Unleashing AI in Ethical Hacking

Haitham S. Al-Sinani[1][0009−0005−0453−3335], Chris J. Mitchell[2][0000−0002−6118−0055], Nabil Sahli[3][0000−0002−9805−6859], and Mohamed Al-Siyabi[4]

[1] Department of Cybersecurity and Quality Control, Diwan of Royal Court, Muscat, Oman. hsssinani@diwan.gov.om
[2] Department of Information Security, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, UK. C.Mitchell@rhul.ac.uk
[3] Department of Computer Science, German University of Technology in Oman, Muscat, Oman, nabil.sahli@gutech.edu.om
[4] Military Technological College, Muscat, Oman, Mohamed.Al-Siyabi@mtc.edu.om

Abstract. This paper explores the potential use of generative Artificial Intelligence (GenAI) to enhance the effectiveness and efficiency of ethical hacking, and outlines a proof-of-concept implementation. It briefly reviews the fundamentals of GenAI with a focus on ChatGPT, and then summarises the concept and phases of ethical hacking. The paper also critically assesses risks such as misuse of AI, data biases, and the danger of over-dependence on technology, emphasising the importance of a collaborative human-AI partnership. The paper concludes with a discussion of possible future directions, including use of AI in strengthening cyber defences. This research contributes to the ongoing dialogue around the ethical and innovative application of AI to bolster security.

Keywords: AI · Ethical Hacking · Generative AI · ChatGPT.

1 Introduction

Ethical hacking, a key proactive security measure, traditionally requires a high level of skill and constant updating of knowledge to work effectively. This paper examines the use of ChatGPT, a state-of-the-art generative AI (GenAI) model, to assist ethical hacking. The significance of this approach lies in the potential of GenAI to enhance the capabilities of ethical hackers, allowing for more sophisticated and efficient security assessments and, ultimately, stronger defence against cyber threats. In considering the adoption of GenAI in ethical hacking, this paper not only contributes to current discourse but also aims to spark further research, fostering a more secure digital environment in an age where cyber threats continue to evolve rapidly.

The remainder of the paper is organised as follows. Section 2 explores ChatGPT and GenAI, and section 3 highlights the ethical hacking landscape. Section 4 presents the GenAI-ethical hacking interoperation model, and section 5

introduces a prototype implementation. Section 6 discusses the potential benefits and risks. Section 7 reviews related work, and, finally, section 8 summarises the conclusions and outlines our plans for future work.

2 Generative AI and ChatGPT

The advent of GenAI, with models like ChatGPT[5] [4] being prominent, is a transformative shift in the AI landscape. These systems, moving beyond traditional AI's focus on pattern recognition and decision-making, excel in content creation, including text, images, and code. The ability to learn from extensive datasets and produce outputs that mimic human creativity is a major advance.

Central to this revolution is the GPT (Generative Pre-trained Transformer) architecture, the basis of models like ChatGPT. Developed by OpenAI, GPT models are built on deep learning techniques using *transformer* models, designed specifically for handling sequential data. These models undergo pre-training, where they learn from a wide array of Internet texts, followed by fine-tuning for specific tasks. This process enables models to grasp not just the structure of language but also its context, essential for generating human-like text.

Each iteration of ChatGPT has demonstrated enhanced sophistication, contextual understanding, and relevance of its outputs. Its primary function lies in interpreting user prompts and generating coherent, contextually appropriate responses. This versatility extends from conducting conversations to performing complex tasks, including coding, content creation, and, as we propose in this paper, ethical hacking. The GPT model family, including ChatGPT, owes much of its success to the transformer model, introduced in the 2017 paper: 'Attention is All You Need [11].' This architecture revolutionises sequence processing through attention mechanisms, enabling the model to focus on different parts of the input based on its relevance to the task.

As we explore the intersection of AI and cybersecurity, understanding ChatGPT's foundational aspects becomes vital. Its generative nature, contextual sensitivity, and adaptive learning capacity can lead to innovative approaches in cybersecurity practices. Our focus will be on how these qualities of ChatGPT can be harnessed to revolutionise methodologies in ethical hacking, exploring the technical, ethical, and practical implications of this integration.

3 The Ethical Hacking Landscape

Ethical hacking, or penetration testing [10], involves applying hacking techniques to help organisations strengthen their security posture. It is a discipline that requires a blend of technical expertise, creativity, and a deep understanding of potential threats. Unlike illegal hacking, ethical hacking is the authorised practice of bypassing system security to identify potential vulnerabilities and thereby help protect against real cyber attacks. Ethical hacking typically follows a structured approach involving five phases, as follows.

