on manual detection and resolution processes, which inevitably led to slow response times and inconsistent handling approaches. As enterprise systems grew more sophisticated, the limitations of these manual methods became increasingly apparent, driving the need for more advanced ways to detect and resolve system disruptions [1]. The introduction of monitoring tools and alerting systems marked significant progress, yet these often created new problems with excessive false alarms and insufficient context for prioritization.

1.2 Challenges in Contemporary Incident Management

Today's technology landscape features cloud services, microservice architectures, and distributed systems that create intricate webs of dependencies, making root cause analysis and impact assessment exceptionally challenging [2]. Adding to this complexity, the accelerating pace of software deployment through DevOps and continuous integration practices introduces additional variables that must be considered during incident investigation. These factors, combined with growing expectations for near-perfect service availability, have pushed traditional incident management approaches to their breaking point.

1.3 The Promise of AI-Enabled Approaches

Artificial Intelligence brings revolutionary capabilities that address many fundamental limitations in traditional incident management. Machine learning techniques can process vast quantities of system telemetry data, identifying patterns and anomalies that might signal developing incidents before users notice any impact [1]. Natural language processing capabilities enable systems to extract meaningful insights from unstructured sources like user complaints, support tickets, and recorded meetings. Furthermore, AI-powered predictive models can forecast potential system failures based on historical patterns, enabling proactive intervention rather than reactive responses.

1.4 Scoping the AI Integration Framework

This paper offers a thorough examination of how AI technologies can be woven throughout the incident management lifecycle to improve detection, classification, documentation, and resolution capabilities. We analyze both the theoretical foundations of AI applications in incident management and the practical realities of implementing these technologies within existing enterprise environments [2]. By examining current research and emerging practices, we develop a framework for understanding how organizations can harness AI to transform their incident management capabilities and build greater operational resilience in our increasingly complex technological world.

2. Automated Incident Detection and Creation

The transformation of incident management has accelerated dramatically with the introduction of artificial intelligence capabilities specifically designed for automated detection and incident creation. Traditional detection methods relied heavily on preset thresholds and manual user reports, which frequently led to delayed responses and overlooked incidents. AI technologies have revolutionized this landscape by enabling proactive monitoring, sophisticated pattern recognition, and automated incident creation workflows that minimize human involvement while dramatically improving detection accuracy [3].

2.1 AI-Driven Monitoring Systems

Today's AI-powered monitoring systems go well beyond traditional threshold-based approaches by simultaneously analyzing multiple data streams throughout enterprise environments. These advanced systems harness sophisticated algorithms to process system logs, application performance metrics, network traffic patterns, and user behavior signals all at once [3]. By building comprehensive monitoring frameworks, organizations achieve near-instantaneous visibility into system health and potential problems. AI monitoring excels at managing the enormous data volumes generated by modern environments, processing information at scales that would completely overwhelm human analysts. This capability enables holistic monitoring strategies that consider both technical measurements and business impact indicators when evaluating potential incidents [4].

2.2 Anomaly Detection Through Machine Learning

Machine learning models have completely transformed anomaly detection by moving past rigid thresholds toward flexible pattern recognition. These models learn baseline behaviors for systems, applications, and networks by studying historical performance data and understanding normal operating patterns [3]. After training, the models detect subtle variations from established patterns that might indicate developing incidents, often long before traditional threshold-based alerts would trigger. Supervised learning techniques use labeled incident data to train classification models that recognize specific known incident types, while unsupervised learning methods excel at discovering new anomalies without previous examples. Reinforcement learning models enhance detection capabilities further by continuously refining their accuracy based on feedback about detection results [4].

