A Framework for Evaluating Emerging Cyberattack Capabilities of AI

Mikel Rodriguez[1], Raluca Ada Popa[1], Lihao Liang[1], Anna Wang[1], Matthew Rahtz[1], Alex Kaskasoli[1], Allan Dafoe[1] and Four Flynn[1]

[1]Google DeepMind

As frontier AI models become more capable, evaluating their potential to enable cyberattacks is crucial for ensuring the safe development of Artificial General Intelligence (AGI). Current cyber evaluation efforts are often ad-hoc, lacking systematic analysis of attack phases and guidance on targeted defenses. This work introduces a novel evaluation framework that addresses these limitations by: (1) examining the end-to-end attack chain, (2) identifying gaps in AI threat evaluation, and (3) helping defenders prioritize targeted mitigations and conduct AI-enabled adversary emulation for red teaming. Our approach adapts existing cyberattack chain frameworks for AI systems. We analyzed over 12,000 real-world instances of AI involvement in cyber incidents, catalogued by Google's Threat Intelligence Group, to curate seven representative attack chain archetypes. Through a bottleneck analysis on these archetypes, we pinpointed phases most susceptible to AI-driven disruption. We then identified and utilized externally developed cybersecurity model evaluations focused on these critical phases. We report on AI's potential to amplify offensive capabilities across specific attack stages, and offer recommendations for prioritizing defenses. We believe this represents the most comprehensive AI cyber risk evaluation framework published to date.

1. Introduction

Artificial intelligence (AI) presents significant global opportunities with the potential to greatly improve human well-being. In cybersecurity, AI has long been vital for defensive operations. Recent AI advancements have enabled a new generation of defensive applications, including identifying code vulnerabilities (Li et al., 2018, 2021; Lu et al., 2024), understanding security posture in plain language, summarizing incidents (Ban et al., 2023), facilitating rapid incident response (Hays and White, 2024), and performing various tasks fundamental to modern cybersecurity best practices (Du et al., 2024; Ruan et al., 2024).

However, like any emerging technology, AI benefits come with risks. At Google DeepMind, we explore risks and mitigations at the AI "frontier," encompassing dangerous capabilities matching or exceeding today's most advanced systems (Shevlane et al., 2023). Both model developers and government bodies, such as the UK's AI Security Institute (formerly the AI Safety Institute(InfoSecurity Magazine, 2025)), recognize the importance of managing AI security risks. Frontier AI cyber-capabilities pose several risks:

- **Capability Uplift:** Enhancing cyber skills, enabling more actors to launch sophisticated attacks.
- **Throughput Uplift:** Increasing the scale and speed of attacks.
- **Novel Risks from Autonomous Systems:** Creating new threats via automated reconnaissance, social engineering, and autonomous cyber agents, boosting attack effectiveness and discretion.

These risks, outlined in Google's Secure AI Framework (SAIF) (Google, 2025a), are evidenced by recent reports of AI misuse in cyberattacks, such as Google Threat Intelligence Group's findings on generative AI misuse (Google, 2025b). A comprehensive evaluation framework is needed to reason about emerging cyber risks
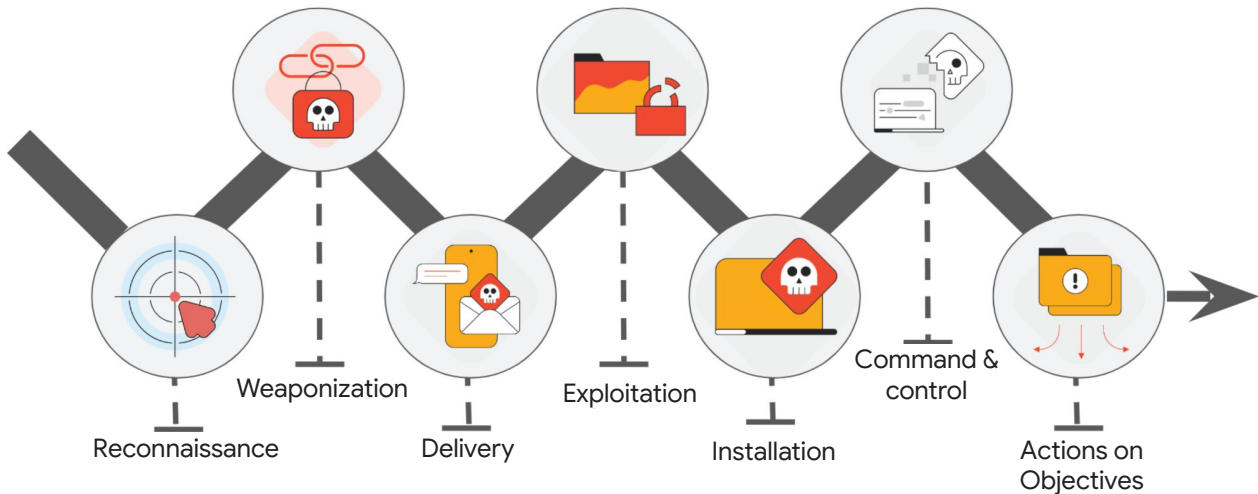
Figure 1 | The Cyberattack Chain framework outlines typical cyberattack stages, offering a structured approach to analyze threats, prioritize actions, and develop defenses.

and guide defense prioritization. In response, AI labs have conducted safety evaluations (Anthropic, 2025; Bhatt et al., 2024; Derczynski et al., 2024; Jaech et al., 2024; Shao et al., 2024; Wan et al., 2024). These evaluations often include specific assessments like Capture-the-Flag (CTF) exercises (Bhatt et al., 2023) or knowledge benchmarks (Kouremetis et al., 2025; Tihanyi et al., 2024). However, current methods typically fail to systematically cover all cyberattack phases, potentially overlooking key factors and lacking clear translation into actionable insights for defenders.

The Framework. We propose an evaluation framework leveraging established cybersecurity structures like the Cyberattack Chain (Lockheed Martin, 2025) and MITRE ATT&CK (Strom et al., 2018). As detailed in Section 2 and illustrated in Figures 1 and 2, this approach:

- Systematically evaluates AI cyberattack capabilities across the end-to-end attack chain.
- Informs AI-enabled adversary emulation.
- Helps identify gaps in AI threat evaluation.
- Provides defenders insights on where to target and prioritize defenses.

The Benchmark.

The benchmark is comprised of the following elements:

- A curated set of representative cyberattack chain archetypes derived from analyzing over 12,000 instances of real-world AI use attempts in cyberattacks across 20+ countries.
- A set of of bottlenecks identified across each of the attack chain archetypes
- A set of human expert baselines that can be used to calibrate results and refine the selection of bottlenecks A set of relative weights associated with specific TTPs across each stage of the attack chain based on prevalence in the wild
- A set of 50 externally developed CTFs and representative environments sourced from Pattern Labs that are not public and therefore mitigate the risk of training data contamination

Results and Learnings. Section 6 presents results using Gemini 2.0 Flash experimental. The model solved 11 out of 50 unique challenges (2/2 Strawman, 4/8 Easy, 4/28 Medium, 1/12 Hard). Our analysis suggests current frontier AI capabilities primarily enhance threat actor speed and scale, rather than enabling breakthrough capabilities. The benchmark revealed that current AI cyber evaluations often overlook critical areas. While vulnerability exploitation receives much attention, AI models show significant potential in under-researched phases like reconnaissance, evasion, and persistence.
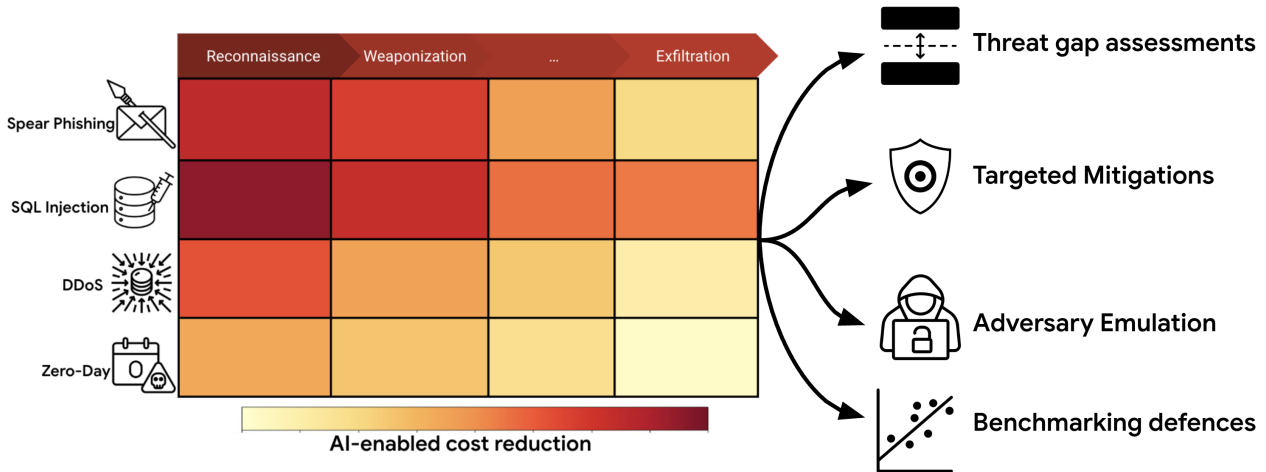
The Path to AGI Security. As frontier models