[5] `https://openai.com/blog/chatgpt`

1. **Reconnaissance:** This involves gathering detailed information about the target system, including public data, network structures, IP address ranges, live hosts, and system configurations. The objective is to understand the target's environment and identify potential vulnerabilities.
2. **Scanning and Enumeration:** In this phase, ethical hackers use various tools to detect open ports, and exploitable system weaknesses, such as vulnerable services. The goal is to map out the target's network and identify entry points for potential attacks.
3. **Gaining Access:** During this stage, hackers leverage identified vulnerabilities to penetrate the target system. The aim is to gain an initial foothold, demonstrating how an attacker could potentially breach security measures and obtain unauthorised access.
4. **Maintaining and Elevating Access:** This stage involves researching and developing ways to gain elevated access and re-enter the system undetected. Techniques such as backdoors are implemented to enable long-term, potentially privileged access. Additionally, tactics like pivoting and lateral movement are employed to navigate through the network, accessing various systems to increase control and maintain persistence within the environment.
5. **Covering Tracks and Reporting:** This stage focuses on erasing any traces of the hacking activity to avoid detection and restoring the target system to its original state. Following this, ethical hackers prepare a comprehensive report detailing their findings and providing recommendations for improving the system's security posture.

While we have chosen to follow the *Penetration Testing Execution Standard* (PTES [2]) for its comprehensive and systematic approach, we acknowledge the existence of other reputable standards in the field. For instance, the *Open Source Security Testing Methodology Manual* (OSSTMM [8]) offers a robust framework for security testing, including physical and human security vectors. The *Information Systems Security Assessment Framework* (ISSAF) [1] provides a wide-ranging protocol for information security assessment, and NIST SP 800-115 [10] gives detailed guidelines for testing information systems, integrating well with risk management frameworks. Each of these standards has its strengths and can be suitable for different testing environments or organisational needs. Additionally, the OWASP (*Open Web Application Security Project*) Top 10 [9] serves as an essential reference for web application security.

4 Generative AI-Ethical Hacking Interoperation

We now propose a conceptual model for using the capabilities of GenAI to support ethical hackers across the five stages of ethical hacking. The model involves employing GenAI's advanced natural language processing and generation skills to enhance and streamline the tasks undertaken by pen-testers in each phase. We next summarise the proposed model's application in each stage of ethical hacking. We are currently undertaking a series of hands-on, research-driven experiments to empirically validate the approach.

4.1 Reconnaissance

The model is built on the assumption that GenAI can significantly enhance the reconnaissance phase of ethical hacking by automating collection and analysis of information. It can be programmed to gather data from multiple open-source intelligence (OSINT) resources, social media platforms, and public databases to create comprehensive profiles of a target system. By processing natural language, GenAI can sift through large volumes of text to identify useful information, such as potential entry points, system configurations, and network infrastructure details. GenAI could also be set to perform domain-specific reconnaissance, such as identifying tech stack details from developer forums or technical documentation, enabling a more targeted approach to vulnerability assessment. GenAI can create believable pretexts for social engineering, e.g. by generating messages that appear legitimate, to gather information without arousing suspicion.

4.2 Scanning and Enumeration

During this phase, GenAI interacts with network scanning tools and software to interpret outputs and suggest next steps. Via integration with existing network scanners and vulnerability assessment tools, it allows ethical hackers to ask questions in natural language about the target system's security posture. Additionally, GenAI's ability to learn from new data could be used to recognise novel patterns of vulnerabilities or calibrate scanning tools more effectively based on the system being assessed. GenAI interprets outputs from tools like Nmap and Wireshark, prioritises vulnerabilities, identifies anomalies, and cross-references findings with CVE databases to guide penetration testing.

4.3 Gaining Access

To help gain access, GenAI could aid ethical hackers by suggesting known exploitation techniques relevant to discovered vulnerabilities. It could simulate an attacker's reasoning by proposing attack vectors and generating proof-of-concept code snippets for exploiting vulnerabilities. GenAI could also provide insights into the latest exploit databases and frameworks, such as Metasploit modules, ensuring that ethical hackers are equipped with cutting-edge knowledge to identify and demonstrate risks effectively.

