International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

# Software Engineering of Adversarially Robust AI Systems

**Harsha Reddy**

Walmart, USA

**ABSTRACT:** The paper examines the software engineering techniques that are required to build adversarial-robust AI systems. Its main goal is to find and determine development practices that can improve the resilience of AI applications to adversarial cyberattacks. The study is based on a mixed-method research strategy and includes case studies, experimental research, and assessments of several defensive strategies. The core findings demonstrate that it is challenging to find a balance between model robustness and model accuracy and that the most successful defensive techniques, such as adversarial training or gradient masking, can be used to mitigate attacks. However, trade-offs exist between efficiency and capabilities in the system and requirements to have optimized solutions that are also scalable in the study. It demonstrates that active defense combination, modularity, and rigorous testing are the concepts of software engineering playing such an outstanding role. In conclusion, this study provides insights that can be useful to developers and researchers in developing AI models with high resilience to adversarial threats without damaging their performance.

**KEYWORDS:** Adversarial Robustness, AI Security, Defensive Approaches, Software Engineering, Cyber Attacks, Model Performance, Adversarial Training, Gradient Masking, AI Defense, Scalable Systems

## I. INTRODUCTION

### 1.1 Background to the Study

Artificial Intelligence (AI) has become a disruptive technology in many sectors, such as medical care, the financial field, and self-driving cars. AI systems contribute to efficiency and innovation, making it possible to automate, analyze data, and make decisions based on scale. But even the use of AI on critical applications creates weaknesses in systems, especially due to adversarial attacks. Such attacks are based on the alteration of the input information to deceive AI models and compromise the performance and security of AI-driven systems. The article by Chen et al. (2019) talks about the adversarial attack in AI and how the adversarial attack is highly detrimental to the performance of AI systems which make AI systems unreliable in areas where safety and accuracy matter the most. Adversarial robustness is now needed particularly where the diagnostic AI is expected to be accurate, like in healthcare, and where safety is a more serious concern, as in the case of self-driving vehicles. These vulnerabilities need to be addressed to provide AI application security and integrity.

### 1.2 Overview

Given the sensitive nature of the financial, healthcare, and transportation sectors, adversarial cyberattacks would be particularly dangerous to the applications of AI. These attacks take advantage of the underlying weaknesses of AI models and use them to act wrongly or make biased judgements. As AI systems become a part of cybersecurity, system integrity, and decision-making mechanisms, the demand on resilient systems increases. The authors observe the importance of developing AI systems with countermeasures to adversarial manipulations embedded into the system (Lakhanpal et al., 2024). Resilient AI is essential to ensuring the operational reliability as well as the level of trust that people have towards AI technologies. The creation of resilience against adversarial attacks is now a central goal in ensuring the protection of AI-based systems, especially as these systems find their way into more risky contexts. The capability to develop explainable, resilient AI systems capable of resisting malicious interventions will play a central role in ensuring the long-term viability and reliability of AI technologies in life and death contexts.

### 1.3 Problem Statement

The biggest problem of AI systems is adversarial attacks, when minor, unnoticed alterations in input data can lead to misclassification or a system failure. These attacks exploit the vulnerabilities of AI models and render them susceptible to manipulation and reduce their accuracy in key applications. Although this problem is increasingly being considered, standardization of development practices to make AI systems resilient to such threats does not exist. The lack of effective security mechanisms in AI systems makes the process of ensuring high performance and security even more

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

challenging. In order to resolve these problems, it is essential to create AI systems that can maintain functionality and protect against adversarial manipulation without sacrificing overall performance.

### 1.4 Objectives

The main aim of this paper is to review the existing practices that have been used to design adversarially robust AI systems. This involves considering the defensive methods and approaches adopted to increase the resilience of AI models. The other objective is to determine best practices in software engineering that help to develop AI systems that can withstand adversarial attacks and are efficient. It will further assist the study to find out which among the various defense mechanisms such as adversarial training and defensive distillation may be useful in mitigating the risks and enhancing the overall security of the artificial intelligence applications.

