AI-generated Papers and Manipulation of Academic Metrics: A Case Study

Haitham S. Al-Sinani(ORCID: 0009-0005-0453-3335)
Department of Cybersecurity and Quality Assurance
Diwan of Royal Court, Muscat, Oman
hsssinani@diwan.gov.om

Chris J. Mitchell(ORCID: 0000-0002-6118-0055)
Department of Information Security
Royal Holloway, University of London, UK
c.mitchell@rhul.ac.uk

***Abstract*—This paper investigates the presence and impact of questionable, AI-generated academic papers on widely used preprint repositories, with a focus on their role in citation manipulation. Motivated by suspicious patterns observed in publications related to our ongoing research on GenAI-enhanced cybersecurity, we identified clusters of questionable papers and profiles. These papers frequently exhibited minimal technical content, repetitive structure, unverifiable authorship, and mutually reinforcing citation patterns among a recurring set of authors. To assess the feasibility and implications of such practices, we conducted a controlled experiment: generating a fake paper using generative AI (GenAI), embedding citations to suspected questionable publications, and uploading it to one such repository (ResearchGate). Our findings demonstrate that such papers can bypass platform checks, remain publicly accessible, and contribute to inflating citation metrics like the H-index and i10-index. We present a detailed analysis of the mechanisms involved, highlight systemic weaknesses in content moderation, and offer recommendations for improving platform accountability and preserving academic integrity in the age of GenAI.**

***Index Terms*—ResearchGate, AI-generated Articles, GenAI, Academic Integrity, H-index, Citation Manipulation**

I. Introduction

Online academic platforms such as ResearchGate have transformed scholarly communication by enabling researchers to freely share publications, track impact, and engage with a global audience. Their open-access nature and user-driven upload features have democratised knowledge dissemination and expanded the visibility of scientific output. However, this openness has also introduced vulnerabilities. Our investigations suggest that large quantities of AI-generated content are being uploaded to widely used preprint repositories, possibly with the goal of artificially inflating key metrics for academics, including citation counts, the H-index and the i10-index.

Citation-based metrics such as the *H-index* and *i10-index* are commonly used to assess a researcher's academic impact. The H-index is defined as the maximum value h such that the researcher has published h papers each of which has been cited at least h times. It attempts to capture both productivity and citation impact in a single number. The i10-index, introduced by Google Scholar, is a simpler metric that counts the number of a researcher's publications with at least ten citations. While these metrics are convenient and widely adopted in academic evaluations, they are also vulnerable to manipulation—especially in open publishing environments where content vetting is minimal.

This emerging malpractice is especially troubling in the current technological landscape, where GenAI systems, such as ChatGPT[1], Gemini[2] and DeepSeek[3], can effortlessly produce texts that mimic academic style, structure, and tone. With minimal oversight, these synthetic documents can appear convincing enough to be mistaken for legitimate scholarship, especially on platforms where submissions are not peer-reviewed or institutionally validated.

Our investigation is motivated by a fundamental question: *Are there fully AI-generated papers on ResearchGate, particularly in the area of GenAI-enhanced Cybersecurity, and if so, how are they being used to manipulate academic metrics?* Unlike earlier studies [1], [2] that focused on peer-reviewed journals or high-profile conferences, we focus on ResearchGate—a popular but lightly moderated platform where researchers frequently upload preprints, reports, and presentation materials. The consequences of citation manipulation here may be less visible than in traditional publishing, but they are no less significant—particularly because ResearchGate documents are assigned DOIs (Digital Object Identifiers) and often indexed by popular systems like Google Scholar.

Our interest in this topic arose through our own line of research on applying GenAI in the field of ethical hacking. While monitoring related publications via Google Scholar alerts, we encountered a series of papers that appeared, on the surface, to be relevant. However, closer inspection revealed they were strikingly similar in content, cited many of the same individuals, and lacked empirical foundations. These observations raised suspicions that the papers were generated using GenAI and published with the intent of artificially inflating citation metrics.

Motivated by these observations, we conducted a multi-phase investigation to explore this phenomenon more systematically. We began by mapping out and reviewing a cluster of suspicious papers and the ResearchGate profiles associated with them. Building on these findings, we designed and carried out a controlled experiment: generating a fake academic paper

[1]https://chatgpt.com/
[2]https://gemini.google.com/app
[3]https://chat.deepseek.com/

using GenAI, embedding citations to several known suspect papers, and uploading it to ResearchGate under fictional authorship. By observing whether our paper is accepted by the platform and whether it contributes to boosting citation metrics of the referenced papers, we assess the ease with which such manipulation can occur. Our final analysis includes a comparison of multiple ResearchGate profiles suspected of participating in citation inflation schemes.

