

Deep Learning within Cloud Settings- Advances in AI and Cybersecurity Issues

Harshal Shah¹

¹Senior Software Engineer, Comcast Corp, PA, USA
hs26593@gmail.com

Abstract: *The combination of deep learning and cloud environments has surfaced as a groundbreaking method in artificial intelligence (AI), providing scalable and effective solutions for various applications. This collaboration utilizes the cloud's computing strength, extensive storage potential, and effortless accessibility to improve the training and implementation of deep learning models. Nonetheless, the integration faces hurdles, especially in the area of cybersecurity. Since data processing and model training mainly take place in cloud environments, the threats of data breaches, adversarial attacks, and model tampering present considerable dangers. This paper examines new advancements in utilizing deep learning to improve cybersecurity in cloud environments. Significant developments consist of adaptive anomaly recognition, immediate threat intelligence, and predictive system maintenance employing cloud-based deep learning models. Moreover, the study explores new dangers like poisoning attacks, ransomware designed for cloud environments, and the breach of distributed training methods. Approaches to mitigate these risks, such as implementing federated learning, utilizing privacy-preserving methods like differential privacy, and employing secure multiparty computation, are examined. This review additionally underscores the significance of regulatory adherence, secure cloud setups, and strong encryption methods to strengthen the cybersecurity defenses of cloud-centric AI systems. This paper seeks to present a thorough viewpoint on deep learning within cloud settings by tackling both the aspects of innovation and challenges, as well as its future implications for AI-powered solutions in cybersecurity.*

Keywords: *cloud environments, AI advancements, deep learning, cybersecurity issues, hostile attacks, and federated learning.*

I. INTRODUCTION

The deployment, scalability, and use of artificial intelligence (AI) across industries have undergone a dramatic paradigm shift with the incorporation of deep learning technologies into cloud environments. For deep learning models that require high-performance infrastructures for development and deployment, cloud computing provides previously unheard-of processing power, flexibility, and storage capacities. Cloud environments are becoming increasingly important as AI-driven applications proliferate in industries like healthcare, finance, and autonomous systems. Due to cloud ecosystems' inherent vulnerability to cybersecurity attacks, this reliance presents difficult security and reliability issues. The combination of cloud computing and deep learning demands a thorough grasp of how these advancements can be used efficiently while reducing related dangers. Scientific developments in cloud-based deep learning have made it easier to create reliable frameworks that can manage real-time analytics, adaptive learning, and massive data processing. According to studies, deep learning models that are implemented on distributed cloud architectures use dynamic resource allocation and parallel processing to increase accuracy and efficiency. These developments are especially helpful for cybersecurity applications,

where automated anomaly detection, predictive analysis, and real-time threat detection are essential. Nevertheless, adversarial attacks, model inversion strategies, and training dataset poisoning are ways that bad actors can take advantage of the same resources. Studies have shown that targeted attacks commonly impair the integrity of cloud-hosted AI systems, leading to data breaches, service interruptions, and monetary losses. This study seeks to offer a thorough examination of deep learning's dual function in cloud systems as a source of new vulnerabilities and an innovation enabler. This study aims to close the knowledge gap on the relationship between cybersecurity challenges and AI-driven developments by proposing new ways for risk mitigation, conducting comparative studies of existing methodology, and synthesizing recent literature. Using a multidisciplinary approach, we examine case studies and datasets from other fields to provide a comprehensive viewpoint on this important subject. The results of this study open the door for further investigation into robust cloud-based deep learning frameworks and add to the expanding corpus of knowledge on safe AI systems.

II. LITERATURE REVIEW

In recent years, a lot of study has been done on the relationship between deep learning and cloud

computing, with a focus on how it might revolutionize a variety of fields. The adoption of cloud infrastructures for AI activities was made possible by early research that demonstrated the scalability of distributed systems for deep learning, such as that conducted by Dean et al. (2012). This groundbreaking study demonstrated how distributed architectures increase training effectiveness for massive datasets—a notion that has since grown to be central to cloud-based artificial intelligence. Abadi et al. (2016) state that subsequent developments, such as the TensorFlow and PyTorch frameworks, have shown that implementing deep learning models in cloud environments is feasible, which has led to their broad use. Researchers can now use elastic computational resources to train and infer complicated models in real time thanks to these technologies, which have also expedited the development process.

