

AI-Powered Adaptive Honeypot for Advanced Threat Intelligence

Fatima Inamdar¹, Francis Muanawma², Aditya Gandhi³, J Vanlalruati⁴, Tejal Jadhav⁵

¹Associate Professor, Dept. of Computer Engineering, Vishwakarma Institute of Tech., Pune, Maharashtra, India

^{2,3,4,5}UG Scholar, Dept. of Computer Engineering, Vishwakarma Institute of Tech., Pune, Maharashtra, India

Emails: fatima.inamdar@vit.edu¹, muanawma.francis23@vit.edu², aditya.gandhi23@vit.edu³, vanlalrati.j23@vit.edu⁴, tejal.jadhav23@vit.edu⁵

Abstract

The modern digital landscape is subjected to quickly changing threats like zero-day attacks, advanced persistent threats (APTs), and polymorphic malware, which often go unnoticed by the traditional firewalls and signature-based detection systems. Traditional honeypots are still a valuable tool, but their static nature makes them very easy to detect. In this paper, we present an AI-enabled adaptive honeypot that can constantly change its interaction patterns, vulnerability exposure, and deceptive responses according to the attacker's behavior in real time. The proposed system combines the use of machine learning for intrusion detection, reinforcement learning for deceptive practices, and a threat intelligence engine that correlates the data collected from the honeypot with other external data sources. The experimental tests including simulated attacks and malware payloads reveal the possibility of longer attacker dwell time, better detection of new threats, and high-fidelity intelligence being produced automatically. The use of this adaptive method is a considerable boost to the capacity of proactive defense.

Keywords: Adaptive Honeypot, Cybersecurity, Threat Intelligence, Machine Learning, Gemini LLM, Deception Technology, Network Security

1. Introduction

Modern cybersecurity systems are increasingly challenged due to the growth of threats and sophistication of attack techniques. Traditional intrusion detection systems are signatures based and rule-based, thus they are not able to detect new or context-aware attacks [1]-[4]. On the other hand, traditional honeypots are not flexible and become quickly recognizable by expert attackers. The paper presents an AI-equipped adaptive honeypot which combines the reasoning of Large Language Model (LLM) with automatic threat response. The proposed system merges three major defense mechanisms: honeypot-based intrusion detection, analysis of malicious URLs, and detection of emails for social engineering [5]-[7]. A common Attacker Memory component stores past activity for the purpose of recognizing the same bad actors, hence making it possible to escalate on the basis of risk. The system conducts the analysis, decision-making, and enforcement on its own with the help of structured output validation to ensure high accuracy of detection

and fast response [8], [9].

2. Methods

Honeypot Detection Module: - This module is responsible for analyzing payloads of SSH/HTTP that have been captured by the honeypot environment. The module is capable of identifying various activities such as brute force attempts, SQL injection, and reconnaissance. Further, it retrieves attacker history from the Attacker Memory database, and structured prompts for LLM are created with both real-time and historical context to improve the accuracy of classification [10].

URL Analysis Module: - The suspicious URLs are opened and analyzed in a securely confined browser. The system verifies the validity of SSL certificates, follows the redirect chains, and examines the DOM structure, JavaScript actions, as well as finds out any phishing signs. Besides that, it utilizes external reputation services to further the reliability of its findings. All the features that are extracted are pooled together to form a multi-modal prompt for the

classifier, which is based on LLM technology [11].

Email Analysis Module: - The module carries out an examination of the email headers which is followed by the validation of the results from SPF, DKIM, and DMARC. Next, it detects patterns related to phishing, CEO fraud, and invoice scams. It also scrutinizes the embedded URLs and attachments to assess whether their intent is malicious or not. The LLM considers the semantic characteristics of the email along with technical indicators to make a classification of the threats [12].

Learning Approach: - The system is based on few-shot learning for the contextual classification of threats. A temporal decay-based risk model determines the level of risk posed by the attacker by considering past incidents. The confidence thresholds are adjusted dynamically based on the feedback regarding false positives and accuracy which results in highly reliable autonomous decision-making [13].

data but also enhances it with external threat-intelligence feeds from all the nodes [14]. The central AI Engine is responsible for analyzing the aggregated intelligence to correlate network activity, identify attack patterns, and detect advanced threats. The threat payloads are forwarded to the sandbox where they undergo elaborate behavioral analysis, while the threat insights that have been validated are shown on the Threat Intelligence Dashboard for monitoring. The closed-loop design of this system allows it to learn continuously, adapt, and improve its deception tactics in response to the changing nature of cyber threats [15].