The goal of this work is not only to document the mechanisms of citation manipulation, but also to raise awareness and prompt a broader discussion on academic integrity in the GenAI age. We present a detailed case study of how generative tools and platform design can interact in ways that compromise the credibility of scholarly metrics. In doing so, we contribute to the growing body of work examining the risks of AI-generated content in academia, and we offer preliminary recommendations for mitigating such threats.

This work is timely in that it addresses a growing and underexamined threat to research evaluation systems—namely, the deliberate misuse of GenAI and academic platforms to fabricate publications and manipulate citation metrics. By combining empirical observation with an experimental case study, we provide concrete evidence of how such manipulation can be carried out on ResearchGate. Our findings not only expose systemic gaps in platform oversight but also underscore the need for greater scrutiny of citation-based indicators that are often used in hiring, promotion, and funding decisions.

The remainder of this paper is organised as follows. Section II introduces the phenomenon of questionable papers, and, section III outlines the core motivations behind their generation. Section IV presents our experimental methodology, and, section V discusses the broader implications of our findings and provides policy-level recommendations. Section VI reviews relevant prior studies. Finally, section VII concludes the paper and outlines directions for future research.

II. The Phenomenon of Questionable Papers

Over the past couple of years, we have been conducting authentic research on the use of GenAI in ethical hacking. Our ultimate goal has been to enhance and automate the often manual and tedious penetration testing processes. We have published our findings and shared our advancements with the research community [3]–[7].

Interestingly, while following up on similar topics via Google Scholar alerts, we were recommended a group of articles that appeared —on the surface— to be highly relevant. However, upon close inspection, these papers struck us as being extremely thin in content and lacking technical depth or academic rigour. Based on their language, structure, and lack of empirical evidence, we strongly believe that these papers were entirely AI-generated and, in essence, fake.

To better understand the nature of the suspected potentially AI-generated papers, we performed a close textual and structural analysis of representative articles recommended to us via Google Scholar. At first glance, many of these papers appeared relevant to our ongoing work on AI and ethical hacking. They were well-formatted, readable, and addressed familiar themes. However, upon deeper examination, several red flags became apparent, strongly suggesting that these papers were either entirely AI-generated or lacked genuine academic rigour.

A. Key Concerns

1) **Many Published in a Single Snapshot in Time.**
   A striking observation was the publication date clustering. For instance, a significant number of these papers were released in the same month and year (e.g. February 2025 [8]–[22]), despite being attributed to different authors. This synchronicity implies the likelihood of batch generation and coordinated mass upload rather than independent scholarly efforts.
2) **Strikingly Similar Content.**
   The papers displayed remarkably similar language, structure, and narrative flow. Phrases such as "AI-powered penetration testing enhances efficiency, accuracy, and scalability" were frequently reused across multiple documents, often with minimal variation. This degree of overlap is characteristic of AI-generated output or templated production rather than distinct academic voices.
3) **Recycled and Dubious References.**
   The reference sections across these papers were nearly identical, containing overlapping citations with questionable relevance. Names like "Mohammed, A." and "Żywiołek, J." appeared recurrently, often in contexts that bore little relation to the main topic. Several citations mimicked DOI-style formatting but could not be verified via trusted academic databases, raising concerns about citation fabrication.
4) **Absence of Methodology or Empirical Work.**
   The analysed papers contained no experimental sections, case studies, or data-driven evaluations. Their claims —such as "AI can automate ethical hacking tasks"— were presented without evidence or validation. The lack of methodology undermines their academic contribution.
5) **No Engagement with Existing Literature.**
   Unlike authentic research papers, which usually position their contributions within the existing body of work, these questionable papers made no attempt to position their arguments within the broader academic discourse. There was no literature review or critical engagement with prior work, indicating a lack of scholarly depth and awareness.
6) **Unverifiable Author Identities.**
   Basic searches for many of the listed authors yielded no academic profiles, institutional affiliations, or presence on reputable platforms like Google Scholar.
7) **Predictable Filename Patterns.**
   We observed that many of these papers follow a noticeable pattern in their file naming conventions on ResearchGate, often using sequential and predictable numbering (e.g. 17, 18, 19, 20). This implies an automated or systematic upload process, further supporting the hypothesis of batch generation and artificial publication activity.

