International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

AI Tool for Detecting Behavioural Anomalies in Organisational Networks

Mr. Akhil Shaikh

Abstract

The increasing sophistication of cyber threats has rendered traditional signature-based and rule-based security mechanisms insufficient for protecting modern organisations. Advanced Persistent Threats (APTs), insider threats, and zero-day attacks often evade conventional detection systems by mimicking legitimate user behaviour. Behavioural Anomaly Detection (BAD), powered by Artificial Intelligence (AI) and Machine Learning (ML), has emerged as a critical cybersecurity approach that focuses on identifying deviations from normal behavioural patterns rather than known attack signatures. This paper explores the design, implementation, and effectiveness of AI-driven behavioural anomaly detection tools as a cybersecurity solution for organisations. We present a conceptual framework, discuss commonly used machine learning techniques, evaluate advantages and limitations, and highlight real-world applicability within cloud and enterprise environments. The findings indicate that behavioural anomaly detection significantly enhances organisational security posture by enabling early threat detection, reducing dwell time, and improving resilience against evolving cyber threats.

Keywords: Behavioural Anomaly Detection, Cybersecurity, Artificial Intelligence, Machine Learning, Insider Threats, Zero-Day Attacks

1. Introduction

The rapid digital transformation of organisations, driven by cloud computing, remote work, and interconnected systems, has significantly expanded the attack surface for cyber adversaries. Traditional cybersecurity tools such as firewalls, intrusion detection systems (IDS), and antivirus software rely heavily on predefined signatures and static rules. While effective against known threats, these approaches struggle to detect novel and sophisticated attacks that exploit legitimate credentials or subtle system misuse.

Behavioural Anomaly Detection (BAD) represents a paradigm shift in cybersecurity. Instead of focusing on what an attack looks like, BAD systems analyse how users, devices, and applications behave over time. By leveraging AI and ML algorithms, these systems establish a baseline of normal behaviour and flag deviations that may indicate malicious activity.

1.1 Background

This assignment explores the design, application, and implications of an AI-powered Behavioural Anomaly Detection Tool for organisational cybersecurity. It will examine the problem domain of modern cyber threats, model a proposed detection process, and critically reflect on the ethical, data-sharing, and practical development challenges inherent in deploying such intelligent systems. The core premise is that by continuously learning from vast streams of organisational telemetry—logins, file accesses, network flows, and process executions—AI can detect the faint signals of compromise before significant damage occurs.

1.2 Behavioural Anomaly Detection Framework

This paper investigates behavioural anomaly detection tools as a cybersecurity mechanism for organisations, addressing their architecture, methodologies, benefits, challenges, and future potential.