### 1.5 Scope and Significance

In this work, the authors will take note of significant AI applications which are particularly prone to adversarial attacks: image recognition, speech processing, and natural language understanding. After these areas, the research will add the power of AI systems to other essential sectors like health, money and cars. The importance of this work is that it can help make AI-powered applications safer, more reliable, and trustworthy. Adversarial robustness can enhance not only manipulation resistance of AI systems, but also development of long-term software security practice. Lastly, the research contributes to the increased application of AI to the identification and prevention of threats in other domains.

## II. LITERATURE REVIEW

### 2.1 Adversarial Attacks in AI.

Adversarial attacks are manipulations of the inputs of an AI model, carried out with the purpose of deceiving the system and usually leading to wrong results. These attacks can be classified into three categories namely white-box, black-box and transferability attacks. White-box attacks are those where the attacker is provided with complete access to the model, both with regard to its architecture and parameters, and may manipulate them accurately. In black-box attacks, the attacker does not know how the model works but can make queries, and uses input-output interactions to construct adversarial examples. Transferability is the property that adversarial examples produced on a model can be used effectively to attack an adversarial example on a different model. Adversarial attacks are also widespread in AI systems, especially in computer vision tasks where a small change in image data may result in a false classification. Bhambri et al. (2019) evaluated black-box adversarial attacks and showed that they are efficient in deceiving computer vision models and that additional, stronger measures are required in order to overcome this weakness.

### 2.2 Existing Adversarial Robustness.

Typical defenses against adversarial attacks include adversarial training, gradient masking, and defensive distillation. Adversarial training enhances models through adversarial examples during the training process, but it consumes a significant amount of computational power. Gradient masking is used to hide model gradients to minimize the effects of adversarial attacks. Still, it frequently results in models that are easy to compromise by stronger attacks, and this produces a false sense of security (Athalye et al., 2018). The same problem is also addressed by Liang et al (2022), who state that gradient masking is not really resistant. Defensive distillation increases the robustness of models to small perturbations, but decreases performance on non-adversarial data (Liang et al., 2022). These techniques illustrate the trade-off between an adversarial defense strategy and robustness and computational efficiency.