8) **Non-Scientific Writing Style.**
While the papers maintained surface-level readability, their tone resembled that of promotional blog posts rather than scientific discourse. They relied heavily on buzzwords like "AI revolutionising cybersecurity," yet offered little technical specificity or conceptual clarity.

B. Final Verdict

Based on these findings, we conclude that many of the papers in question are either AI-generated or strategically crafted with minimal effort for the sole purpose of inflating academic metrics. Under standard peer review, these papers would likely be rejected for lacking originality, methodology, and scientific value. Their presence on platforms like ResearchGate not only dilutes the quality of academic discourse but also highlights systemic vulnerabilities in how such platforms manage, curate, and verify scholarly content.

III. Possible Motivations for Fabricated Paper Generation

We now pose the question: Why are such papers being generated at such an alarming rate? Our hypothesis is that they exist solely to artificially inflate citation-based academic metrics, especially the H-index.

One example of where an academic's publication metrics appear to have benefitted from the upload of large numbers of questionable papers is as follows. We should point out that it may be entirely accidental that the individual concerned has benefitted in this way, since we cannot say for sure why the papers have been created and uploaded; moreover, there are almost certainly many other cases we could have highlighted. Our example case is that of Anwar Mohammed, listed on Google Scholar as affiliated with Singhania University, Rajasthan, India. His profile[4] shows that a number of recently published articles of which he is an author have been cited in a number of apparently questionable preprints uploaded to ResearchGate. These preprints were automatically processed by Google Scholar, significantly changing its computation of metrics such as the H-index and i10-index.

This discovery strongly supports our hypothesis and prompted us to continue the investigation. Using Google Scholar's reverse citation feature, we were able to uncover a chain of related questionable papers uploaded to ResearchGate. Our findings revealed a consistent pattern involving many ResearchGate members. Notably, we identified four illustrative cases —Spunda[5], Pomeroy[6], Gatlin[7], and Roseth[8]— each linked to a series of suspicious publications, many dated 2025 [8]–[22], suggesting a deliberate and repeated engagement in questionable authorship practices. Even more telling was the consistent authorship model:

- **First author:** Typically an unfamiliar name, likely fictitious, often appearing only once or twice across the literature.
- **Second author:** A ResearchGate member with an active profile and prior publication history—frequently the actual uploader of the questionable paper.

To demonstrate this recurring pattern, we compiled a comparison table (see table I) with representative examples. Each case highlights how the second author role is used to anchor the fake paper with a ResearchGate-verified identity, while the likely fictitious first author provides plausible deniability.

TABLE I
Comparison of Suspected Questionable ResearchGate Profiles

Feature	**Rudolf Spunda**	**Jennifer Pomeroy**
ResearchGate Profile	https://www.researchgate.net/profile/Rudolf-Spunda-2/	https://www.researchgate.net/profile/Jennifer-Pomeroy-3/
Total Publications on ResearchGate	10	10
Solo-Authored Papers	4	5
Papers as 2nd Author	6	5
Ephemeral 1st Authors	Fani Sani, Kuldeep Rahul, Abhishek Sharma, Muhammad Shamir, Yasir Nawaz, Haider Ali, Muhammad Nasir, Abrar Ahmed, Mukesh Kumar, Asia Rehman, Sadia Farooq, Halima Sinisa	Umair Zafer, Andre Russell, Ravi Bishnoi, Talaat Hayat, Nimra Fahad, Lennon Zahir
2nd Author Name (ResearchGate Member)	Rudolf Spunda	Jennifer Pomeroy
Profile Upload Activity	Uploaded by Rudolf	Uploaded by Jennifer
Citation Count (on ResearchGate)	1	Not specified
Notes / Suspicious Behaviour	Uses ResearchGate to legitimise papers by collaborating with obscure authors	Repeats similar publication strategy across random names

IV. Preliminary Experiment

Given the gravity of the implications uncovered during our investigation, we designed a controlled experiment to determine whether the suspicious behaviours observed on ResearchGate could be deliberately reproduced and exploited using GenAI. Our aim was not only to test the platform's safeguards but also to empirically demonstrate the mechanisms by which citation manipulation might occur in practice.

