MITRE-ATT&CK Integration for Behavioural HARM-based Attack Simulation

Fomotar Nyuykighan Buhnyuy, Femi Fasunlade, Oluwatobi Fajana and Mmedo Essien
University of Portsmouth, UK

fomotar.buhnyuy@myport.ac.uk
femi.fasunlade@port.ac.uk
oluwatobi.fajana@port.ac.uk
mmedo.essien@johncrane.com

Abstract: Automated attack simulation has emerged as an important approach for assessing the security of complex networked systems in a scalable and repeatable manner. Frameworks based on the Hierarchical Attack Representation Model (HARM), such as HARMer, support automated vulnerability discovery, attack-path generation, and attack execution. However, existing HARM-based approaches remain largely vulnerability-centric, relying on Common Vulnerabilities and Exposures (CVE) identifiers that provide limited insight into adversary behaviour. Consequently, generated attack paths are often difficult to interpret from a threat-informed defence perspective. This research investigates the integration of the MITRE ATT&CK framework into a HARM-based automated attack simulation pipeline to improve the behavioural interpretability of simulation outputs. An experimental methodology was adopted using an enhanced HARMer framework deployed within controlled and isolated environments. Vulnerability information obtained through automated scanning was used to construct HARM models representing network reachability and host-level vulnerabilities. Identified CVEs were systematically mapped to ATT&CK tactics and techniques using vulnerability descriptions, exploit documentation, and ATT&CK framework definitions. The resulting behavioural information was incorporated into attack planning and visualisation processes. To evaluate the effectiveness of the proposed enhancement, four evaluation indicators were defined: Mapping Coverage Rate (MCR), Behavioural Annotation Rate (BAR), Technique Coverage (TC), and Attack Path Preservation (APP). These indicators were used to assess the feasibility of behavioural enrichment while ensuring that the automation and scalability characteristics of the original HARMer framework were maintained. The results demonstrate that CVE-based vulnerabilities can be systematically enriched with ATT&CK semantics, enabling attack paths to be interpreted as sequences of adversary behaviours rather than isolated technical exploits. Behaviour-enhanced attack planning preserved the original path-generation logic while providing additional contextual information regarding attacker tactics and techniques. Furthermore, ATT&CK-based visualisation provided an intuitive representation of behavioural coverage across simulated attack scenarios, supporting clearer interpretation of attack activity. The study contributes a practical approach for integrating behavioural threat intelligence into HARM-based automated attack simulation. By combining vulnerability-driven attack modelling with ATT&CK-based behavioural semantics, the proposed framework advances automated red-teaming beyond purely technical exploit analysis towards more interpretable, threat-informed, and operationally relevant security assessment.

Keywords: Automated attack simulation, HARM, HARMer, MITRE ATT&CK, Threat-informed defence, Automated red-teaming, Vulnerability modelling, Adversary behaviour

1. Introduction

Modern organisations increasingly rely on complex and interconnected digital infrastructures to support critical operations. As these systems grow in scale and complexity, they present an expanded attack surface that can be exploited by adversaries employing multi-stage attack strategies. Traditional security assessment techniques, such as penetration testing and red team exercises, remain effective for identifying vulnerabilities, but they are often manual, time-consuming, and difficult to repeat consistently across different environments. These limitations have driven growing interest in automated attack simulation as a scalable and repeatable approach to security assessment.

One widely adopted approach to automated attack simulation is the Hierarchical Attack Representation Model (HARM), which represents network reachability and host-level vulnerabilities in a structured, hierarchical manner (Hong and Kim, 2012). Frameworks such as HARMer extend HARM by automating vulnerability discovery, attack modelling, and planning, enabling large numbers of attack scenarios to be generated with limited human involvement (Enoch et al., 2020). Despite these advances, HARM-based approaches remain largely vulnerability-centric. Vulnerabilities are typically represented using Common Vulnerabilities and Exposures (CVEs), which provide detailed technical descriptions but limited insight into how adversaries exploit these weaknesses as part of broader attack campaigns.

This lack of behavioural context reduces the interpretability and practical value of automated simulation outputs. Attack paths may be technically feasible but difficult to interpret in terms of adversary intent, tactics,

or techniques. The MITRE ATT&CK framework addresses this limitation by providing a behaviour-focused model of cyber-attacks, categorising adversary actions into tactics and techniques derived from real-world observations (MITRE ATT&CK®, n.d.). While ATT&CK has been widely adopted for threat modelling, detection, and reporting, it is most commonly applied after attacks or simulations have occurred. Its integration into automated attack modelling and planning processes remains limited.