Detection Approach	Detection Mechanism	Time to Detection	False Positive Rate	Contextual Awareness
Traditional Threshold-Based	Static thresholds on individual metrics	Delayed until threshold breach	Higher due to static thresholds	Limited to predefined metrics
AI-Driven Anomaly Detection	Machine learning pattern recognition	Earlier identification of emerging issues	Lower with properly trained models	Comprehensive across multiple metrics
Rules-Based Correlation	Predefined correlation rules	Moderate, dependent on rule complexity	Moderate, based on rule precision	Limited to anticipated scenarios
Unsupervised Learning	Dynamic baseline establishment	Early detection of novel anomalies	Variable based on sensitivity settings	High across varied data sources
NLP-Based User Report Analysis	Sentiment and keyword extraction	Depends on user reporting speed	Low for explicit reports	High for user experience context

Table 1: Comparison of Traditional vs. AI-Driven Incident Detection Methods [1, 3, 4]

2.3 Natural Language Processing for User-Reported Issues

While system monitoring delivers valuable technical insights, user-reported problems often contain essential contextual details about incident impact and user experience. Natural Language Processing (NLP) technologies allow systems to automatically extract and analyze information from unstructured sources including support tickets, chat conversations, social media posts, and customer feedback [4]. These capabilities enable incident management systems to spot potential incidents from user reports, pull out relevant technical information, gauge severity through sentiment analysis, and connect user feedback with system monitoring data. By bridging technical monitoring with user experience, NLP technologies create a more complete picture of potential incidents and their business consequences.

2.4 Automated Incident Creation and Enrichment

The combination of AI-powered monitoring, anomaly detection, and natural language processing culminates in automated generation of detailed incident records containing comprehensive information about detected problems [3]. Modern incident creation systems use AI to automatically generate tickets with pre-filled fields covering affected systems, likely root causes, severity evaluations, and business impact assessments. These systems also enrich incident records by adding relevant historical information, known issues from similar systems, and documentation from past resolution efforts. By automating incident creation, organizations dramatically reduce the gap between detection and response initiation while ensuring responders have comprehensive contextual information available [4].

2.5 Enterprise Implementation Case Studies

Real-world implementations of AI-powered incident detection and creation systems have delivered impressive benefits across various enterprise environments. Financial services companies have deployed machine learning-based anomaly detection to spot potential security incidents and transaction processing problems with far greater accuracy than traditional approaches [3]. Cloud service providers have built comprehensive monitoring systems that correlate data across thousands of servers to identify developing performance issues before customers notice. Healthcare organizations have used NLP capabilities to pull critical information from clinical system reports and automatically create incidents for potential failures that could affect patient care. Telecommunications providers have implemented end-to-end automated detection systems that cut mean time to detection while also reducing false positive rates [4]. These implementations demonstrate the flexibility and effectiveness of AI-powered approaches across diverse operational environments.

3. Intelligent Severity Classification and Prioritization

Accurately classifying incident severity remains a cornerstone of effective incident management, directly influencing response priorities, resource allocation, and resolution timelines. Traditional severity classification approaches often relied on subjective assessments by IT personnel, resulting in inconsistent categorizations influenced by individual perspectives and varying experience levels. The introduction of artificial intelligence to severity classification has transformed this area by bringing data-driven approaches that improve consistency, accuracy, and responsiveness in incident prioritization [5].

3.1 Leveraging Historical Data for Predictive Severity Assessment

AI-enhanced severity classification systems tap into extensive historical incident data to build predictive models that accurately evaluate new incidents. These systems examine past incident records encompassing resolution times, business impact measurements, customer complaints, and technical specifics to uncover patterns linked to different severity levels [5]. Machine learning algorithms trained on this historical information can identify subtle connections between incident characteristics and their actual impacts, enabling more precise initial severity classifications. This approach reduces the subjectivity inherent in manual assessment while incorporating organizational wisdom gained through previous incident management experiences. By creating a data-driven foundation for severity classification, organizations ensure more consistent prioritization decisions while reducing reliance on individual expertise [6].