A. Objectives

The experiment was guided by the following objectives:

1) To test whether a fake academic paper could be easily generated entirely by GenAI.
2) To test whether a fake, AI-generated academic paper could be successfully uploaded to ResearchGate.

[4]https://scholar.google.com/citations?hl=en&user=gpUs8nkAAAAJ
[5]https://www.researchgate.net/profile/Rudolf-Spunda-2/
[6]https://www.researchgate.net/profile/Jennifer-Pomeroy-3/
[7]https://www.researchgate.net/profile/Kaiser-Gatlin/
[8]https://www.researchgate.net/profile/Tom-Roseth-3/

3) To assess whether this paper, by citing other suspected questionable papers, could contribute to inflating citation metrics such as the H-index for their listed authors.
4) To observe whether ResearchGate has any detection mechanisms or moderation workflows capable of identifying and removing inauthentic or AI-generated content.

B. Paper Generation Using GenAI

We began by using GenAI tools (ChatGPT[9] in particular) to generate a completely fictional academic paper titled *"GenAI in Digital Forensics: Enhancing Timeline Reconstruction and Anomaly Attribution"*[10]. The content was generated with minimal human editing to preserve the authenticity of the experiment and to reflect what a typical misuse of GenAI might look like in a real-world setting. The output was fluent, well-structured, and mimicked the format of academic writing.

To ensure transparency, we embedded two disclaimers in the document: one on the title page; and another at the end. Both disclaimers clearly stated that the paper was generated using GenAI for academic experimentation and research purposes only, with no intent to deceive.

The generated paper lacked empirical evidence, literature review, or technical depth, intentionally mirroring the characteristics we had previously observed in suspected fake publications. The language was generalised, and the conclusions were deliberately vague, echoing the stylistic markers of questionable submissions we had reviewed.

C. Submission Process and Citation Strategy

Once the paper was prepared, we proceeded to upload it to ResearchGate. One key challenge we encountered was that ResearchGate required at least one of the listed authors to be associated with an existing and verified ResearchGate profile. To bypass this requirement without linking the paper to our real identities, we introduced a third fictitious author, *Naif Al-Sinani*, sharing a surname with the uploader. This manoeuvre satisfied the platform's verification checks.

Although the uploaded PDF listed only fictitious names as authors, ResearchGate automatically linked the document to the real profile used for submission, revealing a structural flaw in its authorship verification system. At no point did the platform request institutional credentials or validate author identities against external academic databases.

To simulate real-world citation manipulation strategies, we embedded citations to at least three previously identified questionable papers, each authored or co-authored by ResearchGate users Spunda, Pomeroy, Gatlin, and Roseth. This was done deliberately to test whether our fake paper would successfully contribute to boosting their citation metrics.

D. Initial Results and Observations

At the time of writing, the fake paper[11] remains publicly accessible and has not been flagged, moderated, or removed by ResearchGate. Notably, within 24 hours of publication, we observed an increase in the citation counts of the questionable papers cited, suggesting that ResearchGate's impact metrics are responsive to such fabricated references, thereby enabling H-index manipulation through a feedback loop.

E. Limitations of the Study

This study presents a single-case experiment using one GenAI-generated paper over a short timeframe, limiting generalisability across disciplines or broader contexts. We did not scale the experiment using variations in content or citation strategies, nor did we test self-citation effects, as we avoided referencing our own work for ethical reasons. Observed citation increases may be influenced by factors beyond our control, such as platform indexing algorithms. Additionally, findings are specific to ResearchGate and may not extend to other repositories like arXiv or Academia.edu. Nonetheless, the experiment demonstrates that ResearchGate's current infrastructure can be exploited for citation manipulation, warranting further investigation through broader, longitudinal studies.

F. Ethical Considerations

To remain within acceptable ethical boundaries, we deliberately chose not to expand the experiment beyond a minimal intervention. No real identities, publications, or institutional affiliations were cited, and disclaimers were added for full transparency. Indeed, none of our real publications or identities were cited to avoid personal metric inflation.

Although the experiment was limited in scope, the results are sufficient to confirm that ResearchGate's current platform design can be exploited to generate and circulate fake academic content with measurable impact on citation-based metrics. We plan to inform relevant stakeholders, including ResearchGate and Google Scholar, and will consider further controlled experiments in collaboration with academic institutions and ethical review boards.