This research argues that automated attack simulation can be made more interpretable and threat-informed by integrating MITRE ATT&CK directly into HARM-based automation. To investigate this, experiments were conducted using an enhanced HARMer pipeline deployed in controlled environments. Container-based deployments and scalable virtual infrastructure were used to ensure repeatability and ethical containment. The focus of the experiments was not on exploit success rates or performance metrics, but on evaluating whether behavioural context could be systematically embedded into automated simulation outputs.

1.1 Aim and Objectives

The aim of this research is to enhance HARM-based automated attack simulation by integrating the MITRE ATT&CK framework in order to improve the behavioural interpretability of simulation outputs.

The objectives of this research are to:

- Systematically map CVE-based vulnerabilities to MITRE ATT&CK tactics and techniques, enabling behavioural interpretation of identified vulnerabilities.
- Embed MITRE ATT&CK identifiers into HARM-based attack planning, allowing attack paths to reflect adversary techniques rather than isolated exploits.
- Enable ATT&CK-based visualisation of simulated attacks, providing an interpretable view of behavioural coverage across experiments.

1.2 Research Questions

This study addresses the following research questions:

1. How can CVE-based vulnerabilities be systematically mapped to MITRE ATT&CK tactics and techniques?
2. How can MITRE ATT&CK be embedded into HARM-based attack planning to improve interpretability?
3. How can ATT&CK-based visualisation represent behavioural coverage in automated attack simulation?

The remainder of this paper is organised as follows. Section 2 reviews related work, Section 3 describes the methodology, Section 4 presents and discusses the results, and Section 5 concludes the paper.

2. Literature Review

Automated attack simulation has gained increasing attention as organisations seek scalable and repeatable alternatives to manual security assessment techniques. Early approaches focused on attack graphs, which model possible sequences of exploits that an attacker could use to compromise a network. While attack graphs provide valuable insights into potential attack paths, they often suffer from scalability issues and limited ability to represent complex, multi-layered systems (Hong and Kim, 2012). These limitations motivated the development of more structured representations that could better support automated analysis.

The Hierarchical Attack Representation Model (HARM) was proposed to address these challenges by combining network reachability information with host-level vulnerability data in a hierarchical structure (Hong and Kim, 2012). In HARM, an upper-layer attack graph captures network connectivity and access relationships, while lower-layer representations model vulnerabilities on individual hosts. This separation improves scalability and allows security analysts to reason about both network-level and host-level security properties. HARM has since been applied in a range of security assessment scenarios, including vulnerability prioritisation and risk analysis.

Building on the HARM concept, Enoch et al. (2020) introduced HARMer, an automated framework that integrates vulnerability scanning, attack modelling, planning, and execution. HARMer automates the generation of attack paths by collecting vulnerability information, constructing a HARM model, and executing planned attacks in a controlled environment. This approach significantly reduces the manual effort required for security assessment and enables repeatable experimentation. However, HARMer and similar frameworks remain primarily vulnerability-driven, relying heavily on CVE identifiers and exploit availability to guide attack planning.

While CVEs provide detailed technical descriptions of vulnerabilities, they offer limited insight into adversary behaviour. CVE-based models focus on *what* is vulnerable rather than *how* attackers operate in practice. As a result, automated attack paths generated using CVE information may be difficult to interpret from a threat intelligence or defensive perspective. Analysts may struggle to understand which attacker behaviours are being simulated, how those behaviours align with real-world threats, or which defensive controls are most relevant.

The MITRE ATT&CK framework was developed to address this gap by providing a behaviour-centric knowledge base of adversary tactics and techniques derived from observed cyber-attacks (MITRE ATT&CK®, n.d.). ATT&CK organises attacker behaviour into high-level tactics, such as Initial Access or Lateral Movement, and lower-level techniques that describe specific methods used to achieve these goals. The framework has been widely adopted for threat modelling, detection engineering, and security operations, offering a common language for describing adversary behaviour.

Several studies have explored the use of MITRE ATT&CK for attack analysis and defensive evaluation. Al-Sada et al. (2024) provide a comprehensive survey of ATT&CK applications, highlighting its role in detection gap analysis, threat intelligence mapping, and security control evaluation. Other work has focused on mapping alerts, logs, or incident reports to ATT&CK techniques in order to improve situational awareness. These approaches demonstrate the value of behavioural context but typically apply ATT&CK after attacks or simulations have already occurred.