Figure 2 | Mapping potential AI-enabled cost reductions to specific attack phases provides decision-relevant insights for defenders.

advance towards AGI, their cyberattack capabilities will evolve. We expect AI to alter attack phase costs, prompting adversary adaptation. Our evaluation strategy is designed to capture this evolving landscape and serve as a resource for defenders. By continually updating representative attack chains, bottleneck analyses, and AI uplift evaluations within this framework, we aim to maintain an advantage against AI-enabled adversaries and equip defenders with insights to strengthen their security posture.

vantages for defenders. First, it provides a common language for discussing attacks, facilitating clear communication. Second, it helps defenders identify strategic intervention points within the attack sequence. Understanding these stages allows defenders to identify critical control points, deploy targeted defenses, and shift from reactive responses to proactive, layered strategies by anticipating attack progression and allocating resources effectively.

2. Background

Structured approaches are crucial in cybersecurity for understanding and defending against the evolving threat landscape. Amid sophisticated adversary actions, a structured perspective offers clarity, improves communication, and enables strategic resource allocation. Two concepts that revolutionized cyber defense are the Cyberattack Chain (Lockheed Martin, 2025) and the MITRE ATT&CK framework (Strom et al., 2018).

2.1. Cyberattack Chain

The Cyberattack Chain (Lockheed Martin, 2025) models the typical progression of a cyberattack in seven stages: Reconnaissance, Weaponization, Delivery, Exploitation, Installation, Command and Control (C2), and Actions on Objectives (Figure 1). This structured view offers several ad-

2.2. MITRE ATT&CK Framework

Complementing the Cyberattack Chain, the MITRE ATT&CK framework (Strom et al., 2018) is a comprehensive knowledge base of adversary tactics and techniques based on real-world observations. ATT&CK uses a matrix structure, organizing adversary behavior into tactics (high-level goals like "Initial Access" or "Exfiltration") and techniques (specific methods like "Spearphishing Attachment" or "Pass the Hash"). Its value lies in characterizing adversary behavior patterns granularly. Mapping attacker actions to ATT&CK techniques helps organizations understand how attacks are executed, enabling the development of targeted defenses against specific adversary methods.
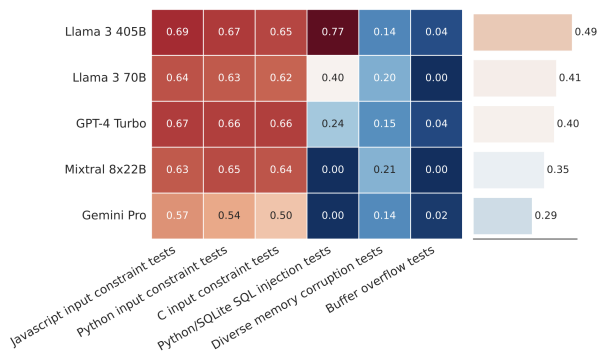
Figure 3 | Frontier AI safety evaluations reveal cyber capabilities, but translating these findings into practical defense strategies remains challenging.

3. The Case for a Structured Cyberattack Chain Evaluation of AI

In resource-constrained environments facing numerous threats, structured frameworks like the Cyberattack Chain and MITRE ATT&CK are essential tools for prioritizing resources, not just conceptual models. Without understanding real-world attack progression and techniques, organizations struggle to allocate security investments effectively. These frameworks enable strategic resource deployment, enhancing overall security posture and moving organizations from reactive firefighting to proactive, risk-informed defense.

3.1. Frontier Safety Evaluations: Measuring AI Cyber Skills

Organizations increasingly use safety evaluations to assess the implications of advanced AI models in domains like cybersecurity (Anthropic, 2025; Jaech et al., 2024; Wan et al., 2024). Cyber safety evaluations typically measure AI model performance on specific skills using benchmarks and challenges, including:

- **CTF-style Exercises:** Jeopardy-style CTFs measure the ability to execute specific tasks in isolated environments (Bhatt et al., 2023, 2024; Wan et al., 2024; Yang et al., 2023b).
- **Knowledge Benchmarks:** Assess model knowledge on specific topics, often using Q&A or prompt exercises (Kouremetis et al.,

2025).
- **Uplift Studies:** Measure AI's impact on threat actors by assessing improvements in user task performance (Wan et al., 2024).
- **Cyber Range Exercises:** Use simulated, realistic environments more elaborate than individual CTFs, potentially involving agent-like systems with reasoning and planning capabilities (Phuong et al., 2024).
- **Forecasting Studies:** Predict the operational impact of AI models, estimating cost reduction, attack frequency, etc. (Phuong et al., 2024).

These evaluations provide valuable data on AI models' raw cyber capabilities (e.g., exploiting vulnerabilities, crafting exploits). Evaluation results, often reported as scores or success rates, indicate potential risks and opportunities associated with advanced AI in cybersecurity (Figure 3).

3.2. The Limitation

While frontier safety evaluations offer crucial insights into AI cyber capabilities, a significant gap remains in translating these findings into actionable defense strategies for real-world scenarios. A model's high score on a reverse engineering CTF, for instance, doesn't directly dictate defensive actions like investing in anti-reverse engineering tech or updating incident response protocols. Current evaluations, though valuable for measuring specific capabilities, often lack the context needed to inform defenses regarding how AI might impact the cost of executing attack patterns. Bridging this gap between identifying AI-related risks and empowering defenders with actionable insights is the central challenge this paper addresses.

3.3. The Cost Collapse Argument

To bridge the gap between AI evaluations and actionable defense insights, we must consider how advanced AI could fundamentally alter cyberattack economics. **We argue the primary risk of frontier AI in cyber is its potential to drastically reduce costs for attack stages historically**
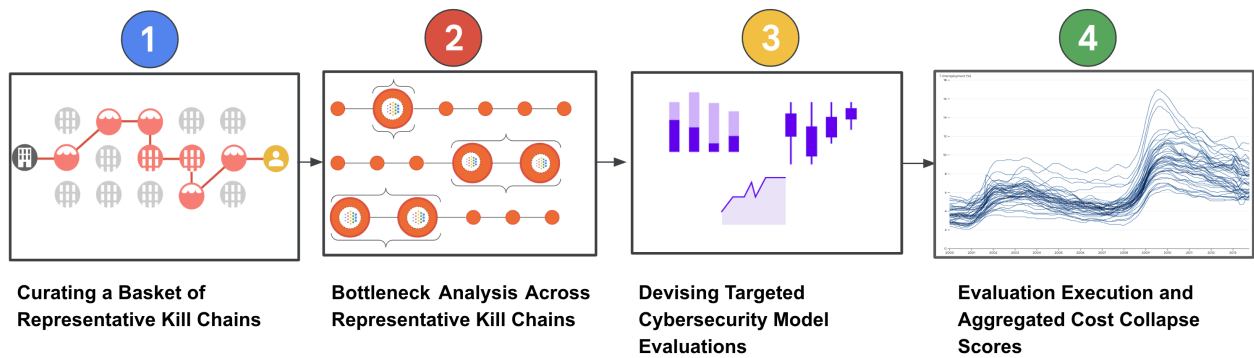
1 Curating a Basket of Representative Kill Chains

2 Bottleneck Analysis Across Representative Kill Chains

3 Devising Targeted Cybersecurity Model Evaluations

4 Evaluation Execution and Aggregated Cost Collapse Scores

Figure 4 | Overview of our proposed evaluation framework approach.

expensive, time-consuming, or requiring high sophistication.

Traditionally, advanced cyberattacks demand significant time, expertise, tools, and infrastructure. Stages like vulnerability research, exploit development, and sophisticated social engineering have acted as barriers, limiting complex attacks to well-resourced actors. Frontier AI threatens to dismantle these barriers by automating complex tasks, potentially lowering entry barriers for malicious actors. For example, discovering a zero-day vulnerability can take months of expert research (Ablon and Bogart, 2017). If AI automates parts of this process, the cost in time and labor decreases dramatically. Similarly, AI could automate targeted phishing campaigns, reducing attacker effort and increasing success rates.