V. Implications and Recommendations

The proliferation of fake academic papers poses a direct threat to the credibility of digital scholarly platforms and the academic community at large. Our findings reveal that individuals are exploiting the goodwill and open-access ethos of platforms like ResearchGate to artificially boost their academic profiles. By generating low-quality or AI-produced content and embedding self-referential citation loops, these actors are able to manipulate key academic metrics —especially the H-index— without contributing any real scholarly value.

Such practices erode the trust that underpins academic publishing. Platforms like ResearchGate, which were created to democratise access to research and foster global collaboration, are now being used as tools for metric-based self-promotion. This undermines fair academic competition and distorts the scholarly record, making it harder for genuine research to stand out amid fabricated content.

[9]https://chatgpt.com/
[10]http://dx.doi.org/10.13140/RG.2.2.35588.23688
[11]http://dx.doi.org/10.13140/RG.2.2.35588.23688

A. Impact on Academic Metrics: H-index Gaming

The central risk lies in the ability of fake papers to distort citation-based impact indicators. Since ResearchGate assigns DOIs to uploaded papers —and these DOIs are often indexed by third-party platforms such as Google Scholar— citations from these questionable documents contribute directly to the H-index and i10-index of the cited authors, regardless of the citing paper's authenticity.

This phenomenon was clearly observed in the case of Anwar Mohammed, whose Google Scholar profile shows a disproportionate spike in citations originating from suspicious ResearchGate uploads, most dated 2024–2025. These documents follow consistent structural and stylistic patterns and frequently cite Mohammed's earlier work, despite lacking any thematic relevance or scholarly depth.

Such gaming of citation metrics not only misrepresents the academic influence of certain individuals but also skews global rankings, hiring decisions, grant evaluations, and editorial opportunities, many of which still rely on bibliometric indicators.

B. Consequences for Academic Integrity and Content Quality

The spread of fake publications has broader consequences beyond individual metric inflation. It compromises the overall quality of accessible research and confuses scholars —especially students and early-career researchers— who may struggle to differentiate legitimate contributions from AI-generated or fraudulent work. The appearance of scholarly legitimacy, conferred by DOIs, academic formatting, and citation counts, can mask a complete absence of methodological rigour or empirical contribution.

Moreover, the presence of unverified and unauthored content on scholarly platforms creates an uneven playing field, devalues peer-reviewed research, and undermines institutional trust in open repositories. If left unchecked, this trend could trigger a larger crisis of credibility for digital knowledge platforms.

C. Ethical Considerations and Platform Responsibility

Our investigation raises serious ethical concerns:

- Should platforms like ResearchGate be permitted to assign DOIs to unverified, unreviewed documents?
- Is it appropriate for citation metrics to be derived from documents that have not undergone any form of academic vetting?
- Do platforms have a duty to proactively prevent the use of GenAI for deceptive academic practices?

While our own experiment was conducted transparently and ethically —with disclaimers and no self-citations— the very fact that it was feasible demonstrates how easily such systems can be abused by bad actors. The onus now falls on platform providers to implement better safeguards, and on the academic community to hold such platforms accountable.

D. Policy Recommendations

In light of the findings presented in this study, we propose a series of platform-level and community-wide reforms:

- **Stricter authorship verification**: Require institutional email addresses or ORCID verification for all listed authors before permitting uploads.
- **Content authentication systems**: Develop automated tools to flag GenAI-generated submissions using linguistic fingerprinting and structural pattern recognition.
- **Transparency labelling**: Introduce tags or visibility filters to identify papers marked as 'experimental,' 'non-peer-reviewed,' or 'AI-generated,' and exclude them from impact metric calculations.
- **Exclude unfiltered preprints from citation metrics**: Google Scholar and other major sites providing information about research publications should stop using unfiltered preprints as a source for its computation of citation counts and metrics such as the H-index. Otherwise they become completely useless.
- **Ethical review for mass uploads**: Require flagged high-volume uploaders to undergo manual review, especially in cases where sequential or templated publication behaviour is detected.
- **Community reporting mechanisms**: Enable and encourage users to flag suspicious documents or profiles for review, with transparent follow-up by platform moderators.
- **Preserving the integrity of DOIs**: Platforms should refrain from assigning DOIs to unvetted or automatically uploaded preprints. The indiscriminate use of DOIs in such cases dilutes their value and creates a false perception of scholarly legitimacy.

VI. Related Work

Recent work has demonstrated that citation counts on Google Scholar can be inflated through paid services, highlighting a new and concerning form of metric manipulation [1]. This practice contaminates databases and misleads those who rely on citation metrics for scholarly evaluation.