3.2 System Dependency Mapping and Impact Correlation

Today's enterprise environments feature complex interconnected systems where incidents affecting one component can ripple through related services. Intelligent severity classification systems incorporate detailed dependency mapping to comprehend these relationships and accurately gauge potential impact scope [5]. AI models examine system architecture diagrams, configuration management databases, and actual interaction patterns to create dynamic dependency maps reflecting current operational relationships. When incidents arise, these models trace potential impact paths through dependent systems, forecasting the extent of service disruption throughout the enterprise architecture. This capability produces more accurate severity classifications by considering not just the immediately affected system but the broader operational context surrounding the incident [6].

3.3 Business Impact and Sentiment Analysis

While technical measurements provide valuable system health information, genuine incident severity must ultimately reflect business impact. Advanced severity classification systems incorporate multiple business context dimensions into their assessment processes [6]. Revenue impact models connect affected systems with business transactions to estimate potential financial consequences. Customer experience analysis examines user feedback, support tickets, and social media sentiment to gauge perception impacts. Contractual obligation evaluation identifies potential service level agreement breaches linked to specific incidents. Reputation risk assessment evaluates public-facing services and potential brand damage. By combining these business dimensions with technical assessments, AI-powered classification systems create comprehensive severity evaluations that align incident prioritization with organizational goals [5].

3.4 Dynamic Severity Adjustment Mechanisms

Incident conditions often change after initial detection, potentially altering their impact and appropriate priority level. Intelligent severity classification systems implement continuous monitoring and dynamic adjustment features to address this reality [6]. These systems analyze ongoing telemetry data, additional user reports, and resolution progress to spot changing conditions that might require severity reclassification. Machine learning models evaluate new information against established patterns to determine when adjustments are needed, automatically updating incident records and alerting relevant stakeholders. This dynamic approach keeps response prioritization aligned with actual impact throughout the incident lifecycle, optimizing resource allocation as situations evolve [5].

3.5 Multi-Dimensional Classification Frameworks

The most advanced severity classification systems employ multi-dimensional frameworks that simultaneously consider various factors when determining appropriate incident prioritization [5]. Technical impact dimensions assess system performance degradation, security implications, and data integrity concerns. Operational dimensions evaluate workforce disruption, process interruptions, and productivity impacts. Customer dimensions analyze user experience effects, visibility to external stakeholders, and satisfaction implications. Strategic dimensions consider alignment with organizational priorities, regulatory compliance risks, and reputational concerns. By combining assessments across these dimensions, AI-powered classification systems create nuanced severity determinations that accurately reflect organizational priorities while ensuring appropriate response allocation [6].

Classification Dimension	Key Assessment Factors	Data Sources	AI Techniques Applied
Operational Scope	Affected systems, dependency propagation	Configuration database, dependency maps	Graph analysis
User Experience	Customer sentiment, visibility, complaint volume	Support tickets, social media	Sentiment analysis, NLP
Temporal Factors	Time of day, business cycle relevance	Business calendars, historical patterns	Temporal pattern recognition
Resolution Complexity	Required expertise, known solutions	Knowledge base, historical incidents	Similarity matching

Table 2: AI-Driven Severity Classification Framework Dimensions [5, 6]

4. Dynamic Impact Assessment and Documentation

Thoroughly assessing incident impact and comprehensively documenting response activities form essential pillars of effective incident management. Traditional impact assessment approaches often depended on static analyses performed at specific moments during incidents, while documentation typically relied on manual note-taking and report writing. Integrating artificial intelligence into these processes has unlocked dynamic, continuous impact assessment and automated documentation features that dramatically improve organizational understanding of incidents and their resolution [7].

4.1 AI-Powered System Dependency Mapping

Accurate impact assessment demands deep understanding of system dependencies and potential propagation paths for incidents across the infrastructure. AI algorithms have revolutionized dependency mapping by examining system interactions, communication patterns, and configuration data to build comprehensive relationship models [7]. These models uncover both explicit dependencies documented in configuration management databases and implicit relationships revealed through operational data analysis. Machine learning techniques can discover subtle interdependencies that manual mapping might miss, delivering a more accurate picture of how incidents might spread. By maintaining continuously updated dependency maps, organizations build a foundation for precise impact assessment that reflects the current state of increasingly fluid infrastructure environments [8].