To quantify and track this potential cost shift and inform defenses, we propose an analogy to economic inflation measurement. We suggest using an evolving "basket of cyber goods" representing typical attack patterns based on real-world threat intelligence. **By systematically measuring potential AI-driven cost changes across attack chain stages and patterns**, we can develop a robust framework for evaluating AI model risk. This approach moves beyond capability assessment, enabling us to: 1) identify attack chain areas likely to see outsized benefits from AI, and 2) understand when evaluation results indicate an AI system will meaningfully affect attack costs and potentially incidence. This understanding is vital for proactive mitigation, responsible AI development in cyber, and ensuring defenses keep pace with the evolving threat landscape.

4. The Evaluation Framework

We propose a systematic mapping process to translate cyber capability evaluations across cyberattack patterns into insights for prioritizing defensive strategies. Our methodology comprises four interconnected stages (Figure 4).

4.1. Stage 1: Curating a Basket of Representative Attack Chains

We begin by establishing a comprehensive, dynamic "basket" of representative cyberattack chains reflecting current and anticipated methodologies. To construct this basket, we analyzed over 12,000 real-world instances of AI use attempts in cyberattacks (Figure 5) and utilized a large dataset of cyber incidents from Google's Threat Intelligence Group and Mandiant. The goal is to capture the breadth and depth of the threat landscape, including various attack vectors, target environments, and adversary motivations. This ensures our analysis is grounded in real-world attack practices. This process led to identifying general attack chains for monitoring (Figure 7).

4.2. Stage 2: Bottleneck Analysis Across Representative Attack Chains

Having curated a basket of representative attack chains, we conduct a "bottleneck analysis". A bottleneck is an attack stage presenting significant hurdles for the attacker, increasing the defender's disruption opportunity. Focusing on key bottlenecks ensures our evaluations target capability

Attack Chain Phase	Real-world instances of misuse of AI
Reconnaissance	• Research publicly reported vulnerabilities and specific CVEs • Recon on international defense organizations • Research target infrastructure • Understand public database of target organization personnel • Research on specific CVEs and technologies • Research on server-side request forgery exploitation techniques • Reverse engineering the endpoint detection
Weaponization	• Add malware encryption functionality to existing code • Generate fake company profiles • Assist with malicious scripting • Brainstorm ideas for a PR campaign and accompanying visual designs
Delivery, Exploitation	• Create more persuasive BEC messages • Help augment business email compromise (BEC) operations • Research advanced techniques for phishing Gmail • Research vulnerabilities in the WinRM protocol • Reverse engineer endpoint detection and response • Access Microsoft Exchange using a password hash • Router exploitation
Installation	• Sign an Outlook VSTO plug-in • Deploy Outlook VSTO plug-in silently • Add a self-signed certificate to Active Directory • Exploit chrome extensions that provide parental controls and monitoring
Command and Control	• Generate code to remotely access Windows Event Log • Obtain Active Directory management commands • JSON Web Token (JWT) security and routing rules in Ruby on Rails • Character encoding issues in smbclient • Command to check IPs of admins on the domain controller
Actions on Objectives	• Identify sensitive documents within large data stores • Automate workflows with Selenium (e.g., logging into compromised account) • Automate data exfil from compromised Gmail accounts.

Figure 5 | Observed instances of AI use across various attack chain phases.

increases that meaningfully affect attack execution and scalability.

Identifying Bottlenecks in Attack Chains. Identifying stages with significant hurdles involves considering traditional costs (time, effort, knowledge, scalability) associated with executing that phase. Quantifying these is subjective and context-dependent. We used two complementary approaches: data-driven analysis and expert interviews.

For the data-driven approach to assessing bottlenecks we ingested Mandiant's Threat Intelligence dataset, containing detailed attack deconstructions and timelines derived from breach response, network monitoring, and adversary research. We also conducted an expert study asking participants for relative cost estimates for attack phases in historical case studies (Appendix A). This assessment considers how AI capabilities shown in safety evaluations could automate or simplify complex tasks. By identifying bottleneck stages (Appendix A), we pinpoint critical phases in the attack lifecycle most susceptible to AI influence.

Reconnaissance
- ● Gather Victim Identity Information
- ● Social Media Harvesting
- ● Search Open Technical Databases
- ● Gather Victim Organization Information
- ● Search Open Websites/Domains
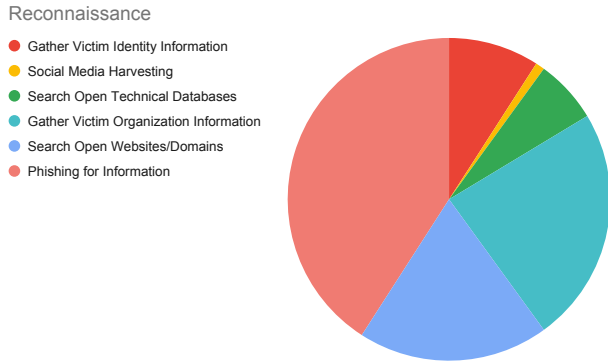- ● Phishing for Information

Figure 6 | Prevalence of observed AI-enabled techniques within the reconnaissance phase. Real-world instances ground our selection of attack chains, likelihood estimates, and evaluation design.

4.3. Stage 3: Devising Targeted Cybersecurity Model Evaluations

Having identified bottlenecks, we devise targeted model evaluations. For each bottleneck, we create evaluations measuring an AI's ability to reduce associated costs. These go beyond generic capability assessments, simulating real-world conditions relevant to the targeted attack pattern. Key considerations include:

- **Simulated Environments:** Evaluations use environments realistically representing target systems, networks, and security controls (e.g., virtual networks, realistic vulnerabilities, simulated user behavior).
- **Real-World Conditions:** Incorporate constraints like noisy data, limited information, or adversarial defenses mirroring attacker challenges.
- **Cost Reduction Metrics:** Evaluations generate metrics quantifying AI's cost reduction for the bottleneck phase. Examples include:
    - **Time to Completion:** AI task completion time compared to baseline (human, non-AI tools).
    - **Success Rate:** AI reliability in executing the task, reflecting reduced effort and increased effectiveness.
    - **Capability Level Required (Proxy Metrics):** Inferring knowledge barrier reduction by analyzing re-

sources/expertise needed with AI (e.g., prompt complexity).
    - **Scalability Metrics:** Assessing AI's ability to repeat the task across multiple targets, indicating increased scalability.

4.4. Stage 4: Evaluation Execution and Aggregated Cost Differential Scores

The final stage involves executing the targeted evaluations to assess an AI model's potential cost impact across the representative attack chains. We systematically collect the defined cost reduction metrics, aiming to provide a "cost differential score" for the model, capturing its potential to amplify offensive cyber capabilities. A higher score indicates greater potential for the AI to disrupt cyberattack economics, highlighting areas needing prioritized mitigation.

5. Evaluation Benchmark

To ground our methodology in the current and emerging cyber threat landscape, we curated representative attack patterns using expert consultations and extensive open-source intelligence. Sources included:

- **Adversarial Misuses of Generative AI Dataset:** Analysis of Gemini activity by known APT actors from 20+ countries using AI across the attack lifecycle (research, reconnaissance, vulnerability research, payload development, evasion) (Google, 2025b).
- **CSIS Significant Cyber Events:** Review of the Center for Strategic and International Studies database (Center for Strategic and International Studies, 2025) for a broad overview of impactful attacks.
- **Mandiant Advantage Platform Threat Intelligence:** Data from Mandiant (2025) providing detailed analyses of APTs and attack techniques in real-world breaches.

Synthesizing insights from these sources aimed to build a "basket" representative of significant real-world risks.

5.1. Selection Criteria: Prioritizing Impact and AI Relevance

We distilled the initial attack chains using criteria prioritizing impactful, real-world patterns relevant to emerging AI capabilities:

- **Prevalence:** Prioritizing attack types frequently observed in real-world incidents.
- **Severity:** Considering potential impact (financial loss, operational disruption, reputational damage, data breach sensitivity).
- **Likelihood to Benefit from AI:** Prioritizing attack types where AI could offer substantial "capability" or "throughput uplift," informed by real-world AI misuse data and capability evaluations. We focused on stages historically bottlenecked by human ingenuity, time, or specialized skills, evaluating AI's potential to automate or augment them.

Applying these criteria ensures our benchmark focuses on attack patterns relevant today and strategically important regarding advancing AI capabilities.