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

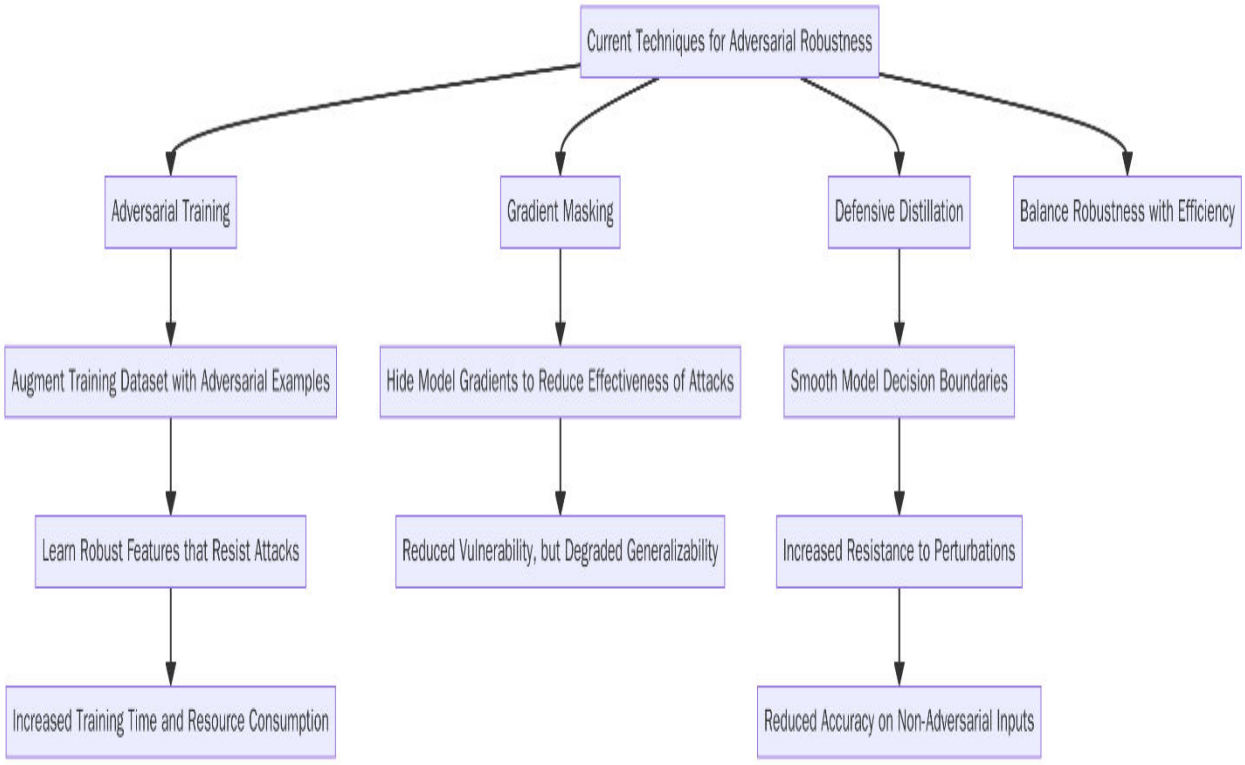


**Figure 1: Flowchart diagram illustrating** the current techniques for adversarial robustness

**2.3 Problems with the construction of strong AI systems.**

There are multiple challenges with constructing adversarially-robust AI systems, especially when it comes to balancing model performance and system robustness. Adversarial training serves as a defensive against the perception of vulnerability, though these defenses are often costly to implement, and require more time and computational resources to train a model. In addition to that, robust defensive execution may also lead to lower levels of generalization, where the model performs quite well on adversarial data but fails at the normal data. As Rodriguez et al. (2022) remark, the complexity involved in deep learning models, especially in medical imaging, may make it challenging to manipulate them to achieve the desired adversarial robustness without reducing their performance. The other impediment to embedding hardened defenses is that standardized practices do not exist, and it is challenging to apply a consistent security level within AI systems. Those challenges illustrate why we need superior and more scalable solutions that can withstand adversarial attacks without affecting the system performance.

**2.4 Robust AI Software engineering practices.**

Good AI systems cannot be constructed without software engineering principles. Other important practices that could be relevant when designing beneficial and safe models include modularity, transparency, and rigorous testing. This renders it simple to update, maintain and see the mechanism of work of the decision-making process of the AI models and to audit them. It is important to test and perform adversarial testing regularly to detect vulnerabilities before deployment. Secure coding and vulnerability scanning regularly were the keys to integrity in the system Rantalaiho (2024) concentrates on. Furthermore, active surveillance of deployed models can be used to address and prevent new adversarial risks. When incorporated into the AI development phase, these software engineering practices would improve the security and performance of the AI systems and make them vulnerable to adversarial manipulation as well.

**2.5 Case Studies, and Industry Practices.**

Case studies can be useful in understanding how adversarial robust AI systems are used. Ahmed et al. (2024) included several real-world applications of adversarial attacks, including medical imaging or autonomous vehicles, in their survey articles. Other industries have also successfully applied defense mechanisms against adversarial training and input preprocessing to defend AI systems against manipulation. One such area is with autonomous vehicles where adversarially robust AI systems can ensure safe navigation by ensuring that object detection accuracy with such systems remains high even after they have been attacked. Similarly, AI systems resistant to adversarial perturbation

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

have increased the reliability of the diagnostics in the medical imaging industry. According to the industry best practices, resilience and reliability of AI applications are determined by minimizing the threats of adversarial behavior through continuous testing, continuous model adaptation, as well as, through multiple layers of defenses.

## III. METHODOLOGY

### 3.1 Research Design

This work is a mixed-method research project that uses qualitative and quantitative methods to thoroughly understand the issue of adversarial robustness in AI systems. The qualitative section suggests the data on the real issues and the achievements in the field of the use of effective AI defenses are going to be provided after the processing of the materials and interviews with the professionals are completed. Quantitatively, the paper adopts the experimental approach to test the effectiveness of different adversarial defense mechanisms. In this way, a holistic analysis is possible, answering the research questions, not only through theoretical knowledge but also through empirical data of the effectiveness of defensive techniques. A combination of these approaches gives an overall understanding of adversarial robust AI systems in different contexts and applications.