4.4 Maintaining and Elevating Access

While maintaining access is typically not a focus for ethical hackers beyond demonstration purposes, we argue that GenAI can still contribute by outlining how attackers could establish persistence. GenAI can offer guidance on elevating access as well as establishing persistent system access, all while remaining undetectable. It could generate hypothetical scenarios and suggest methods for creating backdoors or rootkits based on the vulnerabilities identified. This information can be valuable for understanding the risks associated with APTs

(Advanced Persistent Threats) and formulating defences against them. In addition, GenAI can significantly assist ethical hackers by simulating the techniques of pivoting and lateral movement. This capability empowers the ethical hackers to comprehend and execute the essential strategies required to navigate through and expand their influence within a network.

4.5 Covering Tracks and Reporting

GenAI can assist in covering tracks and restoring the target system to its original state, crucial for avoiding detection during ethical hacking. It can identify traces left by attackers and propose effective methods for their removal, thus aiding in testing the robustness of a system's logging and monitoring capabilities. Furthermore, GenAI is likely to be adept at analysing system logs to detect forensic evidence of an attack and recommending strategies to conceal these traces, enhancing our knowledge of anti-forensic techniques. Additionally, we anticipate that GenAI will prove highly effective in documenting the ethical hacking process, summarising key findings, and generating comprehensive reports. These reports would typically include detailed descriptions of the techniques employed, vulnerabilities uncovered, and strategic recommendations for remediation.

5 A Proof-of-concept

To validate the proposed approach, we conducted an experimental study in a controlled virtual environment. We assessed ChatGPT's effectiveness in aiding penetration testing, including testing Windows virtual machines with Kali Linux as the attack machine. A comprehensive account of the experiment is provided in a separate publicly available technical report [3]. However, we will include here the key details for the completeness of this paper.

5.1 Laboratory Setup

Physical Host The experiment utilised a MacBook Pro with 16 GB RAM, a 2.8 GHz Quad-Core Intel Core i7 processor, and 1 TB of storage, providing sufficient computational capabilities for virtualisation.

Virtual Environment Configuration The virtualisation of the network was achieved using VirtualBox 7, a reliable tool for creating and managing virtual machine environments. The setup included the following virtual machines (VMs).

1. **Kali Linux VM:** this machine functioned as the primary attack platform for conducting the penetration tests. It is equipped with the necessary tools and applications for ethical hacking.
2. **Windows VM:** this machine, running a 64-bit version of Windows Vista with a memory allocation of 512 MB, was the principal target for penetration testing within the experiment.

3. **Linux VM:** this machine, operating on a 64-bit Linux Debian system and allocated 512 MB of memory, was reserved for future testing and analysis.

The network configuration was established in a local NAT (Network Address Translation) network setup, allowing for seamless communication between the VMs and simulating a realistic network environment suitable to penetration testing and cybersecurity research.

5.2 Generative AI Tool

The experiment leveraged ChatGPT-4[6] for its advanced AI capabilities and efficient response time. The selection of ChatGPT-4 was primarily based on its prominent status as a leading GenAI tool, offering cutting-edge technology to enhance the ethical hacking process. It's crucial to note, however, that other GenAI tools are also available, e.g. Google's Bard[7] and GitHub's Co-Pilot[8], which could potentially be utilised in similar contexts. The methodologies and processes described are applicable to both the paid and free versions of ChatGPT, with the paid version chosen for improved performance in this study.

5.3 Objective and Methodology

Our objective is to leverage ChatGPT's capabilities to assist in the ethical hacking process, aiming to gain unauthorised access to the target Windows VM. The experiment followed the five structured phases of ethical hacking (see section 4), with ChatGPT's guidance integrated at each step, as follows.

Reconnaissance In this paper, our emphasis is on active reconnaissance (recon), which necessitates engaging with the target to stimulate responses for observation. Therefore, during this phase, we have followed the steps listed below.

- As an integral part of the initial reconnaissance phase, our aim is to identify active machines within the target network in order to select our target. To achieve this, we posed the following question to ChatGPT: "I'm currently in the initial stage of ethical hacking, known as Reconnaissance. Could you please provide a list of the top 4 commands I can use on my Kali machine to find out which devices are currently active on my local network?". As depicted in technical report [3], ChatGPT responded with a useful compilation of potential Kali terminal commands, including **nmap, netdiscover, and arp-scan,** along with examples on their utilisation.
- Next, we have transitioned to our Kali 'attack' machine and applied the recommendations provided by ChatGPT. As a result, we have successfully identified the active devices within the target network.