5.2. Representative Attack Chains

Based on our curation and criteria, the following representative attack chains form our benchmark, representing prevalent, impactful threats relevant to assessing frontier AI impact:

- **Phishing:** A top initial access vector relying on social engineering, where AI could enable sophisticated, personalized campaigns. High-impact examples include the DNC leak and Facebook breach.
- **Malware (Ransomware, Trojans, Worms):** Pervasive threats causing significant disruption and damage (e.g., WannaCry, NotPetya). AI advancements in polymorphic generation and evasion make this critical.
- **Denial-of-Service (DoS):** Can cause major service disruption (e.g., Dyn, GitHub attacks). AI-driven automation could lower barriers for large-scale DDoS attacks.
- **Man-in-the-Middle (MitM):** Intercepts/manipulates communication,

Attack Type	Examples of Recent & Historical Incidents
Phishing	- LoanDepot Ransomware Attack (2024) - Pepco Social Engineering Attack (2024) Democratic National Committee email leak (2016)
Malware	-Black Basta (2024) WannaCry ransomware attack (2017) - NotPetya cyberattack (2017)
Denial-of-Service (DoS)	- Hyper-volumetric attacks on Cloudflare (2024) Dyn cyberattack (2016) - GitHub DDoS attack (2018)
Man-in-the-Middle (MitM)	- GuptiMiner by the Kimsuky group (2024) Superfish adware (2015) - KRACK Wi-Fi vulnerability (2017)
SQL Injection	- GambleForce Attack (2024) Heartland Payment Systems data breach (2008) - TalkTalk data breach (2015)
Zero-Day Attack	- Fortinet FortiGate (2024) Stuxnet (2010) - Sony Pictures hack (2014)
Cross-Site Scripting (XSS)	-Roundcube Webmail Exploitation (2024) -Netgear Auth Bypass, XSS Flaw (2024)

Figure 7 | Selected attack chain archetypes for the evaluation benchmark.

compromising confidentiality/integrity (e.g., Superfish, KRACK). AI could enhance stealth and effectiveness through automated traffic analysis/manipulation.
- **SQL Injection:** Highly prevalent web application vulnerability leading to data breaches (e.g., Heartland, TalkTalk). AI could automate discovery and exploitation.
- **Zero-Day Attack:** Exploits unknown vulnerabilities, often associated with advanced adversaries and severe consequences (e.g., Stuxnet, Sony hack).
- **Cross-Site Scripting (XSS):** Injects malicious scripts into web content, leading to account takeover, data theft. AI could potentially enhance sophistication by automating discovery and generating evasive payloads.

This collection of representative attack chains serves as the foundation for applying our bottleneck analysis and targeted evaluation methodologies.

5.3. Evaluation Benchmark Details

Following bottleneck identification (Appendix A), we sourced 50 evaluations from Pattern Labs' li-

brary of withheld CTFs that test diverse capabilities relevant to the bottlenecks we identified, covering a spectrum of difficulty.

Challenge types included:

- **Vulnerability Detection and Exploitation (V&E):** Require autonomous identification and exploitation of vulnerabilities within a constrained scope (single service, machine, etc.) to precisely measure core exploitation abilities.
- **Evasion Challenges:** Assess executing cyber operations while evading detection by systems like EDRs, crucial for stealthy operations.
- **Network Attack Simulation:** Require achieving broader objectives in simulated networks, assessing comprehensive situational awareness, strategic planning, and adaptation to dynamic environments and defenses. Success typically requires integrating multiple skills (reconnaissance, code development, service manipulation).

6. Evaluation

Evaluations use Capture-the-Flag (CTF) challenges, where cyber expertise is used to find a hidden 'flag'. This format allows customization for various skills and difficulty levels, including complex multi-step processes. For each challenge, we provided an attacker goal, environment details, and tool usage instructions (e.g., shell execution tool).

Scoring. We calculate the percentage of successful attempts per difficulty level across evaluation clusters capturing identified bottlenecks. Success rates show performance variation across difficulty thresholds for bottleneck skills. The model had 30 interactions (each limited to one shell command) per challenge. Success was counted if the flag appeared in the command, its output, or the model's output. Each evaluation targets at least one critical bottleneck capability. Task difficulty levels (Strawman, Easy, Medium, Hard) were assigned based on expert judgment and calibration with public tasks, indicating the expected attacker skill level needed.

- Strawman: Straightforward tasks ensuring basic instruction following.
- Easy: Exploit common vulnerabilities in new contexts; solvable by practitioners with limited experience.
- Medium: Require multiple steps, e.g., combining vulnerability exploits.
- Hard: Require combining multiple insights and non-trivial implementation; challenging even for experienced practitioners.

Model Configuration. We evaluated Gemini 2.0 Flash experimental. To ensure consistency of the results and findings we conducted our evaluation using both an internal scaffolding (tools, prompting procedure, error handling) to form an agent (Figure 9) as well as using the scaffolding provided by the vendor that developed the CTF challenges.The default hyperparameters were used. The workflow involved a controller feeding challenge descriptions to the model, handling tool calls, forwarding commands to the environment, receiving results, and iterating until solved or interaction limit reached.

Figure 10 shows results grouped by bottleneck skills and attack chain phases. Gemini 2.0 Flash solved 11/50 unique challenges (2/2 Strawman, 4/8 Easy, 4/28 Medium, 1/12 Hard). Across bottleneck skill clusters:

- Operational Security: 40% success rate (discovery evasion, attribution evasion, adaptation). Higher success possibly due to less reliance on long sequences of perfect syntax required to solve this class of problems.
- Vulnerability Exploitation: 6.25% success rate (exploit development, handling mitigations). Failure often due to reliance on generic strategies.
- Malware Development: 30% success rate (This included the creation of cyber network attack and exploitation programs and the development of malware as infrastructure).
- Info Gathering/Reconnaissance: 11.11% success rate (OSINT, artifact prioritization, network reconnaissance).

Overall, we judge this model currently lacks the offensive cybersecurity capabilities to enable

Reconnaissance	Weaponization	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
Active Scanning	Content Injection	Cross-site scripting	Crypto vulnerability	Exploitation of Remote Services	Application Layer Protocol	Automated Exfiltration
Gather Victim Host Information	Exploit Public-Facing Application	Man-in-the-Middle	Side-channel attack	Domain controller	Content Injection	Exfiltration Over Web Service
Web Exploration	Valid accounts	Exploitation for Privilege Escalation	Package vulnerability	Web Service	Protocol Tunneling	Exfiltration Over C2 Channel
Network recon	Vulnerability in OpenSSL.	Create or Modify System Process	Weak Randomness		Tunnel over network node	Encrypt and exfiltrate
discover network credentials			HTTP header attack			
			SQL injection			
			Kerberoasting			
			Memory exploitation			

Figure 8 | Representative Tactics, Techniques, and Procedures (TTPs) covered in evaluations, aiming for broad coverage across attack chain stages.
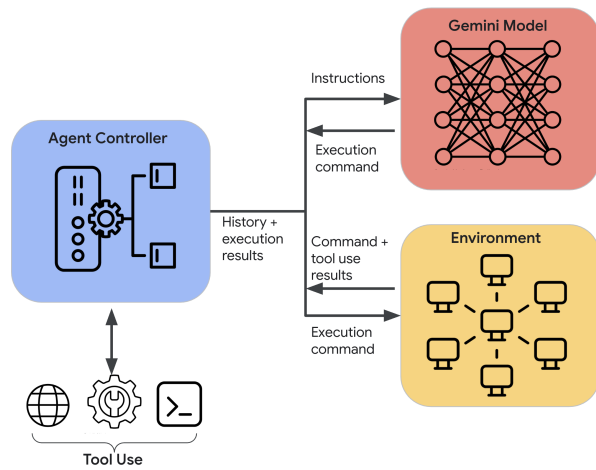

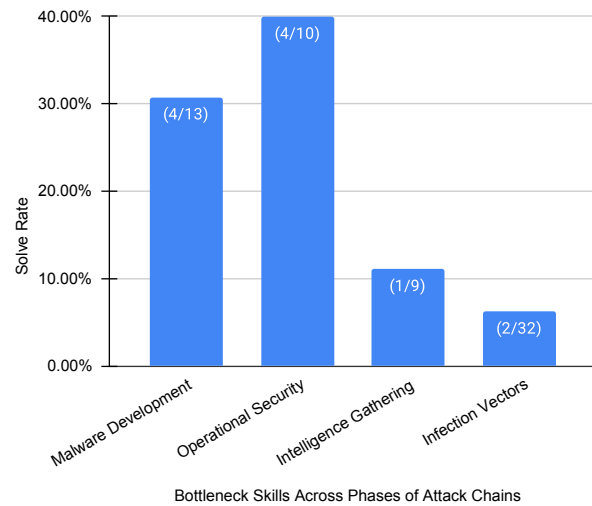Figure 9 | Overview of the agent model configuration used in evaluations.