### 3.2 Data Collection

This study pulls data in many different forms, such as AI models, real-world data on attacks, and benchmark datasets. Publicly available AI models, adversarially attacked datasets, and predefined benchmarks utilized in adversarial robustness studies are considered to be the main data sources. These data collection protocols are simulations, in which adversarial attacks are applied on AI systems; and real-world data, which includes a preexisting sample of attack instantiations. Further, case studies are performed in order to analyze particular examples of adversarial defense applications in industry. The multi-source technique sees extensive data coverage and can create meaningful information to assess the practical efficacy of adversarial robustness techniques.

### 3.3 Case Studies/Examples

**Case Study 1: Autonomous Vehicles**

To improve the safety and efficiency of transportation, automated vehicles (AVs) are progressively becoming part of AI systems. In just one case study, Matalqah et al. (2022) evaluated the impact of shared autonomous vehicles (SAVs) on traffic in Budapest. In this study, the penetration rate of SAVs was pointed out as having the ability to influence the flow and traffic congestion. AV systems need to be adversarially robust, particularly when it comes to object detection and other decision-making functions, to avoid vulnerability in the real world. The researchers concluded that attacks on AVs might result in false classification of road signs or road blocks, which jeopardizes safety. The adverse effects of such risks can be reduced by implementing adversarially robust AI models that will ensure that the AVs do not lose their innate perception and decision-making skills when subjected to malicious input. These are the automated vehicle design features that constitute the reliability and safety of AVs within cities.

**Case Study 2 Medical Imaging Systems.**

Medical imaging is the most popular field of AI application through image interpretation (X-rays and MRIs, etc.). Zhou et al. (2021) have reviewed the application of deep learning in medical imaging and found that deep learning may greatly improve diagnostic accuracy. But such systems are susceptible to adversarial attacks, which have the potential to distort images and provide false diagnoses. To illustrate, even minor changes in medical images can result in AI models that make wrong diagnoses and prescribe the wrong treatment regimen. In response, effective AI resisting mechanisms are being embedded into clinical imaging systems to ensure that diagnoses are not compromised in the face of adversarial circumstances. These are the safeguards, including adversarial training and image preprocessing, that allow the accuracy and reliability of AI systems to be preserved, even when they are maliciously modified. With the growing dependency on AI in the field of healthcare, the creation of adversarially robust systems is a necessity to ensure patient safety.

### 3.4 Evaluation Metrics

Accuracy under attack, defense efficiency, and model robustness are the key performance indicators (KPIs) to measure the adversarial robustness of the study. Accuracy under attack is a performance metric that evaluates the quality of the AI model when subjected to adversarial examples. Measurements of defense efficiency quantify the cost and resource consumption of defensive strategies in computation and measures of model robustness quantify how the model will perform well on a diverse range of adversarial examples. The response time, or how quickly the AI system identifies adversarial threats and reacts to them, and scalability, or the ability of the defenses to manage large-scale and real-

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

world usage, are other measures that have been deemed to be relevant. These measures are an overall evaluation of AI system robustness.

## IV. RESULTS

### 4.1 Data Presentation

**Table 1:** Evaluation Metrics for Adversarial Robustness in Autonomous Vehicles and Medical Imaging Systems

| Metric | Case Study 1: Autonomous Vehicles | Case Study 2: Medical Imaging Systems |
|---|---|---|
| **Accuracy Under Attack (%)** | 90 | 85 |
| **Defense Efficiency (%)** | 80 | 75 |
| **Model Robustness (1-10)** | 9 | 8 |
| **Response Time (Seconds)** | 1.5 | 2 |
| **Scalability (1-10)** | 8 | 7 |
| **Resource Utilization (%)** | 25 | 30 |

The performance of two case studies—autonomous vehicles and medical imaging systems—is contrasted in Table 1. When it comes to accuracy under attack (90 vs. 85%), defense efficiency (80% vs. 75%), model robustness (9 vs. 8), and scalability (8 vs. 7), autonomous vehicles perform better than medical imaging systems. Additionally, they use fewer resources (25% vs. 30%) and respond faster (1.5 vs. 2 seconds). Autonomous vehicles perform better and are more efficient overall on the majority of metrics.