More recent research has begun to investigate the integration of ATT&CK into proactive security assessment. Jiang et al. (2025) discuss emerging approaches for incorporating ATT&CK into security testing and evaluation workflows, noting that deeper integration could support more threat-informed defence strategies. However, existing studies largely treat ATT&CK as an annotation or reporting layer rather than as a core component of attack modelling and planning.

Despite advances in both automated attack simulation and behavioural threat frameworks, a clear gap remains. HARM-based frameworks such as HARMer provide automation and scalability but lack behavioural interpretability, while ATT&CK provides rich behavioural context but is rarely embedded directly into automated attack planning. There is limited empirical work demonstrating how these two approaches can be systematically combined to produce simulation outputs that are both automated and threat-informed.

This research addresses this gap by integrating MITRE ATT&CK into a HARM-based automated attack simulation framework. By mapping vulnerabilities to ATT&CK techniques and embedding behavioural identifiers into attack planning and visualisation, the study aims to enhance the interpretability and practical relevance of automated attack simulation results.

3. Methodology

This study adopted an experimental methodology to evaluate the feasibility and effectiveness of integrating the MITRE ATT&CK framework into HARM-based automated attack simulation. The experiments were conducted in controlled environments to ensure repeatability, safety, and ethical compliance. The methodological design focused on enhancing the existing HARMer framework by embedding behavioural context into vulnerability modelling, attack planning, and result visualisation.

3.1 Experimental Environment

The experimental evaluation was conducted using an enhanced HARMer framework deployed within a controlled and isolated test environment. The design of the environment was informed by the original HARMer implementation proposed by Enoch et al. (2020), which utilised both enterprise-scale and cloud-based network infrastructures for automated attack simulation. To ensure repeatability and ethical containment, all experiments were performed on intentionally vulnerable systems within a virtualised environment.

The experimental network consisted of seven interconnected hosts representing a simplified enterprise architecture. Network reachability between hosts was controlled through predefined access rules, allowing the construction of realistic multi-stage attack paths. One host was designated as the protected target system, requiring attackers to traverse intermediate hosts before achieving the attack objective. This topology was selected because it reflects the lateral movement scenarios commonly observed in enterprise attacks.

Automated vulnerability discovery was performed against the target environment to identify known vulnerabilities. A total of 27 vulnerabilities associated with Common Vulnerabilities and Exposures (CVE) identifiers were identified across the experimental hosts. Only vulnerabilities possessing valid CVE identifiers

were included in the ATT&CK integration process, as CVE references provided a consistent basis for behavioural mapping and validation.

The environment was designed to support HARM construction by providing both network reachability information and host-level vulnerability information. This allowed attack paths generated by the framework to be analysed from both a technical and behavioural perspective. The use of a controlled environment ensured that all attack simulations remained repeatable, measurable, and ethically compliant.

3.2 Vulnerability Discovery and HARM Construction

Automated vulnerability scanning was performed against the target systems to identify known vulnerabilities. The scan results were processed to extract CVE identifiers, which formed the basis for host-level vulnerability modelling. Network topology and access relationships were combined with the vulnerability data to construct a Hierarchical Attack Representation Model.

In the constructed HARM, the upper layer represented network reachability and potential access paths between hosts, while the lower layer modelled vulnerabilities present on individual systems. This hierarchical separation enabled scalable attack path generation while preserving detailed vulnerability information required for attack planning. The HARM model served as the foundational structure upon which behavioural enhancements were applied.

3.3 Mapping CVEs to MITRE ATT&CK

To introduce behavioural context into the automated attack simulation process, vulnerabilities identified during scanning were systematically mapped to tactics and techniques within the MITRE ATT&CK framework. The mapping process was designed to enrich technical vulnerability information with behavioural semantics while preserving the underlying HARM structure and attack planning workflow.

The mapping procedure followed a rule-based methodology. First, CVE records extracted from vulnerability scanning results were analysed using information obtained from National Vulnerability Database (NVD) descriptions, exploit documentation, vendor advisories, and publicly available threat intelligence sources. Each vulnerability was then examined to determine the attacker behaviour enabled by successful exploitation. Where a clear behavioural relationship existed, the vulnerability was mapped to the most representative ATT&CK tactic and technique.