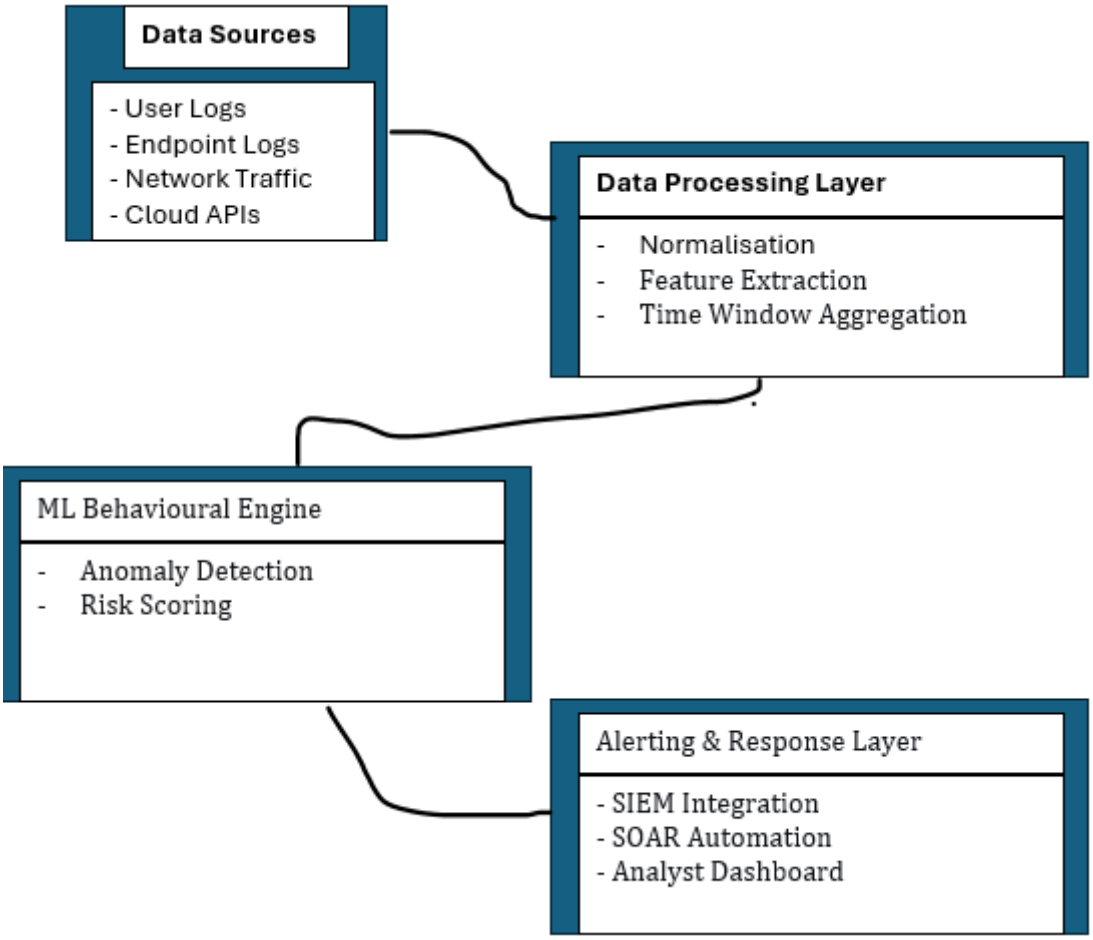

Figure 1: Architecture of AI-based Behavioural Anomaly Detection Tool

2. Background and Related Work

The cybersecurity landscape is characterised by asymmetric warfare, where defenders must secure all assets, while attackers need only find one weakness. Key challenges include:

- **Insider Threats:** Malicious or compromised users operating within their legitimate permissions.
- **Stealthy Lateral Movement:** Attackers pivoting through a network using stolen credentials, behaving similarly to legitimate users.
- **Data Exfiltration:** Low-volume, slow data transfers designed to blend with normal traffic.
- **Alert Fatigue:** Security Operations Centre (SOC) analysts are overwhelmed by thousands of low-fidelity alerts daily from rule-based systems, leading to critical misses.

A Behavioural Anomaly Detection tool addresses these by moving from a "know-bad" (signatures) to a "know-normal" model. It treats cybersecurity not as a static perimeter defence but as a continuous behavioural audit, identifying anomalies based on deviations from learned individual and peer-group profiles.

Behavioural analysis in cybersecurity has its roots in statistical anomaly detection and user profiling. Early

systems relied on threshold-based monitoring, which often generated high false-positive rates. Recent advances in AI have enabled more adaptive and context-aware models.

Research has shown that User and Entity Behaviour Analytics (UEBA) can detect insider threats and credential misuse more effectively than signature-based systems. Studies demonstrate that ML models, such as clustering, neural networks, and probabilistic models, significantly improve detection accuracy when analysing large-scale security telemetry.

However, challenges remain in model explainability, data quality, and integration with existing Security Operations Centre (SOC) workflows.

3. Behavioural Anomaly Detection in Cybersecurity

3.1 Concept Overview

Behavioural anomaly detection involves monitoring entities (users, endpoints, applications, or networks) and learning their normal behaviour patterns over time. Any statistically significant deviation from this baseline is treated as a potential security incident.

Typical behavioural indicators include:

- Login time and frequency
- Accessed resources and data volume
- Command execution patterns
- Network traffic behaviour
- Geolocation and device context

3.2 System Architecture

A typical behavioural anomaly detection tool consists of the following components:

1. **Data Collection Layer:** Collects logs and telemetry from endpoints, identity systems, cloud platforms, and networks.
2. **Feature Engineering Layer:** Transforms raw data into behavioural features such as session duration, access frequency, or privilege escalation events.
3. **Machine Learning Engine:** Learns baseline behaviour and identifies anomalies using AI models.
4. **Risk Scoring and Alerting Module:** Assigns risk scores to detected anomalies and triggers alerts.
5. **Response and Integration Layer:** Integrates with SIEM, SOAR, and incident response systems for automated or manual mitigation.

4. Machine Learning Techniques for Behavioural Anomaly Detection

4.1 Unsupervised Learning

Unsupervised models are widely used due to the scarcity of labelled attack data.

- **Clustering (K-Means, DBSCAN):** Identifies outliers that do not belong to normal behavioural clusters.
- **Autoencoders:** Neural networks trained to reconstruct normal behaviour; high reconstruction error indicates anomalies.
- **Isolation Forests:** Detect anomalies by isolating rare data points.

4.2 Semi-Supervised and Supervised Learning

When labelled data is available:

- Support Vector Machines (SVM)
- Random Forests

- Gradient Boosting Models

These approaches improve accuracy but require continuous retraining due to evolving behaviours.