[6] `https://openai.com/gpt-4`
[7] `https://bard.google.com/`
[8] `https://github.com/features/copilot/`

- To determine the IP address of our Kali 'attack' machine, we employed the 'hostname' command with the '-I' option.
- To make well-informed estimations regarding potential target machines, we can exclude both our Kali IP address and the standard default gateway. We can seek guidance from ChatGPT, wherein it analyses the 'arp-scan' command output, listing active network nodes, and the 'hostname' command output, specifying the Kali IP address. ChatGPT performs this analysis and offers educated insights via a question-answer chat communication.
- As a result of the analysis presented above, we can pinpoint the VMs with the IP addresses 192.168.1.6 and 192.168.1.7 as potential targets. This allows us to proceed to the second scanning stage.

Scanning and Enumeration During this stage, ethical hackers use automated tools to scan the target system or network for vulnerabilities. This can include port scanning, vulnerability scanning, and more. In our specific scenario, the system that demands our scanning attention is the Windows machine identified by the IP address: '192.168.1.6'.

Analogous to preparing for a physical break-in, where determining the house address (reconnaissance) is the initial step, now that we have acquired the address (192.168.1.6), we can proceed to conduct a thorough inspection of the house to determine if any windows (ports) left open that may represent vulnerabilities.

We initiated this phase by consulting ChatGPT for a compilation of key commands suitable for gathering comprehensive information about a specific target (192.168.1.6) using our Kali machine. We explained to ChatGPT that our objective was to acquire extensive knowledge about this particular system to prepare for a forthcoming attack. ChatGPT thus provided us with a brief list of potential scanning commands, prominently featuring 'nmap -A -T4 192.168.1.6'. ChatGPT further clarified that the 'nmap -A -T4 192.168.1.6' command performs an aggressive scan, including OS detection, version detection, script scanning, and traceroute. The '-T4' option speeds up the scan with an aggressive timing template. Nmap, by default, scans the top 1,000 TCP ports; however, to scan all 65,535 ports on the target, we simply add the '-p-' option.

Subsequently, we employed the ChatGPT-recommended key command, 'nmap -A -T4 -p- 192.168.1.6' to perform an exhaustive scan of our target machine. The nmap scan results, unequivocally identifying our Windows target VM, are presented in our technical report [3]. We tasked ChatGPT with the analysis of these results and solicited suggestions for potential unauthorised access routes. ChatGPT highlighted that the nmap scan unveiled potential avenues for attack, most notably the open SMB (Server Message Block) ports 139 and 445, which may harbor vulnerabilities, including the infamous EternalBlue (MS17-010) exploit for remote code execution. We will then consult ChatGPT on exploiting the EternalBlue vulnerability, progressing to the next phase (see below).

Gaining Access In this phase, we seek guidance from ChatGPT to gain access to the Windows VM with the IP address 192.168.1.6 using our Kali at-

tack machine. To simplify the process, we have chosen to exploit the EternalBlue vulnerability via Metasploit. Our request to ChatGPT involves receiving instructions on utilising Metasploit on our Kali machine to execute the EternalBlue (MS17-010) attack after first confirming the system's vulnerability to this exploit. As depicted in our technical report [3], ChatGPT has provided a step-by-step guide. We begin by launching Metasploit with the 'msfconsole' command, and, then, proceed to search for the EternalBlue module using 'search eternalblue'. Next, we select the EternalBlue exploit module with 'use exploit/windows/smb/ms17_010_eternalblue', set the necessary options, including the target host IP address using 'set RHOSTS 192.168.1.6', and, optionally, configure the payload, which is set by default anyway. ChatGPT advises checking the target's vulnerability with the 'check' command, thereby confirming the system's susceptibility. Finally, to execute the exploit, we run the 'exploit' command, resulting in successful system ownership and system (root) access. For a visual representation of this step-by-step process, please refer to the publicly available technical report [3].