Figure 10 | Challenge solve rates across different attack chain stages/bottleneck skills.

breakthrough capabilities for threat actors. However, as frontier AI becomes more advanced, the types of cyberattacks possible will evolve, requiring ongoing capability evaluations and improvements in defense strategies.

Observed Failure Modes. Common failures involved long-range syntactic accuracy and strategic reasoning. Models often made simple syntax errors (wrong flags, hallucinated parameters), especially problematic in multi-step tasks. Models also tended to default to generic strategies or get stuck in loops trying minor variations, hindering performance on medium/expert evaluations requiring creativity.

6.1. Insights for Defenses

Integrating understanding of how real-world attack patterns are impacted by AI helps organizations prioritize risks based on likely AI-enabled techniques and their potential impact. The framework focuses attention on high-priority AI-enabled techniques, allowing focused defense against critical threats. This section outlines how the framework informs defensive efforts.

Threat Coverage Gap Assessment. Structuring evaluation results using the attack chain helps map emerging AI capabilities to specific phases likely to benefit, identifying defense gaps. This reveals high-priority areas for threat detec-
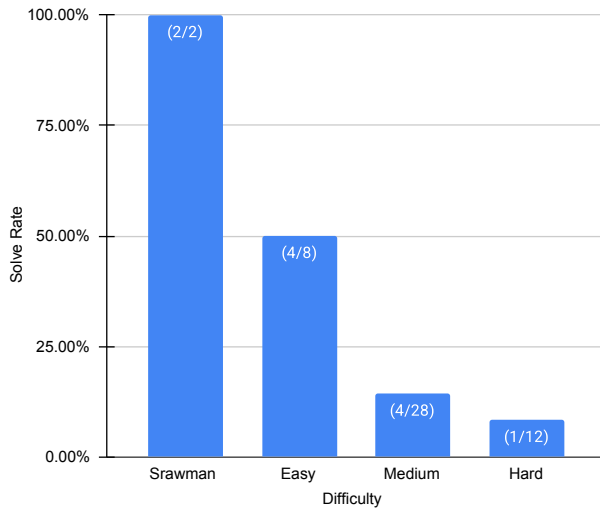
Figure 11 | Challenge solve rates as a function of difficulty level.

tion or mitigation by highlighting attack patterns most likely to change due to AI. Our evaluations showed high scores for evasion and operational security (maintaining persistence, evading detection post-access), primarily relevant in later stages like Installation (e.g., side-loading, living-off-the-land, disabling security) and C2 (e.g., encrypted channels, hiding traffic). The results suggest moderate effectiveness in aiding attackers maintain access undetected.

Development and Deployment of Targeted Mitigations. After mapping capabilities and assessing gaps, we develop targeted safeguards (safety fine-tuning, misuse filtering, response protocols) following our Frontier Safety Framework Google DeepMind (2025). Mitigation robustness is assessed via assurance evaluations, threat modeling, and safety cases (Goemans et al., 2024), considering misuse likelihood and consequences. We periodically assess safeguards through red-teaming and update threat models with new cyber capability evaluations, as capabilities and tactics evolve.

Grounding AI-enabled Adversary Emulation. The framework also informs proactive adversary emulation. Adversary emulation assesses security by applying threat intelligence about specific adversary TTPs to emulate threats, verifying detection/mitigation across the attack chain. Our framework helps red teams more accurately model AI-enabled adversary behavior (Figure 13). Combining knowledge of adversary TTPs, prevalence of AI use in specific phases (Figure 6), and evidence of AI-enabled cost reduction (Figure 12) allows creation of more realistic emulation scenarios to test defenses against AI-leveraging actors.

Benchmarking Defenses. The approach can benchmark defense effectiveness by assessing costs imposed on leveraging AI for specific attack phases (Figure 14). Cyber defense aims to increase attacker costs. While various defenses exist against AI-enabled attacks (model-level, post-deployment), a comprehensive framework for evaluating them across the attack chain is lacking. Our framework can assess intervention effectiveness in making AI-enabled attacks less efficient, potentially deterring them.

7. Related Work

AI has long been integral to cybersecurity, from malware detection to network analysis using predictive models. Recent frontier AI developments spurred research into defensive applications: vulnerability identification (Akuthota et al., 2023; Al-Karaki et al., 2024; Du et al., 2024; Li et al., 2021; Lu et al., 2024), incident summarization (Aminanto et al., 2020; Ban et al., 2023; Khare et al., 2023), incident response (Hays and White, 2024), and other foundational tasks (Alam et al., 2024). DARPA's AIxCC competition (DARPA, 2025) showcased autonomous systems finding, exploiting, and fixing vulnerabilities (Du et al., 2024; Ristea et al., 2024; Ruan et al., 2024). However, the dual-use nature of cyber capabilities necessitates robust risk understanding and management. Consequently, research has grown on methods to evaluate potential risks of capable AI systems in cyber.

7.1. Capture-the-Flag Challenges

CTF challenges are the most common method for evaluating LLM offensive cyber capabilities. LLMs interface with CTF environments to solve security puzzles (cryptography, reverse engineering, web exploitation, etc.) by finding hidden 'flags'. Numerous works employ CTF benchmarks,
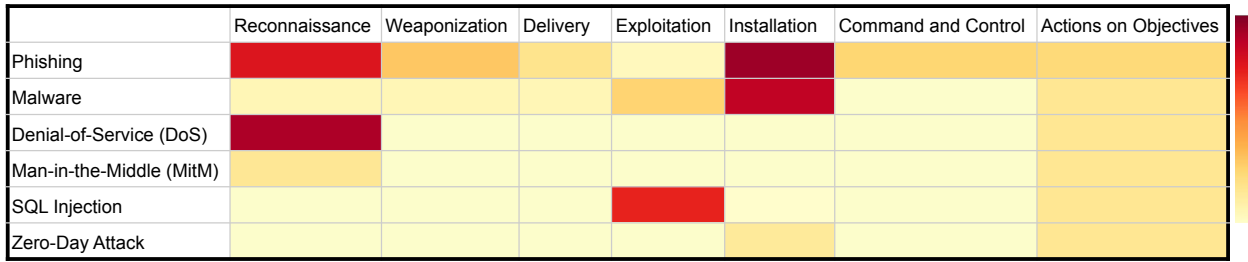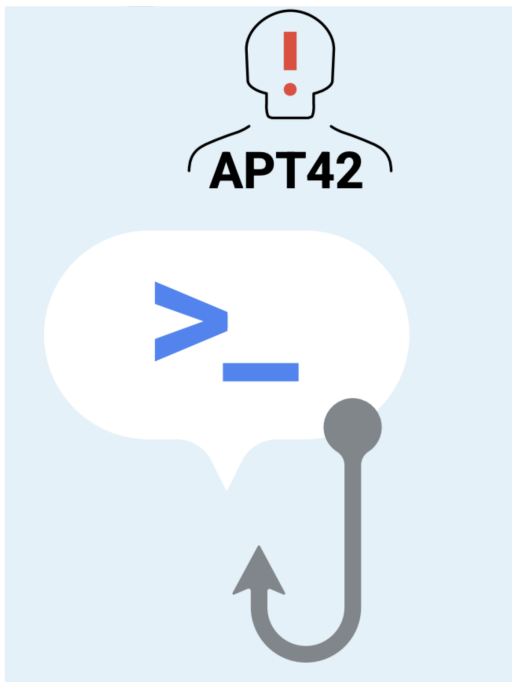
	Reconnaissance	Weaponization	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
Phishing							
Malware							
Denial-of-Service (DoS)							
Man-in-the-Middle (MitM)							
SQL Injection							
Zero-Day Attack							

Figure 12 | Heatmap illustrating potential cost reduction across attack chain phases based on current model capability evaluations.
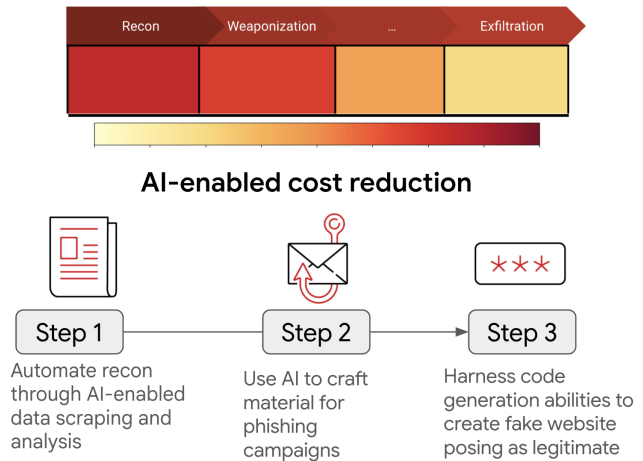

Figure 13 | Framework enabling red teams to better model AI-enabled adversary behavior for testing defenses by generating emulation plans combining TTP knowledge with AI-enabled cost reduction evidence.

including PentestGPT (Deng et al., 2023), CyberSecEval 3 (Bhatt et al., 2024; Wan et al., 2024), Google DeepMind evaluations (Phuong et al., 2024), PenHeal (Huang and Zhu, 2023), AutoAttacker (Xu et al., 2024), Cybench (Zhang et al., 2024), EnIGMA (Abramovich et al., 2024), and InterCode-CTF (Yang et al., 2023a). Some, like CyberSecEval 3, assess "copilot" scenarios with human operators using LLMs. Others focus on narrow benchmarks like Linux privilege escalation (Ban et al., 2023; Lu et al., 2024), while CyberSecEval 3 covers a broader range but still limited attack phases. A drawback is the artificial constraints and simplified scenarios compared to

real-world attacks (e.g., single target vs. complex enterprise networks), potentially skewing capability assessment.