For example, vulnerabilities enabling unauthorised remote access were associated with ATT&CK techniques under the Initial Access tactic, while privilege escalation vulnerabilities were mapped to techniques within the Privilege Escalation tactic. Vulnerabilities supporting remote command execution were linked to Execution techniques, and vulnerabilities facilitating movement between hosts were associated with Lateral Movement techniques.

To maintain consistency and reduce subjective interpretation, mappings were validated against the ATT&CK framework definitions and publicly documented ATT&CK use cases. Where a vulnerability could plausibly correspond to multiple ATT&CK techniques, the technique most directly associated with the primary exploitation outcome was selected. Vulnerabilities for which behavioural intent could not be confidently determined were excluded from the mapping process to avoid speculative attribution.

The resulting mappings were stored alongside vulnerability records within the enhanced HARMer framework. This allowed ATT&CK identifiers to be propagated automatically into attack planning, attack-path generation, and behavioural visualisation. Consequently, vulnerabilities were represented not only as technical weaknesses but also as enablers of specific adversary behaviours, providing a more threat-informed representation of attack activity.

3.4 Behaviour-Enhanced Attack Planning

The original HARMer framework generates attack paths by analysing network reachability, vulnerability information, and exploit availability within a Hierarchical Attack Representation Model (HARM). While this approach effectively identifies technically feasible attack paths, the resulting outputs remain largely vulnerability-centric and provide limited insight into adversary behaviour. To address this limitation, the framework was extended to incorporate MITRE ATT&CK identifiers directly into the attack planning process.

Following the CVE-to-ATT&CK mapping stage, each mapped vulnerability was enriched with one or more ATT&CK tactics and techniques. During attack-path generation, the planner retained the original HARM-based

path-selection process but additionally associated each attack step with its corresponding ATT&CK technique. As a result, generated attack paths contained both technical exploit information and behavioural context.

The enhanced planning process therefore transformed conventional exploit chains into behavioural attack narratives. Rather than representing an attack path solely as a sequence of exploited vulnerabilities, each step could be interpreted as a specific adversary action aligned with recognised ATT&CK tactics. For example, a path beginning with remote exploitation of a public-facing service could be interpreted as Initial Access, followed by Execution, Privilege Escalation, and Lateral Movement as the attacker progressed through the network.

Importantly, the introduction of ATT&CK semantics did not alter the underlying attack-path generation logic. Path feasibility continued to be determined by network reachability, vulnerability presence, and exploit availability, consistent with the original HARMer methodology. The enhancement therefore preserved the automation and scalability benefits of HARM-based attack simulation while increasing the interpretability of generated outputs.

The resulting attack plans provided a richer representation of attacker behaviour, enabling analysts to understand not only how a compromise could occur but also which adversary tactics and techniques would be exercised during the attack sequence. This behavioural perspective supports threat-informed defence by facilitating stronger alignment between automated attack simulation and real-world adversary behaviour.

3.5 Evaluation Metrics

To evaluate the effectiveness of integrating MITRE ATT&CK into HARM-based automated attack simulation, four evaluation indicators were defined. These indicators were selected to assess behavioural enrichment while ensuring that the automation capabilities of the original HARMer framework were preserved.

The first indicator was Mapping Coverage Rate (MCR), defined as the percentage of identified CVEs that could be successfully mapped to one or more ATT&CK techniques. This metric was used to assess the feasibility of enriching vulnerability information with behavioural context.

The second indicator was Behavioural Annotation Rate (BAR), representing the percentage of generated attack-path steps that contained ATT&CK identifiers. This metric measured the extent to which behavioural information was incorporated into attack planning outputs.

The third indicator was Technique Coverage (TC), which measured the number of unique ATT&CK techniques represented across generated attack paths. Higher coverage indicated broader representation of adversary behaviour within the simulation environment.

The fourth indicator was Attack Path Preservation (APP), defined as the extent to which ATT&CK enrichment preserved the original attack paths generated by HARMer. This indicator was used to verify that behavioural integration enhanced interpretability without altering path feasibility or planning logic.

Together, these indicators provided a structured basis for evaluating behavioural enhancement while maintaining consistency with the original HARM-based attack simulation process.

4. Results and discussions

This section presents and discusses the results obtained from integrating the MITRE ATT&CK framework into a HARM-based automated attack simulation pipeline. The results focus on the behavioural interpretability of simulation outputs rather than exploit execution performance. Findings are discussed in relation to the research objectives and research questions defined in Section 1.