4.2 Real-Time Impact Monitoring and Refinement

Incident impact often evolves as conditions shift, mitigation efforts advance, and cascading effects develop across dependent systems. AI-powered impact assessment systems implement continuous monitoring features that track multiple indicators throughout the incident lifecycle [7]. Performance metric analysis measures ongoing system degradation or improvement across affected components. User experience monitoring evaluates changing impacts on customer interactions and satisfaction. Business transaction tracking quantifies evolving financial implications as incidents progress. By combining these continuous data streams, AI systems refine impact assessments in real-time, giving incident responders accurate, current information about situation scope and severity. This dynamic approach marks a major advancement beyond traditional point-in-time impact evaluations [8].

4.3 Self-Learning Impact Assessment Models

Impact assessment accuracy improves dramatically when systems learn from historical incidents and their actual outcomes. Self-learning models examine past incidents, comparing initial impact predictions against eventual documented effects to uncover

assessment patterns and improvement opportunities [8]. These models determine which indicators provided reliable impact signals during previous incidents, adjusting their weighting algorithms to emphasize valuable predictors. Post-incident analysis feedback mechanisms let response teams flag inaccurate assessments, supplying additional training data for ongoing improvement. Through this iterative learning process, impact assessment models achieve increasing accuracy over time, adapting to changing infrastructure environments and evolving incident patterns. This learning capability helps organizations benefit from accumulated incident management experience while reducing dependence on individual expertise [7].

4.4 Automated Documentation Through AI-Driven Transcription

Thorough incident documentation delivers critical information for current resolution efforts while building knowledge repositories for future incident management. AI-powered transcription systems convert audio content from incident response meetings, conference calls, and verbal updates into searchable, analyzable text records [7]. Advanced speech recognition algorithms optimized for technical terminology achieve impressive accuracy even with specialized vocabulary and multiple speakers. Speaker identification features attribute statements to specific participants, preserving accountability and context. Real-time transcription during active incident response meetings enables immediate reference to earlier discussions while ensuring comprehensive documentation without manual note-taking burden. By capturing complete verbal communication records, these systems preserve valuable context and details that manual documentation might miss [8].

4.5 Context-Aware Summarization and Knowledge Extraction

The information volume generated during incident management can overwhelm manual analysis capabilities, making AI-powered summarization crucial for effective knowledge use. Context-aware summarization systems process transcripts, chat logs, system updates, and other documentation to extract essential information while filtering unnecessary content [8]. These systems pinpoint key decisions, action items, and milestone events from extensive documentation. Technical detail extraction captures troubleshooting steps, configuration changes, and resolution approaches. Root cause identification algorithms highlight discussions about underlying issues and contributing factors. By creating structured summaries from voluminous documentation, AI systems make incident knowledge more accessible for both current resolution activities and future reference. This capability enhances organizational learning while ensuring critical insights aren't buried in information overload [7].

5. Implementation Challenges and Considerations

While artificial intelligence promises transformative benefits for incident management, successful implementation requires navigating several significant hurdles. Organizations must address data quality issues, bias concerns, integration difficulties, and ongoing governance requirements to fully realize AI's potential in incident management [9]. Understanding these challenges and developing effective strategies to overcome them proves essential for successful AI deployment within incident management frameworks.