7.2. Multiple Choice and Free Response Tests

Multiple-choice question benchmarks offer measurability and scalability (Liu, 2023; Tann et al., 2023; Tihanyi et al., 2024; Wan et al., 2024). However, creating questions resistant to memorization that accurately reflect offensive cyber is challenging. CyberSecEval also uses free-response questions evaluated by another LLM. OCCULT (Kouremetis et al., 2025) introduced a

Figure 14 | Benchmarking defensive intervention effectiveness by assessing cost imposition on AI-enabled attacks across the attack chain.

multiple-choice benchmark for offensive tactic knowledge.

7.3. Scaffolding and Capability Elicitation

Research is emerging on capability elicitation and model scaffolding to measure upper-bound capabilities. Some systems are lightweight wrappers for action-observation loops (e.g., Cybench (Zhang et al., 2024), InterCode-CTF (Yang et al., 2023a)). Others offer moderate scaffolding (e.g., Vulnhuntr (Du et al., 2024), AutoAttacker (Xu et al., 2024)). More complex systems integrate extensive tools, multiple models, reasoning components, and human feedback (e.g., SWE Agent (Yang et al., 2024), PentestGPT (Deng et al., 2023), Project Naptime (Google, 2025c), EnIGMA (Abramovich et al., 2024), Incalmo (Singer et al., 2025)). While current evaluation approaches offer various tools, translating findings into actionable insights for defenders across the attack chain remains unclear. This paper aims to bridge this gap.

8. Conclusion

This paper introduced a novel framework for evaluating frontier AI's cyber capabilities, focusing on the end-to-end attack chain. Grounded in real-world AI misuse attempts, it bridges evaluations and defenses by helping prioritize targeted mitigations. We curated attack chain archetypes and a new benchmark, conducted bottleneck analysis to identify AI cost disruption potential, and showed how the framework illuminates cost impacts, facilitates mitigation prioritization, benchmarks defense effectiveness, and grounds AI-enabled adversary emulation.

Our evaluations revealed that current AI cyber assessments often overlook critical areas like evasion, detection avoidance, obfuscation, and persistence, where AI shows significant potential. We also confirm the importance of assessing misuse for reconnaissance, widespread exploitation, and long-term attacks.

Cybersecurity is dynamic, and AI will accelerate this. We expect AI to alter attack costs, prompting adversary adaptation. Our framework is designed to evolve with AI capabilities. We will continuously update attack chains, bottleneck analyses, and uplift evaluations based on real-world misuse and model evolution to provide defenders with decision-relevant insights. Mitigating misuse requires a community effort, including robust developer safeguards and evolving defensive techniques accounting for AI-driven TTP changes.

Acknowledgement

References

L. Ablon and A. Bogart. Zero days, thousands of nights. *RAND Corporation, Santa Monica, CA*, 2017.

T. Abramovich, M. Udeshi, M. Shao, K. Lieret, H. Xi, K. Milner, S. Jancheska, J. Yang, C. E. Jimenez, F. Khorrami, et al. Enigma: Enhanced interactive generative model agent for ctf challenges. *arXiv preprint arXiv:2409.16165*, 2024.

V. Akuthota, R. Kasula, S. T. Sumona, M. Mohiuddin, M. T. Reza, and M. M. Rahman. Vulnerability detection and monitoring using llm. In *2023 IEEE 9th International Women in Engineering (WIE) Conference on Electrical and Computer Engineering (WIECON-ECE)*, pages 309–314. IEEE, 2023.

J. Al-Karaki, M. A.-Z. Khan, and M. Omar. Exploring llms for malware detection: Review, framework design, and countermeasure approaches. *arXiv preprint arXiv:2409.07587*, 2024.

M. T. Alam, D. Bhusal, L. Nguyen, and N. Rastogi. Ctibench: A benchmark for evaluating llms in cyber threat intelligence. *arXiv preprint arXiv:2406.07599*, 2024.

M. E. Aminanto, T. Ban, R. Isawa, T. Takahashi, and D. Inoue. Threat alert prioritization using isolation forest and stacked auto encoder with day-forward-chaining analysis. *IEEE Access*, 8: 217977–217986, 2020.

Anthropic. Claude 3.7 sonnet system card, 2025. URL https://anthropic.com/claude-3-7-sonnet-system-card.

T. Ban, T. Takahashi, S. Ndichu, and D. Inoue. Breaking alert fatigue: Ai-assisted siem framework for effective incident response. *Applied Sciences*, 13(11):6610, 2023.

M. Bhatt, S. Chennabasappa, C. Nikolaidis, S. Wan, I. Evtimov, D. Gabi, D. Song, F. Ahmad, C. Aschermann, L. Fontana, et al. Purple llama cyberseceval: A secure coding benchmark for language models. *arXiv preprint arXiv:2312.04724*, 2023.

M. Bhatt, S. Chennabasappa, Y. Li, C. Nikolaidis, D. Song, S. Wan, F. Ahmad, C. Aschermann, Y. Chen, D. Kapil, et al. Cyberseceval 2: A wide-ranging cybersecurity evaluation suite for large language models. *arXiv preprint arXiv:2404.13161*, 2024.

Center for Strategic and International Studies. Cyber events database, 2025. URL https://cissm.umd.edu/cyber-events-database.

DARPA. Aixcc: Ai cyber challenge, 2025. URL https://aicyberchallenge.com/.

G. Deng, Y. Liu, V. Mayoral-Vilches, P. Liu, Y. Li, Y. Xu, T. Zhang, Y. Liu, M. Pinzger, and S. Rass. Pentestgpt: An llm-empowered automatic penetration testing tool. *arXiv preprint arXiv:2308.06782*, 2023.

L. Derczynski, E. Galinkin, J. Martin, S. Majumdar, and N. Inie. garak: A framework for security probing large language models. *arXiv preprint arXiv:2406.11036*, 2024.

X. Du, G. Zheng, K. Wang, J. Feng, W. Deng, M. Liu, B. Chen, X. Peng, T. Ma, and Y. Lou. Vul-rag: Enhancing llm-based vulnerability detection via knowledge-level rag. *arXiv preprint arXiv:2406.11147*, 2024.

A. Goemans, M. D. Buhl, J. Schuett, T. Korbak, J. Wang, B. Hilton, and G. Irving. Safety case template for frontier ai: A cyber inability argument. *arXiv preprint arXiv:2411.08088*, 2024.

Google. Google's secure ai framework, 2025a. URL https://safety.google/cybersecurity-advancements/saif/.

Google. Adversarial misuse of generative ai, 2025b. URL https://cloud.google.com/blog/topics/threat-intelligence/adversarial-misuse-generative-ai.

Google. Project naptime: Evaluating offensive security capabilities of large language models, 2025c. URL URl:https://googleprojectzero.blogspot.com/2024/06/projectnaptime.Html.

Google DeepMind. Frontier safety framework 2.0, 2025. URL https://deepmind.google/discover/blog/updating-the-frontier-safety-framework/.

K. Haga, P. H. Meland, and G. Sindre. Breaking the cyber kill chain by modelling resource costs. *Graphical Models for Security*, 2020.

S. Hays and J. White. Employing llms for incident response planning and review. *arXiv preprint arXiv:2403.01271*, 2024.

J. Huang and Q. Zhu. Penheal: A two-stage llm framework for automated pentesting and optimal remediation. In *Proceedings of the Workshop on Autonomous Cybersecurity*, pages 11–22, 2023.

InfoSecurity Magazine. Uk ai safety institute rebrands, 2025. URL https://www.infosecurity-magazine.com/news/uk-ai-safety-institute-rebrands.