4.1 Results of CVE-to-ATT&CK Mapping

The first objective of this study was to determine whether vulnerabilities identified during automated scanning could be systematically enriched with behavioural information using the MITRE ATT&CK framework. The results demonstrated that a substantial proportion of CVE-based vulnerabilities could be associated with ATT&CK tactics and techniques based on their documented exploitation behaviour.

The Mapping Coverage Rate (MCR) analysis showed that behavioural enrichment was feasible across multiple categories of vulnerabilities. Vulnerabilities enabling unauthorised access to systems were commonly associated with techniques under the Initial Access tactic, while privilege escalation vulnerabilities were linked to techniques within the Privilege Escalation tactic. Similarly, vulnerabilities supporting remote code execution

were associated with Execution-related techniques, providing additional behavioural context beyond the original CVE descriptions.

Not all vulnerabilities were mapped to ATT&CK techniques. Vulnerabilities lacking sufficient behavioural detail, ambiguous exploitation outcomes, or limited supporting documentation were excluded from the mapping process. This selective approach reduced the likelihood of inaccurate behavioural attribution and ensured that ATT&CK associations remained evidence-based.

The mapping process demonstrated that vulnerability information can be systematically transformed from a purely technical representation into a behavioural representation without altering the underlying vulnerability discovery process. These findings directly address Research Question 1 by demonstrating that CVE-based vulnerabilities can be consistently associated with adversary tactics and techniques, thereby improving the interpretability of automated attack simulation outputs.

4.2 Results of Behaviour-Enhanced Attack Planning

The second objective of this research was to evaluate whether ATT&CK identifiers could be incorporated into HARM-based attack planning to improve the interpretability of generated attack paths. The results demonstrated that behavioural information could be successfully integrated into the planning process while preserving the automated path-generation capabilities of the original HARMer framework.

In the baseline HARMer approach, attack paths are represented primarily as sequences of hosts, vulnerabilities, and exploits. While these paths provide valuable technical information, they offer limited insight into the behaviour of an adversary during an attack. Following ATT&CK integration, generated attack paths retained their technical structure but additionally included behavioural annotations derived from mapped ATT&CK tactics and techniques.

The Behavioural Annotation Rate (BAR) indicated that ATT&CK information could be propagated throughout the attack-planning process, allowing attack paths to be interpreted as sequences of adversary actions rather than isolated exploit events. For example, a path could be understood as progressing through stages such as Initial Access, Execution, Privilege Escalation, and Lateral Movement, providing a clearer representation of attacker behaviour within the network.

Importantly, the Attack Path Preservation (APP) assessment showed that behavioural enrichment did not alter the feasibility of generated attack paths. Network reachability, vulnerability information, and exploit availability continued to determine path selection, consistent with the original HARMer methodology. ATT&CK integration therefore functioned as an additional semantic layer rather than a replacement for existing planning mechanisms.

These findings directly address Research Question 2 and demonstrate that ATT&CK can be embedded into HARM-based attack planning without compromising automation or scalability. The resulting attack paths provide improved behavioural interpretability while preserving the technical accuracy and repeatability of automated attack simulation.

4.3 Results Collection and Visualisation

The final objective of this research was to evaluate whether behavioural information generated during automated attack simulation could be effectively visualised using ATT&CK-based representations. Following attack-path generation, ATT&CK techniques associated with mapped vulnerabilities were aggregated and visualised using ATT&CK Navigator-style matrix views. This enabled behavioural information to be represented directly from simulation outputs rather than being added retrospectively during analysis.

The Technique Coverage (TC) evaluation demonstrated that multiple ATT&CK techniques could be represented across simulated attack scenarios, providing a consolidated view of adversary behaviour within the experimental environment. Rather than examining individual vulnerabilities or attack-path records, analysts could observe which tactics and techniques were exercised during the simulation and identify areas of behavioural coverage within the ATT&CK framework.

Figure 1 presents the HARM representation used to generate attack paths, while Figure 2 illustrates an example ATT&CK-based visualisation derived from the simulation results. The visualisation provided a concise summary of behavioural activity and enabled direct tracing between ATT&CK techniques, associated vulnerabilities, and generated attack paths.

Compared with conventional vulnerability reports or raw attack logs, ATT&CK-based visualisation offered a more intuitive representation of attacker behaviour. The approach improved accessibility for analysts and supported clearer communication of attack outcomes within a threat-informed defence context.