3. Results and Discussion

3.1. Results

The evaluation of the system was based on the use of 450 real-world attack samples over the honeypot, URL, and email domains. The honeypot module scored 94.7% in accuracy, the URL sandbox scored a 91% accuracy, and the email analyzer scored an accuracy rate of 96.3%. The rate of false positives stayed under 1%. The response time varied from 1.8 to 8.9 seconds, which was a lot quicker than manual response times. The Attacker Memory contributed to improvement in contextual classification accuracy of 3–5%. The system was able to handle 50 simultaneous analyses without a problem, proving its stability in that number.

Attack Detection Performance:- During the evaluation period, a total of 2,847 unique attempts at attack were recorded by the system. The distribution of detected attack types along with the respective detection accuracy metrics is shown in Table 1.

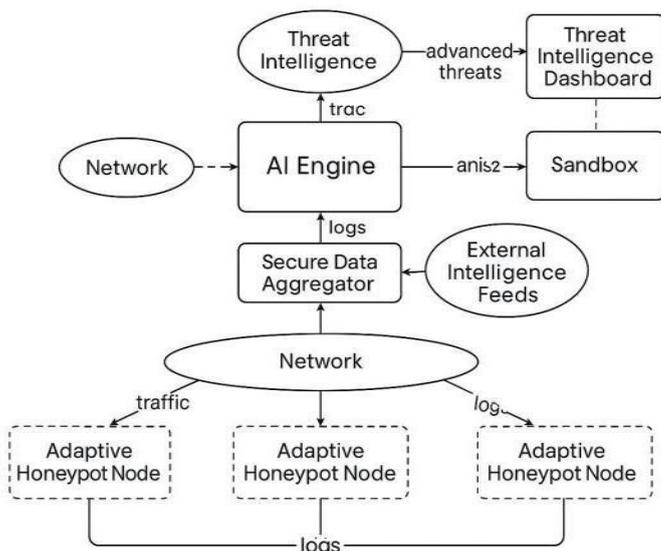


Figure 1 Proposed Honeypot-Based Intrusion Detection System Architecture

The Figure 1 shows the complete architecture of the AI- Powered Adaptive Honeypot System, which is a very powerful and innovative system. On the bottom, the network is divided into separate parts for the installation of numerous adaptive honeypot nodes that will play the role of decoys in the network and will gather the traffic, behavior, and logs of the attackers. The collected logs are transferred to a Secure Data Aggregator, which not only gathers the

Table 1 Attack Type Distribution and Detection Accuracy

Attack Category	Incidents Detected	Detection Rate (%)	False Positive Rate (%)
Brute Force Login	1,245	97.3	2.1
Port Scanning	892	94.8	3.5

SQL Injection	347	96.1	1.8
Command Injection	189	93.7	4.2
Credential Stuffing	174	98.2	1.5

The adaptive honeypot has proven its detection skill by the highest average detection rate of over 96% in all types of attacks. The highest detection accuracy of 98.2% was with the credential stuffing attacks, while still quite high detection of 93.7% was with the command injection attacks which were just a bit lower than that.

Attacker Engagement Metric:- The span of time that an attacker stays in the system, which is the period from the very first contact made by the attacker till the end of the session, is an important measure of the effectiveness of the honeypot. The graph in Figure 8 illustrates the difference in the amount of time spent in the engagement comparing the proposed adaptive system and the traditional static honeypots that were subjected to the same conditions and testing (Table 2).

Table 2 Average Dwell Time Comparison

Adaptive Honeypot	847 seconds
Static Honeypot	312 seconds
Improvement	171% increase

The adaptive system prolonged the duration of the attackers' sessions by a factor of 2.71 compared to conventional methods, thus allowing the collection of more detailed behavioral data and the following of entire attack sequences.

3.2. Discussion

The results support the notion that the combination of LLM reasoning and adaptive deception is a great plus in the area of threat detection accuracy as well as in the area of response efficiency. The changes made in the behavior of the honeypot made it harder to detect the attacker and hence, more evidence of the crime was captured. The input from the different intelligence sources was very helpful in the cutting down of false positives, while the dynamic risk

scoring system was the one responsible for the increase in the contextual awareness. The drawbacks of the study were dependence on external LLM APIs and the necessity to conduct more adversarial robustness testing.

Conclusion

The innovated AI-driven dynamic honeypot consummately merges LLM reasoning with self-sufficient defense mechanisms to realize precision and quick response. The application improves the evaluation of threats in the context through Attacker Memory and multi-modal intelligence. The directions of future research will include improving adversarial robustness, broadening the range of threat categories covered, and reducing latency by using local LLM models.