A. 5.1 Data Quality Requirements for AI Model Effectiveness

The success of AI-powered incident management systems fundamentally depends on the quality of training and operational data. High-quality training datasets must contain sufficient volumes of historical incidents to enable reliable pattern recognition across various incident types and severities [10]. Data completeness becomes crucial, as missing information about incident characteristics, resolution approaches, or impact assessments creates gaps in model understanding. Consistency in historical incident documentation practices significantly affects AI model development, with irregular classification or terminology hampering pattern identification. Organizations must implement thorough data preparation processes including cleaning, standardization, and enrichment to address quality issues in existing incident repositories before AI deployment [9]. Ongoing data governance frameworks become essential to maintain quality standards as new incident data flows into AI systems, ensuring continuous improvement rather than performance degradation over time.

5.2 Addressing Bias and Fairness in Incident Classification

AI models can inherit biases present in their training data, potentially perpetuating or amplifying unfair incident classification patterns. Historical incident data might contain systemic biases in severity ratings, prioritization decisions, or resource allocation that inadvertently transfer into AI models [9]. Technical biases may emerge when certain system types or technologies appear disproportionately in training data, leading to less accurate assessments for underrepresented systems. Organizational biases might surface when past incident management showed preferential treatment toward particular business units or functions. Bias detection techniques during model development become crucial, with fairness metrics and counterfactual testing helping identify potential discrimination patterns [10]. Diverse training data spanning a wide range of incident types, systems, and organizational contexts helps reduce bias risks. Continuous monitoring for fairness in operational incident classification ensures AI systems maintain equitable assessment practices over time, with regular auditing to identify and address emerging bias patterns.

5.3 Integration Complexities with Existing Systems

Integrating AI-powered incident management capabilities with existing IT service management platforms, monitoring tools, and communication systems creates significant technical and procedural challenges [9]. Data integration needs span multiple systems including monitoring platforms, ticketing systems, knowledge bases, and communication tools that often use different data formats and access protocols. Workflow integration becomes crucial to incorporate AI capabilities into established incident management processes without disrupting response procedures. Real-time processing demands for incident detection and impact assessment may strain existing infrastructure capabilities, requiring architectural adjustments. API limitations in legacy systems can restrict data access or action implementation for AI components, potentially requiring custom integration solutions [10]. Organizations must develop thorough integration strategies addressing both technical connection points and procedural handoffs between AI and human-driven components of the incident management ecosystem.

Challenge Category	Specific Challenges	Mitigation Strategies	Organizational Considerations
Data Quality	Incomplete records, inconsistent classifications	Data enrichment, standardization protocols	Establish data governance frameworks
Bias Management	Historical prioritization biases, technical imbalances	Fairness metrics, diverse training data	Regular fairness audits
Technical Integration	Legacy system limitations, real-time processing needs	API development, middleware solutions	Platform modernization
Process Adaptation	Workflow disruption, responsibility transitions	Process mapping, phased automation	Change management programs
Governance	Accountability boundaries, explainability requirements	Decision authority frameworks, explainable AI	Policy development
Skills Development	AI expertise gaps, interpretation capabilities	Training programs, expert partnerships	Workforce development planning

Table 3: Implementation Challenges and Mitigation Strategies for AI-Driven Incident Management [9, 10]

5.4 Governance and Operational Oversight

Strong governance frameworks prove essential for maintaining appropriate control and accountability in AI-powered incident management systems [10]. Clear responsibility boundaries between AI systems and human operators help establish appropriate automation limits while preserving necessary oversight. Explainability requirements ensure AI recommendations and decisions remain understandable and verifiable by human operators, particularly for critical incident classification and prioritization choices. Documentation standards for AI-driven processes enable appropriate audit trails and compliance verification. Escalation protocols must be established for situations where AI systems produce uncertain assessments or encounter novel scenarios requiring human judgment [9]. Regular performance reviews measuring AI system effectiveness against established metrics help identify improvement opportunities and potential issues needing attention.