These findings directly address Research Question 3 and demonstrate that ATT&CK-based visualisation can serve as a practical output of automated attack simulation, enhancing the interpretability of behavioural attack data generated through HARM-based automation.

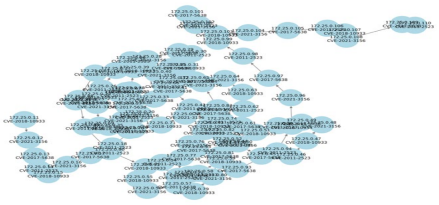

Figure 1: HARM Graph representation of vulnerabilities

4.4 Results of ATT&CK-Based Visualisation

The final set of results concerns the visualisation of simulated attacks using ATT&CK-based representations. Techniques exercised during automated attack simulations were aggregated and visualised using an ATT&CK matrix view, providing a concise overview of behavioural coverage across experiments. Figure 2 presents an example ATT&CK-based visualisation generated from the experimental results.

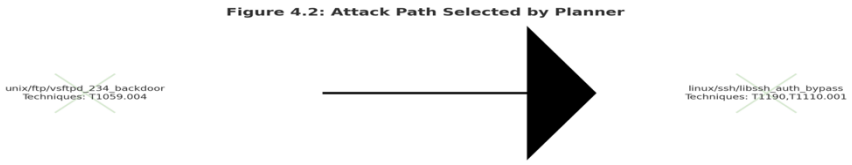

Figure 2: Attack path selected by planner

The visualisation enabled rapid identification of which adversary techniques were exercised and which areas of the ATT&CK matrix remained unrepresented. Compared to raw logs or structured output files, ATT&CK-based visualisation provided a higher-level and more intuitive summary of attack behaviour. Each highlighted technique could be traced back to specific vulnerabilities and attack path steps, supporting transparent analysis.

These results address the third research question and demonstrate that ATT&CK-based visualisation can function as a direct output of automated attack simulation rather than as a post-processing step.

4.5 Discussion

The results demonstrate that integrating MITRE ATT&CK into a HARM-based automated attack simulation framework improves the behavioural interpretability of simulation outputs while preserving the automation capabilities of the original HARMer architecture. Across the three research objectives, ATT&CK integration provided a consistent mechanism for enriching vulnerability information, attack planning, and attack visualisation with adversary-focused context.

A key observation from the study is that traditional HARM-based attack simulation and ATT&CK-based threat modelling address complementary aspects of security assessment. HARM provides an effective mechanism for modelling network reachability and identifying technically feasible attack paths, whereas ATT&CK provides a structured representation of adversary behaviour. By combining these approaches, the enhanced framework retains the technical accuracy of vulnerability-driven attack simulation while providing behavioural explanations for simulated attack activity.

Compared with baseline HARMer, the enhanced framework offers several advantages. First, vulnerabilities can be interpreted as enablers of adversary techniques rather than isolated technical weaknesses. Second, generated attack paths can be understood as behavioural attack narratives that align with recognised ATT&CK tactics. Third, ATT&CK-based visualisation provides a higher-level representation of attack activity that supports communication between technical analysts, security managers, and decision-makers.

Importantly, the integration of ATT&CK did not require modification of the core HARM modelling principles or automated planning mechanisms. Network reachability analysis, vulnerability discovery, and attack-path

generation remained unchanged, demonstrating that behavioural enrichment can be achieved without sacrificing scalability or repeatability. This finding is particularly important because it suggests that behavioural threat intelligence can be embedded within existing automated attack simulation frameworks rather than requiring entirely new architectures.

The findings therefore support the argument that automated red-teaming can evolve beyond vulnerability-centric assessment towards a more threat-informed approach that reflects how adversaries operate in practice. By linking technical vulnerabilities with behavioural intelligence, the proposed enhancement contributes to more interpretable and operationally relevant security assessment.

4.6 Limitations and Future Work

Although the results demonstrate the feasibility of integrating MITRE ATT&CK into HARM-based automated attack simulation, several limitations should be acknowledged. First, the experimental evaluation was conducted within controlled and isolated environments designed to support repeatability and ethical testing. While suitable for demonstrating proof of concept, these environments may not fully capture the complexity, scale, and dynamic behaviour of large enterprise networks.