Maintaining and Elevating Access In this stage of ethical hacking, our objective is to ensure we can re-enter the system in the future, ideally without being detected. Typically, achieving persistent access requires elevated privileges, often in the form of administrator or root access. As a result, we would typically turn to ChatGPT to assist us in elevating our access level. Fortunately, in the previous stage, we successfully exploited the 'EternalBlue' vulnerability, granting us system access (the highest level of privileged access possible). With this in mind, we consulted ChatGPT for guidance on maintaining persistent access. As shown in our technical report [3], ChatGPT provided a list of recommendations for establishing persistent access. These include creating backdoors, utilising scripts for persistence, manipulating services or scheduled tasks, DLL hijacking, and modifying registry keys. For simplicity, we requested a step-by-step guide from ChatGPT on creating a basic backdoor by adding a new user account with administrative privileges. ChatGPT did indeed offer a detailed guide, which involves creating a new user account using the command: 'execute -f cmd.exe -c -H -i -a "/c net user newusername password /add"', adding this user to the administrators group using the command: 'execute -f cmd.exe -c -H -i -a "/c net localgroup administrators newusername /add"', and verifying the new user's addition, such as through the command: 'execute -f cmd.exe -c -H -i -a "/c net localgroup administrators"'. Following ChatGPT's instructions meticulously, we confirmed the successful addition of the new user to the admin group. Subsequently, we also tested this by restarting the Windows target machine and successfully confirmed our ability to log in using the newly created user through the standard Windows login procedure.

Covering Tracks and Documentation This phase comprises two parts:

- **covering our tracks,** which involves erasing or minimising evidence of our activities within the target system, crucial to avoid detection and maintain the system as close to its original state as possible; and
- **documentation,** which involves creating the pen-test report, a topic discussed later.

In the first part, aiming to remain undetected, we asked ChatGPT for a detailed guide on effectively covering our tracks. As shown in our technical report [3], ChatGPT provided a list of actions to achieve this, including:

- the removal of the newly added user account (Haitham) using the command: 'meterpreter execute -f cmd.exe -c -H -i -a "/c net user Haitham /delete"';
- clearing system logs with the command: 'meterpreter > clearev';
- deleting any files created or downloaded onto the target system;
- uninstalling any software or tools;
- resetting system settings;
- removing scheduled tasks for persistence;
- flushing DNS and ARP cache with 'ipconfig /flushdns' and 'arp -d *' to eliminate network activity traces; and, finally,
- gracefully closing the 'Meterpreter' session using the 'exit' command.

While we have implemented some of these recommendations, such as clearing system logs, it's worth noting ChatGPT's caution that clearing logs can raise suspicion in real-world scenarios and might not always be advisable.

As for the second part, the documentation part, it is crucial for ethical hackers to produce a comprehensive and thorough report for each penetration testing assignment. Therefore, we asked ChatGPT to assist us in composing a detailed report for our penetration test (simulation) assignment using all the information that ChatGPT already knows about from our chat-based communication. As shown in our technical report [3], ChatGPT first responded with a template that we can use to structure our report, along with guidance on what to include in each section. Since providing the template was not satisfactory, we asked ChatGPT again to write a comprehensive and detailed report for this penetration testing assignment using the recommended template and the information we have discussed in our chat, incorporating as much detail as feasible and providing supporting evidence where relevant. This time around, ChatGPT responded with a well-written and accurate penetration test report, including writing the 'Executive Summary, 'Introduction', 'Methodology', 'Findings and Results', 'Attack Narrative', 'Conclusions and Recommendations', as well as suggestions for 'Appendices' . In subsequent questions to ChatGPT, we further tweaked and improved the ChatGPT-produced report, including specifying the target organisation, time period, and the date.

6 Discussion: Benefits and Potential Risks

We have introduced a methodical approach to using the capabilities of GenAI to augment each stage of ethical hacking, from initial data gathering to final reporting. The model aims to support provision of a thorough and efficient evaluation

of security vulnerabilities. It assists in the detailed and rapid identification of potential threats, contributes to the development of social engineering scenarios, and aids in formulating defensive measures through simulated attacks. Implementing this model could lead to more streamlined security assessments and potentially improve the training of cybersecurity professionals, thereby equipping them with enhanced skills to identify and address evolving cyber threats.

However, adoption of this model is not without challenges, and requires careful consideration of the accompanying risks and ethical implications. The potential for misuse by adversaries, the risk of over-dependence on automated systems at the expense of human expertise, and the AI's current limitations in processing ambiguous or deceptive information present substantial concerns. Additionally, biases inherent in AI algorithms and the lag in regulatory frameworks to address such rapidly advancing technologies pose further complications. These factors highlight the need for a balanced integration of AI systems, such as ChatGPT, into ethical hacking practices, ensuring that it complements human judgment rather than replacing it, and is governed by a robust ethical framework to maximise its potential while safeguarding against misuse.