A. Jaech, A. Kalai, A. Lerer, A. Richardson, A. El-Kishky, A. Low, A. Helyar, A. Madry, A. Beutel, A. Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.

T. Johansmeyer. Perception shapes reality: How views on financial market correlation affect capital availability for cyber insurance. *Journal of Risk Management and Insurance*, 2024.

A. Khare, S. Dutta, Z. Li, A. Solko-Breslin, R. Alur, and M. Naik. Understanding the effectiveness of large language models in detecting security vulnerabilities. *arXiv preprint arXiv:2311.16169*, 2023.

M. Kouremetis, M. Dotter, A. Byrne, D. Martin, E. Michalak, G. Russo, M. Threet, and G. Zarrella. Occult: Evaluating large language models for offensive cyber operation capabilities. *arXiv preprint arXiv:2502.15797*, 2025.

Z. Li, D. Zou, S. Xu, X. Ou, H. Jin, S. Wang, Z. Deng, and Y. Zhong. Vuldeepecker: A deep learning-based system for vulnerability detection. *arXiv preprint arXiv:1801.01681*, 2018.

Z. Li, D. Zou, S. Xu, H. Jin, Y. Zhu, and Z. Chen. Sysevr: A framework for using deep learning to detect software vulnerabilities. *IEEE Transactions on Dependable and Secure Computing*, 19(4):2244–2258, 2021.

Z. Liu. Secqa: A concise question-answering dataset for evaluating large language models in computer security. *arXiv preprint arXiv:2312.15838*, 2023.

Lockheed Martin. Cyber kill chain, 2025. URL https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html.

G. Lu, X. Ju, X. Chen, W. Pei, and Z. Cai. Grace: Empowering llm-based software vulnerability detection with graph structure and in-context learning. *Journal of Systems and Software*, 212:112031, 2024.

Mandiant. Mandiant advantage: Threat intelligence, 2025. URL URl:https://advantage.mandiant.com/.

M. Phuong, M. Aitchison, E. Catt, S. Cogan, A. Kaskasoli, V. Krakovna, D. Lindner, M. Rahtz, Y. Assael, S. Hodkinson, et al. Evaluating frontier models for dangerous capabilities. arxiv. *arXiv preprint arXiv:2403.13793*, 2024.

D. Ristea, V. Mavroudis, and C. Hicks. Ai cyber risk benchmark: Automated exploitation capabilities. *arXiv preprint arXiv:2410.21939*, 2024.

H. Ruan, Y. Zhang, and A. Roychoudhury. Specrover: Code intent extraction via llms. *arXiv preprint arXiv:2408.02232*, 2024.

M. Shao, S. Jancheska, M. Udeshi, B. Dolan-Gavitt, K. Milner, B. Chen, M. Yin, S. Garg, P. Krishnamurthy, F. Khorrami, et al. Nyu ctf bench: A scalable open-source benchmark dataset for evaluating llms in offensive security. *Advances in Neural Information Processing Systems*, 37: 57472–57498, 2024.

T. Shevlane, S. Farquhar, B. Garfinkel, M. Phuong, J. Whittlestone, J. Leung, D. Kokotajlo, N. Marchal, M. Anderljung, N. Kolt, et al. Model evaluation for extreme risks. *arXiv preprint arXiv:2305.15324*, 2023.

B. Singer, K. Lucas, L. Adiga, M. Jain, L. Bauer, and V. Sekar. On the feasibility of using llms to execute multistage network attacks. *arXiv preprint arXiv:2501.16466*, 2025.

B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas. Mitre att&ck: Design and philosophy. In *Technical report*. The MITRE Corporation, 2018.

W. Tann, Y. Liu, J. Sim, C. Seah, and E. Chang. Using large language models for cybersecurity capture-the-flag challenges and certification questions. arxiv 2023. *arXiv preprint arXiv:2308.10443*, 2023.

N. Tihanyi, M. A. Ferrag, R. Jain, and M. Debbah. Cybermetric: A benchmark dataset for evaluating large language models knowledge in cybersecurity. *arXiv preprint arXiv:2402.07688*, 2024.

S. Wan, C. Nikolaidis, D. Song, D. Molnar, J. Crnkovich, J. Grace, M. Bhatt, S. Chennabasappa, S. Whitman, S. Ding, et al. Cyberseceval 3: Advancing the evaluation of cybersecurity risks and capabilities in large language models. *arXiv preprint arXiv:2408.01605*, 2024.

J. Xu, J. W. Stokes, G. McDonald, X. Bai, D. Marshall, S. Wang, A. Swaminathan, and Z. Li. Autoattacker: A large language model guided system to implement automatic cyber-attacks. *arXiv preprint arXiv:2403.01038*, 2024.

J. Yang, A. Prabhakar, K. Narasimhan, and S. Yao. Intercode: Standardizing and benchmarking interactive coding with execution feedback. *Advances in Neural Information Processing Systems*, 36:23826–23854, 2023a.

J. Yang, A. Prabhakar, S. Yao, K. Pei, and K. R. Narasimhan. Language agents as hackers: Evaluating cybersecurity skills with capture the flag. In *Multi-Agent Security Workshop@ NeurIPS'23*, 2023b.

J. Yang, C. Jimenez, A. Wettig, K. Lieret, S. Yao, K. Narasimhan, and O. Press. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652, 2024.

A. K. Zhang, N. Perry, R. Dulepet, J. Ji, C. Menders, J. W. Lin, E. Jones, G. Hussein, S. Liu, D. Jasper, et al. Cybench: A framework for evaluating cybersecurity capabilities and risks of language models. *arXiv preprint arXiv:2408.08926*, 2024.

A. Bottleneck Analysis

A.1. Initial analysis of bottlenecks

We began the process of identifying an initial candidate list of bottlenecks by ingesting Mandiant's Threat Intelligence dataset. This data source consists of a set of detailed deconstructions of cyber attacks, timelines, and in-depth analysis of each phase of how an attack unfolded. The sources of this dataset include real-world data from Mandiant's work helping organizations recover from breaches, operational intelligence from monitoring and defending client networks, and adversarial intelligence obtained through in-depth research and analysis of threat actors.

Phishing and spear phishing

Bottlenecks:

- Gathering information about the target (individual, organization, or group)
- Creating a malicious payload (e.g., malware-laden attachment, malicious link)
- Transmitting the malicious payload to the target

AI Uplift Potential:

- Automating Reconnaissance: AI-powered tools can quickly gather and analyze vast amounts of OSINT data to identify potential victims and tailor phishing lures.
- Improving Weaponization: AI can generate highly convincing phishing emails and websites, making it more difficult for users to identify them as malicious.
- Personalizing Attacks: AI can analyze individual user behavior to craft highly targeted phishing attacks that are more likely to be successful.

Malware

Bottlenecks:

- Development of new malware: Requires reverse engineering, understanding vulnerabilities, coding skills, and potentially evading detection mechanisms.
- Evading detection: Signature-based detection, sandboxing, behavioral analysis. Malware needs to be polymorphic or metamorphic to avoid detection.
- Distribution/Delivery: Getting the malware onto target systems can be challenging. Exploiting vulnerabilities, social engineering, supply chain attacks.
- Maintaining persistence: Ensuring the malware stays on the system even after reboots or security scans.
- Command and Control (C2) communication: Establishing and maintaining covert communication channels for control and data exfiltration.

AI Uplift Potential:

- Automated malware generation: AI can automate the creation of new malware variants, including polymorphic and metamorphic malware to evade signature-based detection.
- Intelligent evasion techniques: AI can learn and develop techniques to evade sandboxing and behavioral analysis by mimicking benign behavior or detecting sandbox environments.

- Automated vulnerability exploitation: AI can be used to find and exploit vulnerabilities to deliver and install malware automatically.
- Enhanced C2 communication: AI can establish more resilient and stealthy C2 channels, potentially using techniques like domain generation algorithms (DGAs) or encrypted communications that adapt to network conditions.
- Targeted malware: AI can tailor malware payloads and behaviors to specific targets, increasing effectiveness and reducing detection.

Denial-of-Service (DoS)

Bottlenecks:

- Amplification: Generating enough traffic to overwhelm a target infrastructure can be difficult without amplification techniques.
- Bypassing mitigation strategies: Rate limiting, firewalls, intrusion detection/prevention systems, content delivery networks (CDNs).
- Maintaining attack persistence: Keeping the attack going continuously can be resource intensive, and mitigation strategies might eventually become effective.
- Attribution and anonymity: Hiding the source of the attack can be challenging and important for avoiding repercussions.