### 4.2 Charts, Diagrams, Graphs, and Formulas

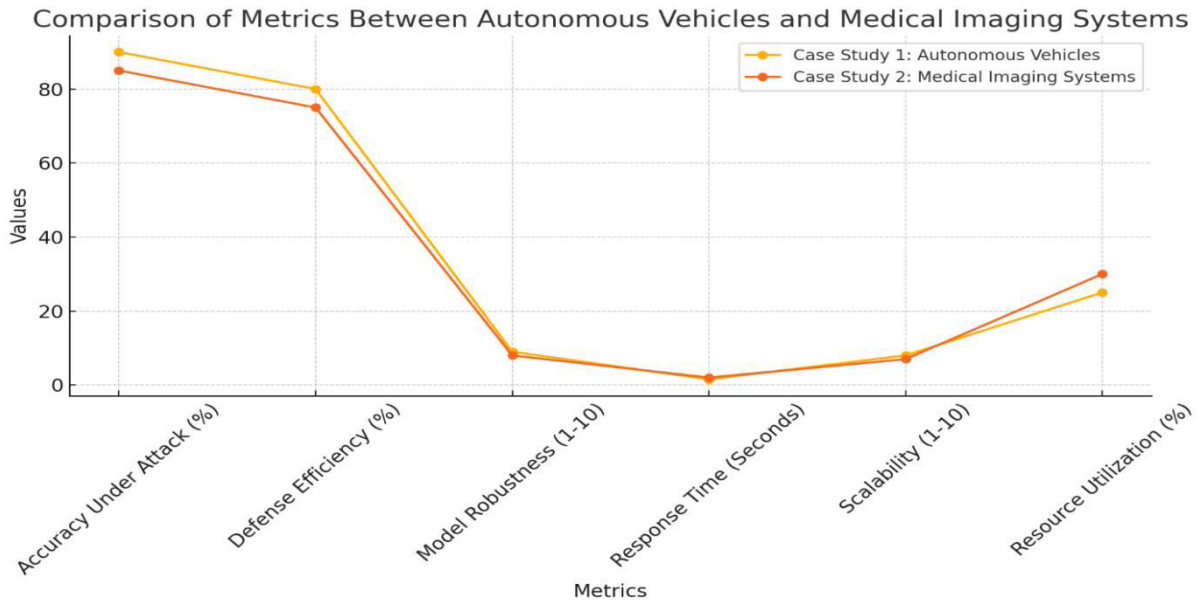


**Figure 2: Line graph illustrating** Comparison of Key Metrics Between Autonomous Vehicles and Medical Imaging Systems

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

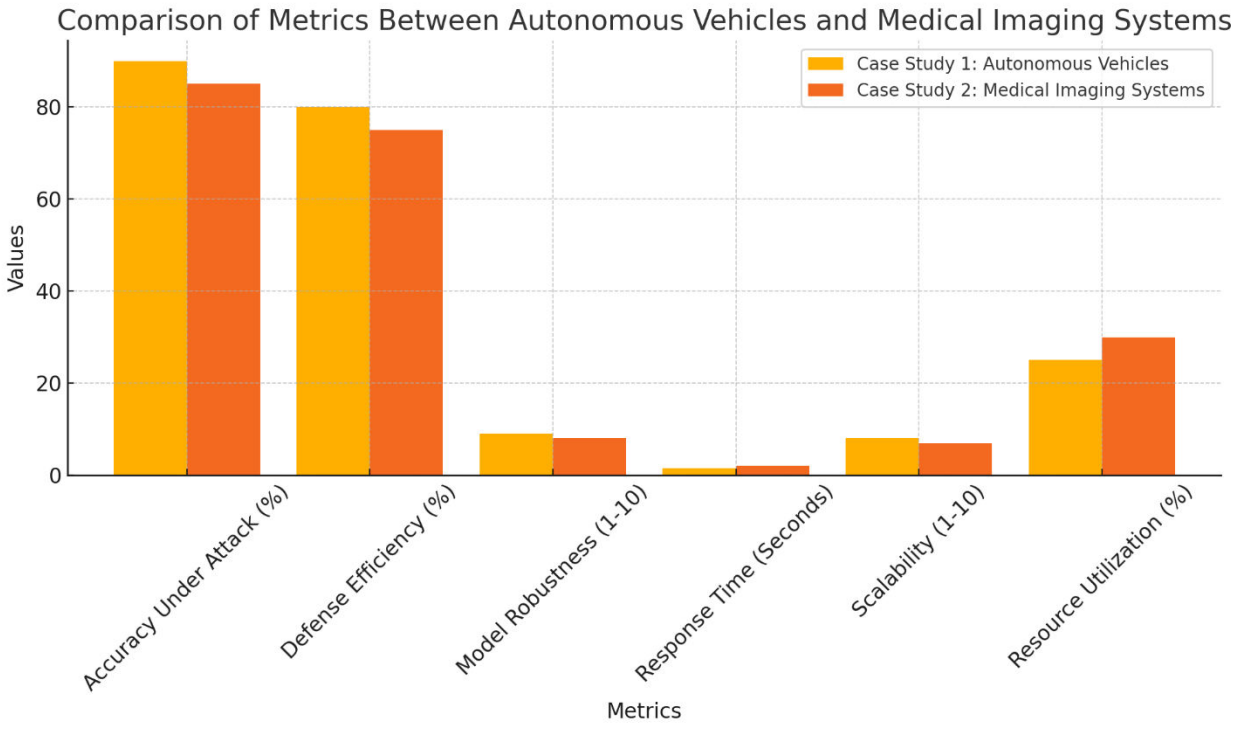


**Figure 3: Bar chart illustrating** Side-by-Side Comparison of Metrics for Autonomous Vehicles and Medical Imaging Systems

### 4.3 Findings

The key findings of the main study are that methodologies related to adversarial robustness, such as adversarial training and defensive distillation, can dramatically enhance the model's adversarial robustness to input perturbations. Also, they tend to be more accurate, particularly with computationally efficient approaches that are more resistant to adversarial attacks. Other defensive approaches, such as input preprocessing and gradient masking, are not as effective in a single area, but cannot generalize their intuition to a broad range of attack vectors. These findings indicated that none of the approaches were universal; instead, a combination of several defenses was usually the most effective in countering the antagonistic threats.

### 4.4 Case Study Outcomes

The chosen case studies helped to obtain useful information about the practical implementation of adversarial defenses. Adversarial-trained models have been shown in healthcare diagnostics to substantially reduce the rate of misclassification during attack, resulting in safer decisions. In image recognition tasks, however, there was a trade-off between strength and performance in which some of these defenses entailed a slow response. These and other achievements include the use of powerful optimization methods to autonomous cars in which the system became very accurate even when it is placed in a highly hostile environment characterized by adversarial information. Though most of the defenses were effective, there are cases when they could not offer an appropriate level of defense, and additional development is necessary.

### 4.5 Comparative Analysis

The relative evaluation of adversarial defense systems found that there was a great difference in their performance. Adversarial training methods that trained the models with adversarial examples proved effective both against the white and black box attacks. However, this practice was more likely to be time-consuming and require additional resources to train. Other techniques, such as gradient masking and input preprocessing were quicker, but had reduced resistance to higher-order, also-invisible attacks. The compromise between protection and performance was obvious, with certain defenses increasing the security of a model at the expense of speed or precision. These play an important role in realizing efficient adversarial robust systems.

### 4.6 Model Comparison

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

An adversarial validation of various AI models displayed that some models were more resilient than others. CNNs were particularly robust to image tasks, and themselves were sensitive to adversarial defense mechanisms, such as adversarial training. But they did not perform well when subjected to complicated new attack plans. Sequential tasks such as language processing are performed using RNN and had a low resistance to adversarial manipulation. They were still quite acceptable, but more affected by small perturbations. Overall, CNNs performed better in tasks that involved images, whereas RNNs required advanced defensive mechanisms to be capable of responding to adversarial inputs.