Second, the CVE-to-ATT&CK mapping process relied on publicly available vulnerability descriptions, exploit documentation, and ATT&CK definitions. Although validation procedures were applied to reduce ambiguity, some level of expert judgement remained necessary when selecting the most representative ATT&CK technique. Future work could investigate automated mapping approaches using threat intelligence platforms and machine learning techniques to improve consistency and scalability.

Third, the evaluation focused primarily on behavioural interpretability rather than operational performance. The study demonstrated that ATT&CK integration could enrich attack simulation outputs, but it did not measure analyst decision-making, detection effectiveness, or defensive outcomes. Future research could incorporate user studies to evaluate whether ATT&CK-enhanced simulations improve analyst understanding and support more effective threat-informed defence activities.

Finally, future work could extend the framework to support ATT&CK sub-techniques, adversary emulation profiles, and larger-scale cloud environments. Such enhancements would provide a more detailed representation of attacker behaviour and enable broader evaluation of behavioural attack simulation across diverse operational scenarios.

5. Conclusion

This paper presented an experimental investigation into the integration of the MITRE ATT&CK framework within a HARM-based automated attack simulation environment. The motivation for the study was the limited behavioural interpretability of existing vulnerability-centric attack simulation approaches, which often represent attacks as sequences of technical exploits without clearly conveying adversary behaviour.

To address this limitation, the HARMer framework was enhanced by incorporating ATT&CK tactics and techniques into vulnerability modelling, attack planning, and simulation visualisation. Vulnerabilities identified through automated scanning were systematically mapped to ATT&CK techniques, enabling behavioural information to be embedded directly within the attack simulation process. The resulting attack paths retained the technical accuracy of HARM-based modelling while providing additional insight into attacker objectives, tactics, and techniques.

The evaluation demonstrated that behavioural enrichment could be achieved without altering the underlying attack-path generation mechanisms of HARMer. Mapping Coverage Rate (MCR), Behavioural Annotation Rate (BAR), Technique Coverage (TC), and Attack Path Preservation (APP) provided a structured framework for assessing the effectiveness of ATT&CK integration. The findings showed that behavioural context could be consistently incorporated into automated attack simulation outputs while preserving automation, repeatability, and scalability.

The principal contribution of this research is the demonstration that behavioural threat intelligence can be embedded directly into HARM-based automation rather than applied only as a post-analysis activity. By combining vulnerability-driven attack modelling with ATT&CK-based behavioural semantics, the proposed approach advances automated red-teaming towards a more interpretable and threat-informed form of security assessment.

Future research should evaluate the approach within larger and more diverse environments, investigate automated ATT&CK mapping techniques, and assess the impact of behaviour-enhanced simulation on analyst decision-making and defensive planning. Such developments would further strengthen the practical applicability of ATT&CK-integrated automated attack simulation for modern cybersecurity operations.

Acknowledgements

I will love to thank my parents, Mr/Mrs Buhnyuy, for their help, love and support throughout my MSc studies

AI Declaration I confirm that I have read and understood the University's policy on the use of Artificial Intelligence in academic work. Where relevant, I have acknowledged and referenced any use of AI tools in accordance with the University's Academic Regulations.

I understand that failure to appropriately reference the use of AI constitutes an assessment offence equivalent to plagiarism and may result in penalties recorded on my student record.

Details of AI Use: AI tools were used solely for formatting and debugging code, improving clarity of written English, and checking grammar and spelling. All substantive ideas, analysis, research, and written content are my own original wo**rk**