5.5 Future Developments and Emerging Capabilities

As AI technologies continue advancing, incident management systems move toward increasingly sophisticated capabilities that will further revolutionize organizational practices [10]. Predictive incident management represents a major leap forward, using historical patterns and real-time system signals to forecast potential incidents before they occur. This early warning capability enables proactive intervention that can prevent or minimize impact before users experience problems. Autonomous remediation functionality emerges for well-understood incident types, with AI systems implementing proven resolution procedures automatically without human intervention for routine issues. Collaborative AI approaches balance task distribution between artificial intelligence and human operators, leveraging each party's strengths [9]. Cross-organizational intelligence sharing enables broader learning from incident patterns across enterprises, accelerating improvement while maintaining appropriate

privacy and security boundaries. These emerging capabilities point toward a future where incident management becomes increasingly proactive and effective through deeper AI integration.

6. Conclusion

The integration of artificial intelligence technologies into incident management processes marks a fundamental shift in how organizations detect, classify, document, and resolve system disruptions. By applying AI capabilities across automated incident detection, intelligent severity classification, dynamic impact assessment, and comprehensive documentation, enterprises can achieve substantial improvements in operational resilience while dramatically reducing mean time to resolution. The capabilities discussed throughout this paper—from anomaly detection models that spot emerging issues before users feel the impact to self-learning impact assessment systems that grow more accurate with experience—demonstrate the breadth of AI applications possible within incident management frameworks. While implementation challenges around data quality, bias mitigation, and system integration demand careful attention, organizations that successfully navigate these obstacles position themselves to capture significant benefits in incident response efficiency and effectiveness. As AI technologies continue maturing, the future of incident management shifts increasingly from reactive to proactive approaches, with predictive capabilities and autonomous remediation becoming standard practice. This technological evolution has the potential to fundamentally reshape how organizations think about system reliability and service continuity, ultimately driving both operational excellence and superior customer experiences through smarter, more responsive incident management processes.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1] Advised Skills, "Challenges in AI Implementation and Solutions." Advised Skills Blog. 2025. <https://www.advisedskills.com/blog/artificial-intelligence-ai/challenges-in-ai-implementation-and-solutions>
- [2] Altay Ataman, "Data Quality in AI: Challenges, Importance & Best Practices." AIMultiple. March 28, 2025. <https://research.aimultiple.com/data-quality-ai/>
- [3] Arturo Peralta, et al. "Intelligent Incident Management Leveraging Artificial Intelligence." Mathematics, 13(7), 1055. March 24, 2025. <https://www.mdpi.com/2227-7390/13/7/1055>
- [4] LOUIS AU YEUNG. "Guidance for the Development of AI Risk and Impact Assessments." UC Berkeley Center for Long-Term Cybersecurity. July 2021.
- [5] oham Mehta, et al. "Dynamic Documentation for AI Systems." arXiv.org. March 20, 2023. <https://arxiv.org/abs/2303.10854>
- [6] Pouya Hamadani, et al. "A Holistic View of AI-driven Network Incident Management." ACM SIGCOMM HotNets, 45-52. November 28, 2023. https://conferences.sigcomm.org/hotnets/2023/papers/hotnets23_hamadani.pdf
- [7] Hanzhang Wang, et al. "Anomaly Detection for Incident Response at Scale." arXiv.org. April 24, 2024. <https://arxiv.org/abs/2404.16887>
- [8] Weiwei Pan; Huixin He. "An Intelligent Fault Severity Diagnosis Method based on Hybrid Ordinal Classification." IEEE ITNEC Conference. June 12-14, 2020. <https://ieeexplore.ieee.org/document/9084970>
- [9] Yalan Jiang, et al. "Analyzing Crash Severity: Human Injury Severity Prediction Method Based on Transformer Model." Vehicles (MDPI), 7(1), 5. January 15, 2025. <https://www.mdpi.com/2624-8921/7/1/5>
- [10] Yash Sonawane, "AIOps: Automating Incident Management with AI & Machine Learning." DEV Community. April 3, 2025. https://dev.to/yash_sonawane25/aioops-automating-incident-management-with-ai-machine-learning-4ebg