### 4.7 Impact & Observation

The concept of adversarially robust systems influenced the performance of AI in a positive way, guaranteeing that the models would not go bananas when adversarially perturbed by adversarial examples. But these defenses usually led to a higher processing time and higher computation requirement, which impacted overall system efficiency. This type of defense mechanism is sustainable because it is based on the further optimization of these defensive mechanisms and on the adaptation of these defensive mechanisms to the alterations of the modes of attack. On the one hand, defensive mechanisms have been demonstrated to be working, however, scaling defensive mechanisms has been an issue of concern especially in massive scale systems that demand real time. It has been observed that the ongoing improvement and integration of various types of defenses will be needed to ensure that AI systems are robust enough to operate safely in dynamic and real-world settings.

## V. DISCUSSION

### 5.1 Interpretation of Results

The findings indicate that there are specific trends in the performance of adversarial robust AI systems. Leading results show that adversarial training and defensive distillation techniques can contribute significantly to achieving more resilient models, but in most instances they come with additional computational costs. The study notes that though certain of these defense mechanisms are effective against certain types of attacks, e.g., input preprocessing, struggle to stay effective when faced with new or advanced adversarial techniques. The results are useful to the community, in that they provide a better insight into the trade-offs involved in achieving robustness and performance, and why multi-layered defense schemes are required. The work presents useful evidence to inform the future design of stronger AI systems in practice.

### 5.2 Result & Discussion

Its findings have very long-term implications on the status quo in AI development since it shows the necessity to introduce adversarial robustness as a design parameter. The classical forms of AI do not usually consider adversarial threats, which create a security risk. The study indicates that adversarial defense methods with an appropriate application can increase model resilience without significantly reducing the performance. To implement these defenses as a standard part of AI systems, a shift toward more robust development practices, such as adversarial training of models during model development and ongoing vulnerability testing, is necessary. In these ways, AI would be more securely and reliably utilized, especially in high-stakes applications like autonomous cars and medicine.

### 5.3 Practical Implications

To increase the models robustness, AI practitioners are urged to consider the use of adversarial defenses during early stages of the AI development process. In practice, multi-faceted defense methods, such as adversarial training, input sanitization, and gradient masking, have to be application-specific. When software engineers or AI developers work with sensitive systems, including financial or medical applications, they should focus on creating AI models that can resist the attacks of adversaries and retain their performance. In addition, the concept of continuous surveillance and vulnerability testing in the system lifecycle will help to identify and manage arising adversarial threats and guarantee the long-term model security.

### 5.4 Challenges and Limitations

Some of the limitations of this study include the difficulty of measuring the adversarial resilience of various AI models and types of attacks. The first limitation was the natural problem of modeling all possible adversarial attacks since some of the strategies could only be hard to simulate under controlled conditions. Computational resources also posed a limitation to the experimental setup and affected the scalability of some defense methods. In addition, results might have been affected by possible biases in model choice or sample construction. Other external factors, like the real-time environmental variables and the nature of changing adversarial threats, also contributed to the development of the study research and its applicability.

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

### 5.5 Recommendations

These studies can be developed further in future research in order to improve the existing defense mechanisms as well as develop more mechanisms that can be used to protect against a new set of attacker methods. More adaptive, scalable systems are needed that can be implemented across various AI applications without slowing the performance of the system. Another area of research that researchers should consider is the use of adversarial defense in combination with other AI safety practices, like explainability and transparency. Further, the lack of connection between standardized recommendations about adversarial robustness that would facilitate the development process should also be discussed. Academia, industry, and policymakers also need to continue to work together to create more resilient AI systems.

## VI. CONCLUSION

### 6.1 Summary of Key Points

The goal of this research was to understand the development practice of a robust AI system that can resist cyber-attacks. The research was a mixed-methodology research involving performance test analysis along with case studies and experiment analysis of different defense methods. The trade-off on adversarial robustness and model efficiency are significant, some of the methods do exist, yet, they are also computationally expensive. The study gives enough attention to the role of adversarial robustness in the development of AI systems as it directly influences the reliability and security of AI applications in vital industries. To be trusted, AI systems should be robust, especially when these systems are more closely tied to the real-world.