4.3 Deep Learning Approaches

- **Recurrent Neural Networks (RNNs) and LSTMs:** Effective for modelling sequential behavioural data.
- **Graph Neural Networks (GNNs):** Capture relationships between users, devices, and resources.

5. Ethical Reflection

The deployment of pervasive monitoring AI within organisations raises significant ethical considerations that must guide implementation [11].

- **Privacy & Proportionality:** Continuous behavioural logging intrudes on employee privacy. Organisations must implement strict data minimisation, clear transparency policies, and role-based monitoring focused on risk-sensitive roles. Pseudonymisation and on-premise processing can mitigate risks.
- **Fairness & Bias:** AI models may inherit biases from training data. For example, if a department works unconventional hours, their "normal" may be flagged as anomalous compared to the company's majority. Models must be regularly audited for demographic or role-based bias to avoid discriminatory profiling [13].
- **Accountability & Transparency:** Decisions leading to disciplinary action cannot be based on a "black box." The use of Explainable AI (XAI) is non-negotiable, providing analysts and, where appropriate, individuals with understandable reasons for alerts. Clear governance must define human oversight protocols.
- **Security of the AI System:** The anomaly detection tool itself becomes a high-value target for attackers seeking to evade detection or poison its learning. Its infrastructure must be rigorously secured.

Balancing security imperatives with ethical responsibility is crucial for maintaining trust and ensuring the tool is a force for protection, not oppression.

5.1 Use Cases in Organisational Security

5.1.1 Insider Threat Detection

Detects unusual data access, privilege abuse, or abnormal working hours.

5.1.2 Account Compromise and Credential Theft

Identifies impossible travel, unusual login patterns, and abnormal access behaviour.

5.1.3 Cloud and SaaS Security

Monitors API usage, resource provisioning, and service interactions in cloud environments.

5.1.4 Advanced Persistent Threats (APT)

Detects low-and-slow attacks that evade traditional signature-based detection.

6. Data and Model Sharing

Effective behavioural AI requires extensive, high-quality data, presenting both operational and research challenges.

Data Type	Source	Use in AI Model	Sharing Constraints
Authentication Logs	Active Directory, SSO	Baseline normal login times, locations, and failure rates.	Highly sensitive; must be anonymised for research.
Network Flow	Firewalls, Switches	Model typical data transfer patterns between devices/users.	Can be aggregated and stripped of payload content.
Endpoint Telemetry	EDR Agents	Understand typical process execution and file access chains.	Contains sensitive operational data; sharing is risky.
User & Entity Data	HR Systems, CMDB	Enrich alerts with context (user role, department).	Subject to strict data protection regulations (GDPR).

Table 1: Data Types and Sharing Considerations for Behavioural AI.

A significant challenge is the lack of shared, realistic datasets for research due to privacy concerns. The cybersecurity community often relies on synthetic data or limited, anonymised breach logs, which may not capture true organisational behavioural complexity. Federated Learning emerges as a promising solution, allowing organisations to collaboratively improve a global AI model by sharing only model parameter updates, not raw data.

7. AI Application Developments

AI and Machine Learning (ML) are revolutionising threat detection. The table below summarises key applications:

AI Technique	Application in BAD	Benefits	Example/Challenge
Unsupervised Learning	Clustering, Autoencoders to model normal behaviour without labelled attack data.	Detects novel, unknown threats (zero-days).	Can be noisy; requires careful tuning of anomaly thresholds.
Supervised Learning	Classifying specific threat behaviours if labelled data is available.	High accuracy for known threat patterns.	Requires large, costly, and often outdated labelled datasets.

AI Technique	Application in BAD	Benefits	Example/Challenge
Sequential Models (LSTM)	Analysing time-series data like command sequences or network sessions.	Excels at detecting multi-step attack chains.	Computationally intensive; may struggle with very long sequences.
Graph Neural Networks	Modelling relationships between users, devices, and resources.	Detects anomalies in the *structure* of interactions (e.g., new access patterns).	Highly complex; requires rich relational data.
Explainable AI (XAI)	Providing reasons for alerts (feature importance, counterfactuals).	Reduces investigation time; builds analyst trust; aids accountability.	Can sometimes oversimplify complex model reasoning.