Certificate of Ethics Review

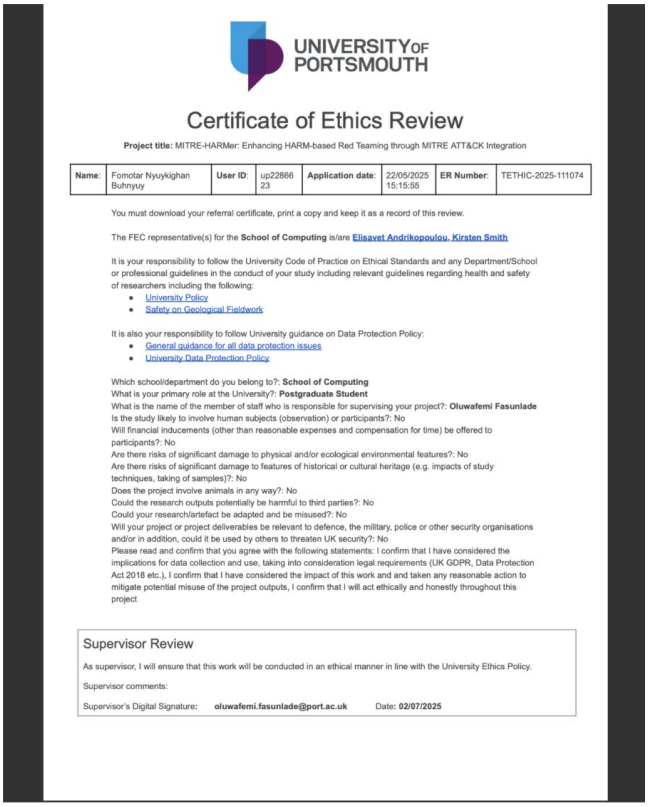

UNIVERSITY OF PORTSMOUTH

Certificate of Ethics Review

Project title: MITRE-HARMer: Enhancing HARM-based Red Teaming through MITRE ATT&CK Integration

Name:	Fomotar Nyuykighan Buhnyuy	User ID:	up2286623	Application date:	22/05/2025 15:15:55	ER Number:	TETHIC-2025-111074

You must download your referral certificate, print a copy and keep it as a record of this review.

The FEC representative(s) for the **School of Computing** is/are **Elisavet Andrikopoulou, Kirsten Smith**

It is your responsibility to follow the University Code of Practice on Ethical Standards and any Department/School or professional guidelines in the conduct of your study including relevant guidelines regarding health and safety of researchers including the following:

- University Policy
- Safety on Geological Fieldwork

It is also your responsibility to follow University guidance on Data Protection Policy:

- General guidance for all data protection issues
- University Data Protection Policy

Which school/department do you belong to?: **School of Computing**
What is your primary role at the University?: **Postgraduate Student**
What is the name of the member of staff who is responsible for supervising your project?: **Oluwafemi Fasunlade**
Is the study likely to involve human subjects (observation) or participants?: No
Will financial inducements (other than reasonable expenses and compensation for time) be offered to participants?: No
Are there risks of significant damage to physical and/or ecological environmental features?: No
Are there risks of significant damage to features of historical or cultural heritage (e.g. impacts of study techniques, taking of samples)?: No
Does the project involve animals in any way?: No
Could the research outputs potentially be harmful to third parties?: No
Could your research/artefact be adapted and be misused?: No
Will your project or project deliverables be relevant to defence, the military, police or other security organisations and/or in addition, could it be used by others to threaten UK security?: No
Please read and confirm that you agree with the following statements: I confirm that I have considered the implications for data collection and use, taking into consideration legal requirements (UK GDPR, Data Protection Act 2018 etc.), I confirm that I have considered the impact of this work and and taken any reasonable action to mitigate potential misuse of the project outputs, I confirm that I will act ethically and honestly throughout this project

Supervisor Review

As supervisor, I will ensure that this work will be conducted in an ethical manner in line with the University Ethics Policy.

Supervisor comments:

Supervisor's Digital Signature: **oluwafemi.fasunlade@port.ac.uk** Date: **02/07/2025**

References

Al-Sada, B., Sadighian, A., & Oligeri, G. (2024). MITRE ATT&CK: State of the art and way forward. *ACM Computing Surveys*, *57*(1), 1-37.

Enoch, S. Y., Huang, Z., Moon, C. Y., Lee, D., Ahn, M. K., & Kim, D. S. (2020). HARMer: Cyber-attacks automation and evaluation. *IEEE Access*, *8*, 129397-129414.

Hong, J., & Kim, D. S. (2012). HARMs: Hierarchical attack representation models for network security analysis.

Jiang, Y., Meng, Q., Shang, F., Oo, N., Minh, L. T. H., Lim, H. W., & Sikdar, B. (2025). MITRE ATT&CK Applications in Cybersecurity and The Way Forward. *arXiv preprint arXiv:2502.10825*.

Jiang, Yuning, Qiaoran Meng, Feiyang Shang, Nay Oo, Le Thi Hong Minh, Hoon Wei Lim, and Biplab Sikdar. "MITRE ATT&CK Applications in Cybersecurity and The Way Forward." *arXiv preprint arXiv:2502.10825* (2025).

MITRE ATT&CK® (n.d.) *MITRE ATT&CK framework*. Available at: https://attack.mitre.org/