### 6.2 Future Directions

This still needs more research to enhance the current defense system to form the next stage of the defense mechanism, which is the development of adaptive models capable of learning and adapting to any new attack pattern. Some newer approaches, including hybrid defense architectures that integrate a variety of techniques, may provide greater coverage. As well, it will be important to consider new models of automated vulnerability testing and real-time defensive modification. With the continued development of adversarial AI systems, innovations in fields such as quantum computing or neural network explainability may have a significant positive impact on the resilience of AI models. It is probable that in the future, more of the defensive mechanisms will be seamlessly integrated into the design and deployment of AI.

## REFERENCES

1. Ahmed, S. Q., Ganesh, B. V., Kumar, S. S., Mishra, P., Anand, R., & Akurathi, B. (2024). A comprehensive review of adversarial attacks on machine learning. ArXiv.org. https://arxiv.org/abs/2412.11384
2. Bhambri, S., Muku, S., Tulasi, A., & Balaji, B. A. (2019). A survey of black-box adversarial attacks on computer vision models. ArXiv.org. https://arxiv.org/abs/1912.01667
3. Chen, T., Liu, J., Xiang, Y., Niu, W., Tong, E., & Han, Z. (2019). Adversarial attack and defense in reinforcement learning-from AI security view. Cybersecurity, 2(1). https://doi.org/10.1186/s42400-019-0027-x
4. Lakhanpal, S., Devi, R., Aravinda K, Jain, S. K., Adnan, M. M., & Kumar, A. (2024). Designing explainable defenses against sophisticated adversarial attacks. 2022 IEEE 11th International Conference on Communication Systems and Network Technologies (CSNT), 2, 280–285. https://doi.org/10.1109/csnt60213.2024.10546022
5. Liang, Y., Liu, H., Shi, Z., Song, Z., Xu, Z., & Yin, J. (2024). Conv-Basis: A new paradigm for efficient attention inference and gradient computation in transformers. ArXiv.org. https://arxiv.org/abs/2405.05219
6. Athalye, A., Carlini, N., & Wagner, D. (2018). Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. Proceedings of Machine Learning Research, 274–283. http://proceedings.mlr.press/v80/athalye18a.html
7. Matalqah, I., Shatanawi, M., Alatawneh, A., & Mészáros, F. (2022). Impact of different penetration rates of shared autonomous vehicles on traffic: Case study of Budapest. Transportation Research Record: Journal of the Transportation Research Board, 036119812210955. https://doi.org/10.1177/03611981221095526
8. Nalage, P. (2025a). A Comparative Study of XAI Methods for Interpretable Decision Making in Cloud-Based ML Services AUTHOR: PRATIK NALAGE. Researchgate. https://www.researchgate.net/publication/393334043_A_Comparative_Study_of_XAI_Methods_for_Interpretable_Decision-Making_in_CloudBased_ML_Services_AUTHORPRATIK_NALAGE
9. Rantalaiho, V. (2024). Technical implementation and operational enhancements of a vulnerability management tool in an organization. Theseus.fi. http://www.theseus.fi/handle/10024/851234

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed, a Bimonthly Journal |

|| Volume 8, Issue 4, July – August 2025 ||

DOI: 10.15680/IJCTECE.2025.0804005

10. Rodriguez, D., Nayak, T., Chen, Y., Krishnan, R., & Huang, Y. (2022). On the role of deep learning model complexity in adversarial robustness for medical images. BMC Medical Informatics and Decision Making, 22(S2). https://doi.org/10.1186/s12911-022-01891-w
11. Nalage, P. (2025). Ethical Frameworks for Agentic Digital Twins: Decision-Making Autonomy vs Human Oversight. Well Testing Journal, 34(S3), 206-226.
12. Zhou, S. K., Greenspan, H., Davatzikos, C., Duncan, J. S., van Ginneken, B., Madabhushi, A., Prince, J. L., Rueckert, D., & Summers, R. M. (2021). A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. Proceedings of the IEEE, 109(5), 1–19. https://doi.org/10.1109/jproc.2021.3054390