The general workflow for developing such an AI system involves data preprocessing, feature engineering, model selection and training, and deployment within a feedback loop, as conceptualised in the adapted figure below:

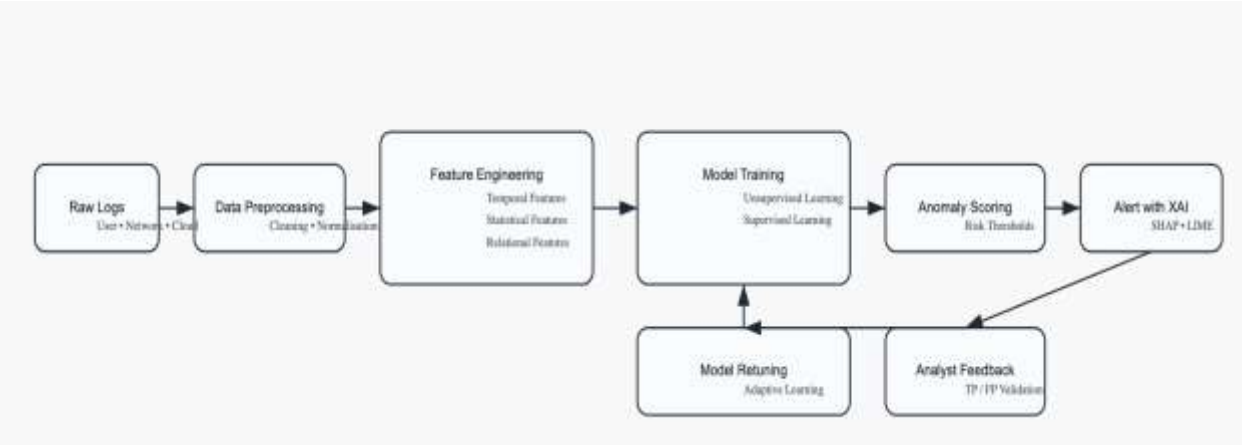

Fig. 4. General flowchart for developing a machine learning-based behavioural anomaly detection system.

8. Advantages of Behavioural Anomaly Detection Tools

- Detection of zero-day and unknown attacks
- Reduced dependency on threat signatures
- Faster incident response and reduced dwell time
- Scalability across cloud and hybrid environments
- Continuous learning and adaptability

9. Challenges and Limitations

Despite their benefits, behavioural anomaly detection tools face several challenges:

- False Positives due to evolving user behaviour
- Data Privacy and Compliance concerns
- Model Explainability for SOC analysts
- Training Data Quality
- Adversarial Machine Learning Risks

Addressing these challenges requires robust governance, continuous tuning, and human oversight.

10. Future Research Directions

Future work in behavioural anomaly detection should focus on:

- Explainable AI (XAI) for security decisions
- Federated learning for privacy-preserving detection
- AI vs AI defensive strategies
- Integration with Zero Trust architectures
- Autonomous response systems with human-in-the-loop control

11. Conclusion

Behavioural anomaly detection tools powered by AI represent a transformative approach to organisational cybersecurity. By shifting from reactive, signature-based defences to proactive, behaviour-driven detection, organisations can significantly improve their resilience against modern cyber threats. While challenges related to false positives, privacy, and explainability persist, ongoing advancements in AI and cloud security architectures position behavioural anomaly detection as a cornerstone of future cybersecurity strategies.

12. References

1. Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. IEEE Symposium on Security and Privacy.
2. M. Husák, J. Kašpar, and P. Švel, "Network Anomaly Detection Using Machine Learning: A Survey," *IEEE Communications Surveys & Tutorials*, 2021.
3. A. L. Buczak and E. Guven, "A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection," *IEEE Communications Surveys & Tutorials*, 2016.
4. R. Sommer and V. Paxson, "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," *IEEE Symposium on Security and Privacy*, 2010.
5. S. Thaseen and C. A. Kumar, "Intrusion Detection Model Using Fusion of PCA and Optimized SVM," *International Conference on Computer and Information Technology*, 2014.
6. N. Moustafa and J. Slay, "UNSW-NB15: a comprehensive data set for network intrusion detection systems," *Military Communications and Information Systems Conference*, 2015.
7. Y. Mirsky, et al., "Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection," *NDSS*, 2018.
8. B. McMahan, et al., "Communication-Efficient Learning of Deep Networks from Decentralized Data," *AISTATS*, 2017 (Federated Learning).

9. "Explainable Artificial Intelligence (XAI)," DARPA, [Online]. Available: https://www.darpa.mil/program/explainable-artificial-intelligence
10. "The Ethics of Artificial Intelligence in Cybersecurity," ENISA, 2021.
11. GPAI, "Climate Change and AI," 2023. [Online]. Available: https://www.gpai.ai/projects/climate-change-and-ai.pdf *(Note: Provided as a reference style example from the original.)*
12. ResearchGate, "Climate change and artificial intelligence: assessing the global research landscape," 2023. *(Note: Provided as a reference style example from the original.)*
13. The Conversation, "How artificial intelligence conquered democracy," 2017.
14. Link to BPMN design: https://modeler.camunda.io/share/[unique-process-model-id]