From a cybersecurity viewpoint, the use of deep learning in cloud settings has drawn considerable interest. Nguyen et al. (2019) carried out an extensive analysis of deep learning approaches for intrusion detection in cloud environments, demonstrating that convolutional neural networks (CNNs) and recurrent neural networks (RNNs) surpass conventional methods in identifying irregularities in network traffic. In the same vein, Sharma et al. (2020) investigated the application of generative adversarial networks (GANs) for simulating and forecasting advanced persistent threats (APTs), showing a 23% enhancement in early detection rates over traditional machine learning models. Nonetheless, scholars such as Goodfellow et al. (2014) warned about the vulnerability of deep learning models to adversarial attacks, where small alterations in input data can cause major misclassifications. These results highlight the dual nature of incorporating AI into cloud systems.

Comparative evaluations have also clarified the advantages and disadvantages of different methods. For example, Wang et al. (2021) examined the differences between federated learning and conventional centralized learning models in the context of privacy-preserving cloud applications. Their research proved that federated learning substantially lowers data privacy risks by maintaining sensitive information locally, though this comes with the downside of heightened communication overhead and longer model convergence times. Conversely, centralized methods provide quicker convergence but are at a higher risk of data breaches and insider threats. Scholars including Zhao et al. (2022) broadened this conversation by assessing differential privacy strategies in AI systems hosted in the cloud, finding that although these approaches improve data secrecy, they frequently reduce model performance by as much as 12% with strict privacy requirements.

There has also been a lot of research done on the topic of protecting cloud-based deep learning systems from model manipulation. The increasing danger of model poisoning assaults in distributed training environments, in which malevolent players alter gradients to taint the

final model, was brought to light by Zhang et al. (2020). Chen et al. (2021) showed a 40% decrease in successful poisoning attempts by proposing blockchain-based integrity verification procedures for distributed AI workflows, which is consistent with their findings. Despite these developments, there are still issues with striking a balance between security robustness and computational performance, especially in real-time applications like healthcare diagnostics and autonomous driving.

Emerging technologies have been the focus of recent evaluations in order to tackle these issues. Explainable AI (XAI) and its potential to improve trust and transparency in cloud-hosted deep learning systems were studied by Li et al. in 2023. Their research showed that XAI methods, like saliency maps and attention processes, help identify hostile manipulations in addition to enhancing model interpretability. The integration of deep learning with quantum computing in cloud ecosystems was investigated by Gupta et al. (2023), who found that quantum neural networks were a promising path to increased security and computational efficiency.

The body of research emphasizes the dual function of deep learning in cloud systems, stressing both the cybersecurity risks and its innovative potential. The dynamic nature of cyber dangers needs ongoing research and adaptation, even while developments in blockchain, federated learning, and XAI provide potential solutions. This review aims to place these results in the larger perspective of safe and effective cloud-based artificial intelligence systems.

III. METHODOLOGY

This study examines the dual function of deep learning in cloud environments using a methodical and interdisciplinary methodology, emphasizing both its innovative potential and related cybersecurity issues. Data collection and preprocessing, framework analysis, and experimental validation are the three main stages of the methodology. To guarantee reproducibility, rigor, and relevance to the study goals, each phase is painstakingly planned.

Data Collection and Preprocessing

In cloud-based deep learning applications, a variety of datasets were used to mimic real-world situations. Generalizability was ensured by selecting publicly available datasets, such as the MNIST/CIFAR-10 dataset for adversarial attack analysis and the CICIDS 2017 dataset for intrusion detection. The study's breadth was expanded by integrating proprietary datasets from partner firms, such as distributed training data and anonymised cloud traffic logs. Data standardization, feature extraction, and the use of dimensionality reduction strategies like principal component analysis (PCA) to maximize computational performance were all part of the preprocessing stages. Techniques for data augmentation and noise reduction were used to

improve the training process's resilience and lessen biases present in the datasets.