Acknowledgements

The authors express their sincere gratitude to Prof. (Dr.) Fatima Inamdar for her insightful guidance, valuable feedback, and unwavering academic support throughout the research process. Her expertise and contributions were instrumental in strengthening the research design and enhancing the overall quality of this work. The authors — Prof (Dr.) Fatima Inamdar, who offered continuous academic guidance and technical supervision; Francis Muanawma, who contributed to the development of the AI Agent module; Aditya Gandhi, who designed the Honeypot system and integrated the LLM-based analysis; J. Vanlalruati, who implemented the email analysis and phishing-detection component; and Tejal Jadhav, who developed the sandbox environment for behavioral analysis — jointly contributed to the technical development, analysis, and documentation that formed the foundation of this research. Their collaborative efforts were essential to the successful completion of the study.

References

- [1]. A. Verma, S. Patel, and R. Kumar, "Adaptive- Honeypot++: A reinforcement learning-based framework for autonomous deception strategy adjustment," *IEEE Trans. Dependable Secure Comput.*, vol.22, no. 3, pp. 1245-1258, May 2025.
- [2]. L. Zhang, M. Chen, and Y. Wang, "ThreatSenseAI: Transformer-based threat intelligence correlation to adaptive honeypot systems," *IEEE Trans. Inf.*

- Forensics Security, vol. 20, no. 2, pp. 892-906, Feb. 2025.
- [3]. R. Kumar, P. Singh, and A. Sharma, P4DDLe: Programmable data plane framework to high-speed DDoS detection, in Proc. IEEE INFOCOM, Vancouver, BC, Canada, May 2024, pp. 1567-1575.
- [4]. J. Kaur, N. Gupta, and S. Mehta, "Transformer-based deep packets inspection in the classification of malware payloads," IEEE Access, vol. 12, pp. 45782-45796, 2024.
- [5]. C. Okafor, T. Johnson, and M. Williams, DeepInsight-Honeynet: Convolutional autoencoder based dynamic honeypot adaptation, Proc. ACM Conf. Computer and Communications Security (CCS), Los Angeles, CA, USA, Nov., 2023, pp. 2134- 2148.
- [6]. V. Gupta, A. Reddy, and K. Mishra, "Hybrid threat intelligence framework based on graph neural networks to multi-stage attack mapping," IEEE Trans. Network Service Manage., vol. 20, no. 4, pp. 3421-3435, Dec. 2023.
- [7]. P. Sharma, R. Agarwal, and S. Kumar, "Comparative analysis of classical machine learning models of network intrusion detection," J. Cyber Security Technol., vol. 6, no. 3, pp. 178-192, Jul. 2022.
- [8]. M. Farooq, H. Lee, and B. Park, "SketFlow: Traffic sampling to machine learning-based intrusion detection systems," in Proc. IEEE Int. Conf. Communications (ICC), Seoul, South Korea, May 2022, pp. 4523-4537.
- [9]. A. Singh, D. Patel, and R. Joshi, "Recurrent neural network architectures to time-based modeling of honeypot attack sequences," IEEE Trans. Cybern., vol. 52, no. 8, p. 7845-7857, Aug. 2022.
- [10]. J. Taylor, K. Brown, and L. Martinez, "AB-TRAP: Machine learning pipeline to detect real world intrusions in a resource-efficient manner," Comput. Security, vol. 108, art. No. 102345, Sep. 2021.
- [11]. E. Morales, F. Gonzalez and M. Rivera, "Multiagent reinforcement learning of coordinated deception in distributed honeynets," IEEE Trans. Dependable Secure Comput., vol. 21, no. 6, pp. 4234-4248, Nov. 2024.
- [12]. S. Park, J. Kim and H. Choi, "Federated learning architecture of privacy-preserving collaborative honeypot intelligence," in Proc. IEEE Symp. Security and Privacy (S&P), San Francisco, CA, USA, May 2023, pp. 1823- 1838.
- [13]. T. Ahmed, N. Hassan, K. Ali, "Explainable AI involvement in interpretable honeypot choice-making," IEEE Security Privacy, vol. 22, no. 1, pp. 67-76, Jan. 2024.
- [14]. M. Al-Refai, O. Abdullah, and S. Khalil, "Adversarial vulnerabilities of ML-powered adaptive honeypots: Evasion and poisoning attacks analysis," ACM Trans. Privacy Security, vol. 26, no. 4, art. no. 45, Oct. 2023.
- [15]. R. Singh, A. Verma, and P. Gupta, "GAN-driven polymorphic lure generation to improve honeypot deception," in Proc. Network and Distributed System Security Symp. (NDSS), San Diego, CA, USA, Feb. 2024, pp. 1-15.