7 Related Work

The intersection of AI and cybersecurity is a vibrant area of research, with studies ranging from AI's role in detecting intrusions to its use in aiding offensive security including ethical hacking. Foundational work by Handa et al. [6] has showcased the value of machine learning in network intrusion detection. The rise of sophisticated language models like GPT-3, introduced by Brown et al. [4], has heralded new research possibilities, such as exploring AI's use in crafting realistic phishing attacks, a topic explored by Zannettou [12]. Contemporary studies, for example Gupta et al. [5], look at the dual role of AI, showing how it could both be employed for cyberattacks and harnessed for cyber defence and ethical guidance. Furthermore, a recent practical study by Harrison et al. [7] shows how advances in AI's deep learning algorithms enhance acoustic side-channel attacks against keyboards, achieving groundbreaking keystroke classification accuracies with common devices and apps like smartphones and Zoom.

This paper seeks to expand on these discussions, offering an exploration of GenAI's role across all stages of ethical hacking — a topic that remains underexplored in the existing literature, necessitating a deeper investigation.

8 Conclusions and Future Directions

This paper has discussed the potential role of GenAI as a tool for ethical hacking, and introduced a framework for its application. We have highlighted the collaborative capabilities of human expertise paired with AI's computational power for cybersecurity. The paper also outlined a proof-of-concept implementation; the initial results of which clearly show that ChatGPT, a state of the art GenAI model, is an effective and impactful tool in the field of ethical hacking. The paper

sets the agenda for future empirical research to further validate the assertions underlying our GenAI-ethical hacking interoperation model.

Ongoing research includes a series of hands-on, research-driven experiments aimed at both substantiating the proposed model and refining it to encompass a wider array of hacking domains. This entails increasing the model's scope across other security disciplines including wireless security, privilege escalation, protection against the OWASP Top 10 (web[9] and mobile[10]) vulnerabilities, and mobile app security. By conducting these experiments, we aim to continuously evolve the model to address the ever-changing landscape of cyber threats, ensuring its effectiveness against increasingly sophisticated future attack vectors.

References

1. Information systems security assessment framework (issaf) (2006), https://www.oissg.org/issaf/
2. Penetration testing execution standard technical guidelines (2014), http://www.pentest-standard.org/
3. Al-Sinani, H., Mitchell, C.: Unleashing AI in ethical hacking: A preliminary experimental study. Technical report, Royal Holloway, University of London (2024), https://pure.royalholloway.ac.uk/files/58692091/TechReport_UnleashingAIinEthicalHacking.pdf
4. Brown, T.B., et al.: Language models are few-shot learners. In: Advances in Neural Information Processing Systems. vol. 33, pp. 1877–1901 (2020)
5. Gupta, M., et al.: From chatgpt to threatgpt: Impact of generative AI in cybersecurity and privacy. IEEE Access (2023)
6. Handa, A., Sharma, A., Shukla, S.K.: Machine learning in cybersecurity: A review. WIREs Data Mining and Knowledge Discovery **9**(4), e1306 (2019)
7. Harrison, J., Toreini, E., Mehrnezhad, M.: A practical deep learning-based acoustic side channel attack on keyboards. In: IEEE European Symposium on Security and Privacy, EuroS&P 2023 - Workshops, Delft, Netherlands, July 3-7, 2023. pp. 270–280. IEEE (2023). https://doi.org/10.1109/EUROSPW59978.2023.00034, https://doi.org/10.1109/EuroSPW59978.2023.00034
8. Institute for Security and Open Methodologies (ISECOM): Open Source Security Testing Methodology Manual (OSSTMM) (2020), https://www.isecom.org/OSSTMM.3.pdf
9. OWASP: Owasp top ten (2021), https://owasp.org/www-project-top-ten/
10. Swanson, M., Bartol, N., Sabato, J., Hash, J., Graffo, L.: Technical guide to information security testing and assessment (NIST SP 800-115). Special Publication 800-115, National Institute of Standards and Technology (2008), https://csrc.nist.gov/publications/detail/sp/800-115/final
11. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
12. Zannettou, S.: "What do you think about...?" — generating cybersecurity question-answer pairs using a contextual question generation model. Cybersecurity **4**(1), 10 (2021)

[9] https://owasp.org/www-project-top-ten/

[10] https://owasp.org/www-project-mobile-top-10/