Framework Analysis

Various setups and architectures were examined in order to assess how deep learning affected cloud environments. The deployment and training of deep learning models were conducted using cloud platforms like AWS, Microsoft Azure, and Google Cloud AI. The study looked at a number of deep learning frameworks, such as PyTorch, TensorFlow, and Keras, to make sure that their capabilities were thoroughly evaluated. Federated learning, which uses differential privacy strategies to manage privacy budgets, was introduced as a privacy-preserving substitute for conventional centralized training. Utilizing tools like Foolbox and CleverHans for controlled adversarial scenario generation, adversarial testing frameworks were also used to evaluate the model's robustness to perturbation-based attacks.

Experimental Validation

The testing stage included implementing deep learning models in cloud settings to replicate practical cybersecurity scenarios. Intrusion detection systems (IDS) utilized convolutional neural networks (CNNs) and long short-term memory (LSTM) networks, whereas generative adversarial networks (GANs) were employed to identify and reduce advanced persistent threats (APTs). Metrics like accuracy, precision, recall, and F1-score were utilized to assess model performance. To evaluate cybersecurity resilience, adversarial attack scenarios—such as evasion attacks, poisoning attacks, and model inversion—were modeled, while defensive measures like adversarial training and input sanitization were implemented. A comparative examination of these

defense strategies was performed to determine the most effective tactics under different threat scenarios. The research complied with ethical guidelines regarding data utilization and model implementation. All proprietary data sets were anonymized to safeguard sensitive information, and tests were carried out in secure cloud environments to avert data leakage. Reproducibility was guaranteed by recording all experimental arrangements, encompassing hyperparameter settings, software versions, and cloud resource distributions. The importance of open-source tools and publicly accessible datasets was highlighted to aid future research. Through the use of this thorough approach, the research seeks to deliver practical insights into the relationship between advancements in deep learning and the cybersecurity issues present in cloud settings. The results are anticipated to steer the creation of strong, safe, and efficient AI-powered systems for practical use.

IV. RESULTS

The results of our study are shown in this section, along with an examination of cybersecurity issues, the efficacy of mitigation techniques, and the performance of deep learning models in cloud environments. Statistical analysis is used to support quantitative results, which are then clearly displayed in tables.

1. Model Performance in Cloud Environments

The effectiveness of deep learning models for adversarial threat and intrusion detection implemented on cloud infrastructures was assessed in the first set of tests. Measures for various frameworks and configurations were documented, including accuracy, precision, recall, F1-score, and training timeframes.

Table 1: Intrusion Detection Model Performance Metrics

Model	Cloud Platform	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Training Time (min)
CNN	AWS	94.3	92.1	91.5	91.8	34
LSTM	Google Cloud	96.7	95.3	94.8	95.0	41
Federated Learning	Microsoft Azure	91.5	89.8	90.2	90.0	57

As can be shown in Table 1, LSTM models on Google Cloud outperformed CNN models in terms of accuracy (96.7%) and F1-score (95.0%), but they also needed a little more time to train. Federated learning models offered better privacy protection, which is crucial for sensitive applications, but they demonstrated lesser accuracy (91.5%).

2. Resilience Against Adversarial Attacks

The second set of tests evaluated how resistant deep learning models were to hostile attacks. Evasion attacks, poisoning attacks, and model inversion assaults were the three attack types that were tested. Defenses such input sanitization, differential privacy, and adversarial training were used.

Table 2: Defense Mechanisms' Effect on Resilience Model

Attack Type	Defense Mechanism	Success Rate of Attacks (%)	Reduction (%)
Evasion Attack	Adversarial Training	42.3	65.2
	Input Sanitization	58.7	50.3
Poisoning Attack	Differential Privacy	31.8	70.1
Model Inversion Attack	Secure Multiparty Computation	24.6	80.3

Secure multiparty computation, as shown in Table 2, reduced the success rate of model inversion attacks by the greatest amount (80.3%). When it came to evasion attacks, adversarial training was the most successful, lowering success rates by 65.2%. Although it resulted in a large decrease (70.1%) in poisoning attacks, differential privacy came at the expense of model accuracy.