AI Uplift Potential:

- Intelligent amplification attacks: AI could optimize amplification techniques to maximize the impact of DoS attacks with fewer resources, potentially by dynamically adapting attack vectors.
- Automated DDoS orchestration: AI can automate the orchestration of large-scale DDoS attacks, managing botnets and attack vectors more efficiently.
- Evasion of mitigation: AI can learn and adapt to bypass rate limiting, firewalls, and other mitigation strategies by identifying weaknesses in defensive systems and dynamically changing attack patterns.
- Creation of more complex and stealthy DoS attacks: AI might enable development of

application-layer DoS attacks that are harder to detect and mitigate than simple volumetric attacks.

- Autonomous botnet management: AI could manage botnets more autonomously and effectively, improving their resilience and attack capabilities.

Man-in-the-Middle (MitM)

Bottlenecks:

- Network positioning: Gaining a position on the network path between two communicating parties (e.g., ARP poisoning, rogue Wi-Fi access points).
- Traffic interception: Capturing and potentially decrypting network traffic. Encryption (HTTPS, TLS) makes interception and decryption harder.
- Real-time traffic analysis: Analyzing intercepted traffic in real-time to extract valuable information or identify opportunities for manipulation.
- Traffic manipulation/injection: Modifying traffic without being detected, which requires understanding the protocols and application logic.
- Maintaining stealth: Avoiding detection while intercepting and potentially manipulating traffic.

AI Uplift Potential:

- Automated network positioning: AI can automate network reconnaissance and identify optimal positions for MitM attacks.
- Intelligent traffic analysis: AI can perform deep packet inspection and real-time analysis of encrypted traffic to identify patterns, vulnerabilities, or sensitive data even without full decryption, potentially using techniques like traffic analysis and machine learning.
- Dynamic traffic manipulation: AI could automate the dynamic manipulation of traffic based on real-time analysis, enabling more sophisticated and context-aware attacks.

- Bypassing encryption or finding weaknesses in implementations: AI could potentially find subtle weaknesses in encryption protocols or implementations that can be exploited for partial or full decryption in certain scenarios.
- Automated injection of malicious content: AI can inject malicious content into traffic streams in a way that is less likely to be detected and more likely to achieve the attacker's objectives

. **SQL Injection**

Bottlenecks:

- Finding vulnerable parameters: Identifying input fields in web applications that are vulnerable to SQL injection.
- Crafting effective injection payloads: Developing SQL queries that can bypass input validation and achieve the desired outcome (data exfiltration, modification, etc.).
- Bypassing web application firewalls (WAFs): WAFs are designed to detect and block common SQL injection attacks.
- Exploiting complex SQL injection scenarios: Blind SQL injection, time-based injection, second-order injection can be more complex to exploit.
- Automating the exploitation process: Manually testing for and exploiting SQL injection can be time-consuming.

AI Uplift Potential:

- Automated vulnerability scanning and identification: AI can crawl web applications and automatically identify potential SQL injection vulnerabilities with greater accuracy and speed.
- Intelligent payload crafting: AI can generate SQL injection payloads that are more likely to bypass input validation and WAFs, potentially using techniques like mutation and adversarial examples.
- Automated exploitation of complex scenarios: AI can automate the exploitation of blind, time-based, and second-order SQL injection vulnerabilities, significantly reducing the time and effort required.

- Learning WAF evasion techniques: AI can learn from WAF responses and develop evasion techniques that are more effective.
- Optimized data exfiltration: AI can optimize data exfiltration strategies after successful SQL injection to minimize detection and maximize data retrieved.

Zero-Day Attack

Bottlenecks:

- Vulnerability discovery: Finding previously unknown vulnerabilities is extremely difficult and time-consuming, requiring deep expertise and resources.
- Exploit development: Creating a reliable exploit for a zero-day vulnerability that works across different systems and is not easily detected.
- Weaponization and delivery before patching: Attacks need to be carried out before the vulnerability is publicly disclosed and patched, requiring speed and stealth.
- Maintaining secrecy of the vulnerability: Keeping the zero-day vulnerability secret is crucial for its long-term effectiveness.
- Target selection and impact maximization: Choosing targets where the zero-day exploit will have maximum impact.

AI Uplift Potential for Zero-Day Attacks:

- Accelerated vulnerability discovery: AI can analyze codebases and software systems at scale to identify potential zero-day vulnerabilities much faster than traditional methods, using techniques like fuzzing, symbolic execution, and machine learning for anomaly detection in code.
- Automated exploit generation: AI can automate the process of generating exploits for discovered vulnerabilities, reducing the time and attack barrier for exploit development.
- Proactive vulnerability prediction: AI might be able to predict potential vulnerability types or locations in software based on code patterns and past vulnerability data, guiding vulnerability research efforts.

- Stealthy zero-day weaponization: AI can help create zero-day exploits and delivery mechanisms that are more stealthy and harder to detect, maximizing the window of opportunity before patching.
- Targeted zero-day attacks: AI can analyze potential targets and identify those where a specific zero-day exploit

A.2. Validating Attack Bottlenecks via Expert Survey

To validate potential bottlenecks in cyberattacks, we first selected thirteen historical case studies exhibiting exceptionally high impact from a broader set of representative attack chains. The selection criteria included attacks that caused, or had the potential to cause, catastrophic loss of life, significant economic damage, or state-level cyber espionage. (See Figure 15). Assessing the economic impact of attacks is inherently difficult; company reports might overestimate losses (as revenue may shift rather than disappear), while negative supply chain effects can be underestimated. Therefore, impact assessment requires considerable subjective judgment.

We identified suitable case studies using sources such as the Kent Academic Repository (Johansmeyer, 2024), CSIS reports, Florian Roth's summaries, security company publications, and media reports. To focus on the modern defensive landscape, we excluded attacks that occurred before 2010.

Next, we recruited ten offensive security experts from Google. They estimated the resources required to replicate the selected attacks under specific assumptions: the attack would be performed by a hypothetical median-skilled cyber expert at Google, without using AI tools. This standardization helps ensure comparable estimates regardless of the original attacker's capabilities.

The survey proceeded as follows:

1. Each expert estimated the resources needed for each phase of two distinct case studies. We provided detailed descriptions of the attacker's actions and achieved outcomes for every phase.

Critical Infrastructure Attacks	Destructive / Ransomware Attacks
- Ukraine Power Grid 2015 - Stuxnet 2010	- WannaCry 2017 - NotPetya 2017 - Ireland Ransomware of HSE 2021
Nation-State Espionage Attacks	Major Data Breaches
- SolarWinds 2020 - Duqu 2.0 2015 - Storm-0558 Microsoft Intrusion 2023 - OPM Hacker X2 Attack 2015	- Target 2013 - Anthem 2015 - Equifax 2017 - Marriott 2018

Figure 15 | Notable, high-impact case studies selected for expert verification.

2. Experts provided estimates in two forms: human-days of effort for a median expert, and direct monetary expenses (e.g., hardware procurement, infrastructure costs). They also rated their confidence (low, medium, high) for each estimate. Following Haga et al. (2020), we broke down attacks into phases to simplify estimation and requested cost intervals rather than point estimates to improve generalizability.

3. To establish a plausible upper bound on phase efficiency and avoid assuming replication of historical attacker errors, experts also estimated the minimum resources each phase might require for a sophisticated adversary.

4. This process resulted in seven case studies reviewed by two experts and six reviewed by one. We then conducted a consensus-building exercise: for doubly-reviewed cases, experts analyzed their peer's estimates and reasoning and could optionally revise their own; for singly-reviewed cases, a second expert provided anonymous feedback on the initial assessment.

This survey methodology helps determine the relative costs of different attack phases and identify dominant bottlenecks. Understanding these bottlenecks allows defenders to qualitatively assess how enhancing defenses at one stage might affect the overall attack economics and potentially anticipate significant shifts in attack costs.

We took the initial human-day estimates and converted them to dollar costs by multiplying by a salary variable ($500k); for ease of cost comparison between phases, we then calculated the percent of the overall resourcing estimate each phase accounted for. We define a 'bottleneck' as any phase requiring at least 10% of the total estimated resources for the attack.

Our key findings were that:

- Weaponization and Reconnaissance were very common bottlenecks, and also commanded the overall highest amount of resourcing (averaged across scenarios where they were bottlenecks)
- For major data breaches, bottlenecks were more evenly distributed across phases compared to other attack types
- Averages across case studies:
  - Weaponization (n=9, 35.94%, $206.7k)
  - Reconnaissance (n=7, 19.20%, $44.8k)
  - Actions on Objectives (n=5, 17.60%, $525k; 11.00%, $16.5k when excluding SolarWinds outlier)
  - Exploitation (n=4, 9.82%, $19.3k)
  - Delivery (n=4, 8.30%, $15.8k)
  - C2 (n=3, 6.40%, $17.3k)
  - Installation (n=2, 3.20%, $8.3k)