The medical community has also explored the risks of AI-generated publications. Májovský et al. [2] showed how large language models (LLMs) like ChatGPT can be used to produce plausible but entirely fraudulent medical research articles. These papers were well-structured and readable, yet lacked any empirical grounding or peer-reviewed validation.

Earlier manifestations of this problem were seen with SCIgen, a program developed to auto-generate academic-looking computer science papers [23]. Though initially satirical, SCIgen's outputs were accepted at real conferences, providing early evidence that even minimal scrutiny could fail to detect synthetically generated nonsense.

This paper complements this prior work by providing a focused, experimental investigation into how AI-generated papers can be used to manipulate citation metrics on ResearchGate. We first identified clusters of questionable, AI-generated papers in a domain closely aligned with our own ongoing research on AI in ethical hacking. Notably, we observed that these papers disproportionately cited the work of a key beneficiary —Anwar Mohammed of Singhania University— whose citation metrics appear to benefit significantly from this

citation pattern. Through a controlled upload and citation strategy, we highlight specific weaknesses in platform moderation and contribute empirical evidence to the broader discussion on academic integrity in the age of GenAI.

VII. Conclusions and Directions for Future Work

This study investigated the presence and implications of questionable, AI-generated academic papers on widely used preprint repositories, such as ResearchGate, with a particular focus on their role in citation manipulation. Through detailed analysis and a controlled experiment, we confirmed that it is both technically and procedurally feasible to upload a synthetic paper generated using GenAI, link it to a legitimate ResearchGate profile, and use it to artificially boost the citation counts of other suspect publications. Our findings highlight structural weaknesses in platform moderation and reveal how these vulnerabilities can be exploited to manipulate bibliometric indicators such as the H-index and i10-index.

By examining patterns of suspicious publication behaviour, mapping citation loops, and engaging with ResearchGate's content submission pipeline, we demonstrated how GenAI tools can be co-opted for metric-based academic fraud. Notably, our study exposes how seemingly legitimate citation activity—when powered by inauthentic documents—can distort scholarly metrics and rankings, ultimately undermining trust in academic evaluation systems.

While limited in scope and conducted ethically, our study highlights wider risks: it exposes gaps in ResearchGate's content validation and reflects a broader trend of GenAI-enabled self-promotion across digital academic platforms. Our findings add empirical weight to growing concerns around GenAI's impact on scholarly publishing.

Future work includes tracking citations of our experimental, fake paper, testing similar uploads across platforms like arXiv and SSRN, and engaging with providers such as ResearchGate and Google Scholar to recommend safeguards —e.g., stronger author verification, metadata checks, and GenAI-content detection. We also aim to collaborate with institutions to define ethical GenAI use in research writing and develop protocols for spotting fabricated content. Finally, we believe that our work contributes to broader efforts to preserve the integrity of digital scholarship and reform how impact is measured.

Disclaimer

This study involved the creation and upload of a fake academic paper solely for experimental and ethical research purposes, and was conducted transparently, with disclaimers and no ill intent.

Acknowledgements

Some portions of this manuscript were refined using GenAI tools (specifically, ChatGPT) to assist with language polishing and structural clarity. Following the use of such GenAI tools, the authors thoroughly reviewed and edited the content as necessary and take full responsibility for the final publication. All core ideas, findings, experiments, and arguments were developed and authored by the named contributors.

References

[1] H. Ibrahim, F. Liu, Y. Zaki, and T. Rahwan, "Citation manipulation through citation mills and pre-print servers," *Scientific Reports*, vol. 15, no. 1, p. 5480, Feb. 2025, https://www.nature.com/articles/s41598-025-88709-7.

[2] M. Májovský, M. Černý, M. Kasal, M. Komarc, and D. Netuka, "Artificial intelligence can generate fraudulent but authentic-looking scientific medical articles: Pandora's box has been opened," *J Med Internet Res*, vol. 25, p. e46924, May 2023. [Online]. Available: https://www.jmir.org/2023/1/e46924

[3] H. Al-Sinani and C. Mitchell, "Unleashing AI in ethical hacking: A preliminary experimental study," Royal Holloway, University of London, Technical Report, 2024, https://pure.royalholloway.ac.uk/files/58692091/TechReport_UnleashingAIinEthicalHacking.pdf.