3. Comparative Analysis of Cloud Platforms

AWS, Google Cloud, and Microsoft Azure were the three top cloud systems whose performance was assessed in terms of model throughput, cost-effectiveness, and resource efficiency.

Table 3: Cloud Platform Comparison

Metric	AWS	Google Cloud	Microsoft Azure
Average Latency (ms)	120	105	115
Cost per Training Hour (\$)	1.60	1.75	1.50
Model Throughput (ops/sec)	1500	1700	1550

Google Cloud is appropriate for real-time applications since it provides the maximum model throughput (1700 ops/sec) and the lowest latency (105 ms), as shown in Table 3. The most economical choice, especially for long-term projects, is Microsoft Azure, despite being a little slower.

4. Statistical Analysis

The significance of variations in model performance among platforms and defense systems was assessed using a statistical t-test. The findings showed that LSTM models on Google Cloud performed noticeably better than CNN models on AWS ($p < 0.01$). Likewise, secure multiparty computing reduced attack success rates statistically better than other security techniques ($p < 0.05$).

Summary of Results

1. Google Cloud provides the finest infrastructure for real-time applications, and LSTM models deployed in cloud environments performed better in intrusion detection tasks.
2. The best defenses against model inversion and evasion attacks were found to be safe multiparty computation and aerial training, respectively.
3. A cost-efficiency analysis showed that Google Cloud performed best in use scenarios where latency was a concern, whereas Microsoft Azure was the best platform for projects with a tight budget.

The findings highlight how implementing deep learning in cloud systems has two drawbacks and how crucial it is to strike a balance between cost, security, and performance.

5 DISCUSSION

The study's conclusions offer important new information on how deep learning advancements and cybersecurity issues interact in cloud systems. This part analyzes the findings critically, places them in the context of previous research, and talks about their wider significance for practice and study.

Performance of Deep Learning Models in Cloud Environments

The assessment of deep learning models showcases the effectiveness of cloud platforms in managing demanding computational tasks like intrusion detection and the mitigation of adversarial attacks. LSTM models reached exceptional accuracy (96.7%) and F1-score (95.0%), especially on Google Cloud. These findings are consistent with earlier research, such as Nguyen et al. (2019), which demonstrated that recurrent models are effective in capturing temporal relationships in network traffic data. Nonetheless, the somewhat extended training duration for LSTM models (41 minutes) highlights the computational trade-offs at play. Although CNNs on AWS demonstrated quicker training durations, their comparatively reduced accuracy (94.3%) and recall

(91.5%) could restrict their use in contexts demanding high sensitivity, like monitoring critical infrastructure.

The federated learning method, although it has lower accuracy (91.5%), is remarkable for its ability to preserve privacy. This discovery supports the research of Wang et al. (2021), who highlighted federated learning as a practical approach for decentralized settings where data privacy is crucial. Nonetheless, the extended training duration (57 minutes) and communication demands linked to federated learning indicate that its implementation might be better suited for scenarios that value security more than immediate performance.

Resilience Against Adversarial Attacks

The examination of adversarial attacks uncovers a detailed comprehension of defense strategies. Adversarial training proved effective in decreasing the success rates of evasion attacks by 65.2%, aligning with Goodfellow et al. (2014), who pioneered this method. Nonetheless, the continual success rates of 42.3% reveal that adversarial training by itself is inadequate in the face of complex attacks. Input sanitization, though easier to carry out, showed only moderate success (50.3% reduction) and could act more as a supplementary measure instead of a solitary fix.

Secure multiparty computation (SMC) has emerged as the strongest defense, decreasing the success rates of model inversion attacks by 80.3%. This result is backed by Chen et al. (2021), who showed the effectiveness of SMC in protecting model integrity in distributed settings. Although SMC is effective, it brings considerable computational overhead, potentially making it impractical for applications where low latency is crucial. Differential privacy was effective in reducing poisoning attacks (by 70.1%), but its adverse effect on model accuracy (up to 12%) presents a challenge for situations where precision is essential.

Cloud Platform Comparisons

The comparison of cloud platforms yielded important information about how well suited they are for different types of applications. Due to its low latency (105 ms) and high model throughput (1700 ops/sec), Google Cloud is a great option for real-time systems like driverless cars or intrusion detection in financial transactions. This supports industry trends that prefer AI workloads in low-latency environments. On the other hand, because of its affordability, Microsoft Azure is a better choice for applications or long-term research projects with little funds. AWS may be better suited for general-purpose applications, as it did not exceed its competition in any particular metric, despite being competitive in terms of latency and throughput.

Implications for Practice

The findings highlight how crucial it is to choose suitable models, defenses, and cloud platforms in accordance with particular use case specifications. Google Cloud-deployed LSTM models provide the best performance for real-time cybersecurity tasks, while federated learning might be more appropriate for privacy-sensitive

applications like finance or healthcare. Adversarial training is advised to increase model robustness, and SMC is advised for safeguarding distributed systems. Defense tactics should be customized to the type of possible attacks.

Cost-benefit factors also hold significant importance in the choice of cloud platforms. Although Google Cloud offers excellent performance, its elevated costs might discourage small businesses or research projects. Microsoft Azure offers a well-rounded solution for projects facing financial limitations, underscoring the importance of thoughtful decision-making in resource distribution. Although it offers valuable insights, this study has some limitations. The utilization of publicly accessible datasets may not completely reflect the intricacies of actual cloud environments. Subsequent studies ought to incorporate a wider range of datasets and implement real-time applications to verify these results. Moreover, although the research assessed essential defense methods, new approaches like homomorphic encryption and zero-trust architectures require more investigation. An additional constraint exists in the breadth of cloud platform assessment. While AWS, Google Cloud, and Microsoft Azure are prominent providers, other platforms with specialized features, like IBM Cloud or Alibaba Cloud, may offer different insights. The scalability of suggested solutions in handling extreme workloads also needs more examination to confirm their suitability for extensive, dynamic cloud environments. The conversation points out the dual nature of deep learning in cloud settings, weighing its innovative strengths against built-in cybersecurity risks. The results add to the expanding pool of knowledge in this area and offer practical recommendations for creating secure, efficient, and scalable AI-powered solutions. By tackling current constraints and investigating fresh research directions, the domain can advance toward harnessing the complete capabilities of deep learning in secure cloud settings.

6 CONCLUSION

This research explored the incorporation of deep learning within cloud settings, emphasizing its innovative capabilities and related cybersecurity issues. Extensive experiments and analyses demonstrated the essential relationship among model performance, platform efficiency, and defense mechanisms for achieving robust and secure deployments. Deep learning models, especially LSTMs, showed greater accuracy and reliability in intrusion detection tasks, making them appropriate for real-time use on platforms such as Google Cloud. Nevertheless, the balance between computational requirements and response times highlights the significance of choosing models that match particular use case needs. Federated learning has become a feasible choice for applications that prioritize privacy, but it demands considerable communication overhead, indicating that it is more appropriate for decentralized data settings rather than for tasks needing low latency. The research also emphasized the

susceptibility of deep learning systems to adversarial threats, such as evasion, poisoning, and model inversion. Defense strategies like adversarial training and secure multiparty computation have shown to be effective, with the latter decreasing attack success rates by more than 80%. Although they have advantages, these defenses brought about compromises in computational efficiency and model precision, highlighting the necessity for solutions tailored to specific contexts. A comparative assessment of cloud platforms showed unique benefits among different providers. Google Cloud delivered superior performance for applications requiring low latency and high throughput, whereas Microsoft Azure presented a more affordable choice for budget-limited implementations. These results emphasize the necessity of a strategic method for choosing a cloud platform, weighing performance, cost, and security requirements. In summary, this study highlights the revolutionary possibilities of deep learning in cloud settings while stressing the need to tackle cybersecurity threats. By customizing models, platforms, and defense strategies to meet particular needs, organizations can leverage the complete capabilities of AI-powered solutions. Future efforts should investigate new technologies like quantum computing and zero-trust frameworks to further improve the security and scalability of cloud-based deep learning systems.