[4] H. S. Al-Sinani and C. J. Mitchell, "AI-augmented ethical hacking: A practical examination of manual exploitation and privilege escalation in Linux environments," *CoRR*, vol. abs/2411.17539, 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2411.17539

[5] ——, "Introducing PenTest++: An AI-augmented, automated, ethical hacking system," in *Proceedings of the 2nd International Workshop on Cybersecurity: Blockchain and Artificial Intelligence Applications (CyBAI'25), part of the 2025 IEEE 8th Congress on Information Science and Technology (CiSt), 4–10 October 2025*. Marrakech, Morocco: IEEE, p. to appear. [Online]. Available: https://doi.org/10.48550/arXiv.2502.09484

[6] H. S. Al-Sinani, C. J. Mitchell, N. Sahli, and M. Al-Siyabi, "Unleashing AI in ethical hacking," in *Security and Trust Management — 20th International Workshop, STM'24, Bydgoszcz, Poland, September 19–20, 2024, Proceedings*, ser. Lecture Notes in Computer Science, F. Martinelli and R. Rios, Eds., vol. 15235. Springer, 2024, pp. 140–151. [Online]. Available: https://doi.org/10.1007/978-3-031-76371-7_10

[7] H. S. Al-Sinani, N. Sahli, C. J. Mitchell, and M. Al-Siyabi, "Advancing ethical hacking with AI: A Linux-based experimental study," in *Proceedings of the Joint National Conference on Cybersecurity (ITASEC & SERICS 2025), February 03–08, 2025, Bologna, Italy*, G. Costa, R. Montanari, M. Carminati, and G. Sciarretta, Eds., vol. 3962. CEUR-WS, 2025. [Online]. Available: https://ceur-ws.org/Vol-3962/paper7.pdf

[8] A. Ahmed and J. Pomeroy, "IoT security in the quantum era: Leveraging AI for predictive threat intelligence," ResearchGate, Feb. 2025.

[9] M. Kumar and T. Roseth, "Ethical hacking in the age of AI and IoT: Proactive cyber defense strategies," ResearchGate, Feb. 2025.

[10] S. Ahmed and T. Roseth, "Quantum computing and blockchain synergy: A new paradigm for information security," ResearchGate, Feb. 2025.

[11] K. Yadav and T. Roseth, "AI-powered cyber security: Enhancing SOC operations with ML and blockchain," ResearchGate, Feb. 2025.

[12] Y. Jaiswal and T. Roseth, "Machine learning for SOC operations: Automating threat detection and response," ResearchGate, Feb. 2025.

[13] T. Hayat and K. Gatlin, "AI-powered ethical hacking: Rethinking cyber security penetration testing," ResearchGate, Feb. 2025. [Online]. Available: https://www.researchgate.net/publication/390128254_AI-Powered_Ethical_Hacking_Rethinking_Cyber_Security_Penetration_Testing

[14] N. Fahad and K. Gatlin, "Leveraging blockchain for decentralized and secure SOC operations," ResearchGate, Feb. 2025.

[15] F. Sani and R. Spunda, "Quantum computing and cybersecurity: Preparing for a post- quantum world," ResearchGate, Feb. 2025.

[16] K. Rahul and R. Spunda, "Advanced ethical hacking techniques using AI and predictive modeling," ResearchGate, Feb. 2025.

[17] A. Sharma and R. Spunda, "SOC optimization through AI-powered automation and blockchain integration," ResearchGate, Feb. 2025.

[18] M. Shamir and R. Spunda, "The role of blockchain in securing IoT devices and critical infrastructure," ResearchGate, Feb. 2025.

[19] U. Zafer and J. Pomeroy, "Blockchain-powered IoT security: Ensuring data integrity and device trustworthiness," ResearchGate, Feb. 2025.

[20] A. Russell and J. Pomeroy, "Strengthening SOC operations with blockchain for tamper-proof incident management," ResearchGate, 2025.

[21] R. Bishnoi and J. Pomeroy, "AI and quantum computing: Transforming information security protocols for the future," ResearchGate, Feb. 2025.

[22] M. Nasir and J. Pomeroy, "Ethical hacking meets AI: Revolutionizing vulnerability assessments and pentesting," ResearchGate, Feb. 2025.

[23] J. Stribling, M. Krohn, and D. Aguayo, "Scigen – an automatic cs paper generator," 2005. [Online]. Available: https://pdos.csail.mit.edu/archive/scigen/