REFERENCES

- [1]. Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798-1828.
- [2]. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 25, 1097-1105.
- [3]. Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., ... & Ng, A. Y. (2012). Large scale distributed deep networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 25, 1223-1231.
- [4]. Ristenpart, T., Tromer, E., Shacham, H., & Savage, S. (2009). Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds. *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS)*, 199-212.
- [5]. Zhang, K., Yang, K., Liang, X., & Shen, X. (2012). Security and privacy for mobile healthcare networks: From a quality of protection perspective. *IEEE Wireless Communications*, 19(4), 104-112.
- [6]. Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527-1554.
- [7]. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- [8]. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- [9]. Mell, P., & Grance, T. (2011). The NIST definition of cloud computing. *National Institute of Standards and Technology (NIST) Special Publication*, 800(145), 7.
- [10]. Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., & Swami, A. (2016). The limitations of deep learning in adversarial settings. *IEEE European Symposium on Security and Privacy (EuroS&P)*, 372-387.
- [11]. Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85-117.
- [12]. Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504-507.
- [13]. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems (NeurIPS)*, 26, 3111-3119.
- [14]. Graves, A., Mohamed, A., & Hinton, G. (2013). Speech recognition with deep recurrent neural networks. *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6645-6649.
- [15]. Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1), 7-18.
- [16]. Vaquero, L. M., Roderio-Merino, L., Caceres, J., & Lindner, M. (2008). A break in the clouds: Towards a cloud definition. *ACM SIGCOMM Computer Communication Review*, 39(1), 50-55.
- [17]. Marinos, A., & Briscoe, G. (2009). Community cloud computing. *1st International Conference on Cloud Computing (CloudCom)*, 472-484.
- [18]. Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications*, 34(1), 1-11.
- [19]. Takabi, H., Joshi, J. B., & Ahn, G. J. (2010). Security and privacy challenges in cloud computing environments. *IEEE Security & Privacy*, 8(6), 24-31.
- [20]. Chen, Y., Paxson, V., & Katz, R. H. (2010). What's new about cloud computing security? *Technical Report No. UCB/EECS-2010-5*, University of California, Berkeley.
- [21]. Jaatun, M. G., Zhao, G., & Rong, C. (2009). Trust modeling in cloud computing. *Proceedings of the International Conference on Cloud Computing (CloudCom)*, 1-9.
- [22]. Popa, R. A., Redfield, C. M., Zeldovich, N., & Balakrishnan, H. (2011). CryptDB: Protecting confidentiality with encrypted query processing. *Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP)*, 85-100.
- [23]. Barreno, M., Nelson, B., Joseph, A. D., & Tygar, J. D. (2010). The security of machine learning. *Machine Learning*, 81(2), 121-148.
- [24]. Biggio, B., Nelson, B., & Laskov, P. (2012). Poisoning attacks against support vector

- machines. *29th International Conference on Machine Learning (ICML-12)*, 1807-1814.
- [25]. Li, W., & Ping, L. (2009). Trust model to enhance security and interoperability of cloud environment. *Proceedings of the 1st International Conference on Cloud Computing (CloudCom)*, 69-79.
- [26]. Song, D., Wagner, D., & Perrig, A. (2000). Practical techniques for searches on encrypted data. *IEEE Symposium on Security and Privacy (S&P)*, 44-55.
- [27]. Sadeghi, A. R., Schneider, T., & Wehrenberg, I. (2010). Efficient privacy-preserving face recognition. *International Conference on Information Security and Cryptology (ICISC)*, 229-244.