The use of AI and ML to improve the resilience of cloud systems.

[1]Solomon Airemen, [2]Ibukunoluwa Olushiyan, [3]Oluwaseyi O. Alabi, 0009-0005-0027-5930

[1]Department of Operations and Production Management, University of Lagos, Nigeria
[2]Department of Electrical/Electronics and Computer Engineering, University of Lagos, Nigeria
Department of Mechanical Engineering Lead City University Ibadan, Nigeria
*Corresponding author Email: alabi.oluwaseyi@lcu.edu.ng

Abstract - The growing complexity and dynamic nature of cloud systems pose significant challenges to their resilience. This paper explores the application of Artificial Intelligence (AI) and Machine Learning (ML) techniques to enhance the resilience of cloud systems. Through a comprehensive review of existing literature, we delve into various aspects where AI and ML play pivotal roles, including storage utilization tracking, and predictive maintenance. Notable studies are discussed, such as the integration of regression and neural network techniques for workload forecasting and the utilization of diverse machine learning models for resource consumption prediction. However, cloud service providers (CSPs) need to meet the quality of service (QoS) requirements of their customers, and doing so in a cost-effective manner can be challenging. The workloads of virtual machines (VMs) are often dynamic, with fluctuations that make it difficult to accurately predict future resource requirements. The paper highlights the potential of AI to improve the resilience of cloud systems by detecting and preventing failures, optimizing resource utilization, and improving overall performance. Insights gleaned from these advancements not only contribute to the academic discourse on cloud system resilience but also provide practical implications for architects, administrators, and developers seeking to fortify cloud infrastructure against emerging challenges. This exploration underscores the transformative impact of AI and ML in fortifying the resilience of cloud systems, making them more adaptive, efficient, and capable of withstanding evolving operational landscapes.

Keywords - Artificial intelligence, resilience system, machine learning, deep learning, big data, Cloud system

I. INTRODUCTION

Cloud computing is a highly popular IT service that has evolved to offer a wide range of capabilities under a single platform. A single cloud service provider can now offer a comprehensive set of services, including web hosting, computation, storage, and more. This consolidated approach provides businesses with a single point of contact for all their IT needs, streamlining the process of acquiring and managing IT resources. In addition, this consolidation allows for greater economies of scale, improving the cost-effectiveness of cloud services. The adaptability of cloud computing, facilitated by service models such as pay-as-you-go, has prompted numerous organizations to transition their operations to the cloud [1]. The elasticity inherent in the service model enables most organizations to realize cost savings, as they are spared the need to construct, oversee, and expand an extensive private infrastructure. In contemporary scenarios, the significance of cloud computing lies in its capacity to allow users to pay for precisely the cloud services they require, precisely when needed. Cloud computing offers the advantage of being able to scale up or down as needed, allowing clients to only pay for the resources they use. This flexibility is valuable because it addresses the common problem of overprovisioning, where clients may pay for resources, they don't need, or under provisioning, where clients may not have enough resources to meet demand. By being able to scale on demand, clients can be assured that they are only paying for what they need, which is important in an era of rapidly changing market conditions and consumer demands [2].

Cloud service providers leverage hardware virtualization to operate multiple virtual machines (VMs) or workloads on a single physical machine (PM) with diverse resource allocations. In the cloud environment, numerous applications find their hosting on these virtual machines. However, the risk of a PM becoming overloaded arises when the resource demands from its VMs surpass the available resources, especially considering the dynamic variations in each VM's load over time. The repercussions of load imbalance within a PM extend to all the VMs and applications operating on that particular machine. A critical aspect of this scenario involves the adherence to Service Level Agreements (SLAs) [3], where any shortfall in resource allocation to customer applications constitutes a breach. An SLA acts as a contractual agreement between a cloud customer and the cloud service provider, ensuring the optimal performance of the customer's applications. To uphold the SLA, a cloud service provider must guarantee that each virtual machine receives the necessary resources and actively prevent PM overload. To achieve this, a prediction model comes into play, retrieving historical data from the resource manager and providing anticipated workload projections for each task, as illustrated in Figure 1. Subsequently, the resource management system allocates each virtual machine to a specific PM based on these anticipated outcomes. This proactive approach helps maintain the delicate balance of resources, ensuring the smooth operation of virtual machines and compliance with SLAs [4].

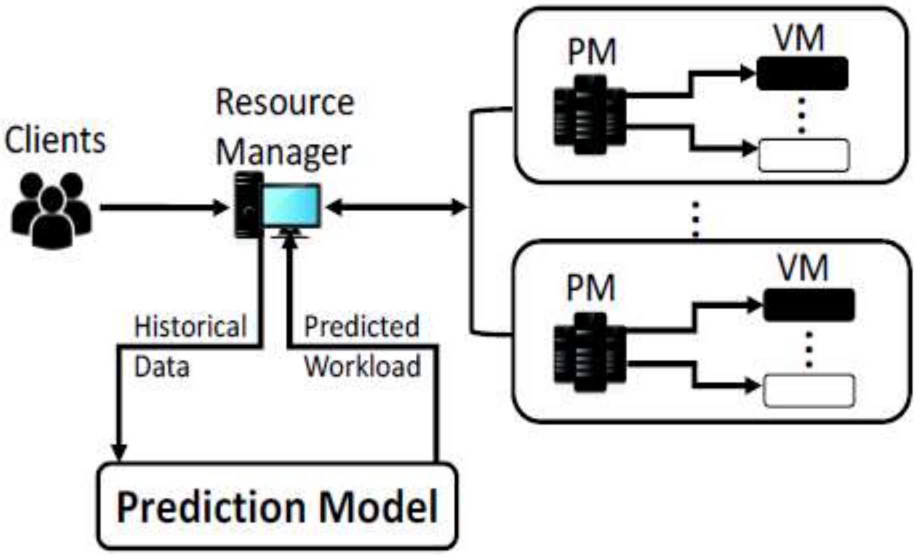

Figure 1. Prediction of workload procedure systems.

In cloud data centers, resource overbooking is a common practice, where CPU, bandwidth, and other resources are shared by multiple tenants. While this allows for efficient use of resources, it also means that those resources are frequently overutilized, leading to poor application performance. This is particularly true when virtual machines with high resource requirements are hosted on the same physical machine, competing for the limited resources available. This issue is further compounded by the fact that VMs are often provisioned with more resources than they need, further reducing the overall efficiency of the system. One of the most challenging aspects of resource provisioning in cloud environments is addressing the problems that arise during both the initial VM allocation and VM migration stages. These issues can be quite resource-intensive, requiring a significant amount of time and effort to resolve. For example, when allocating a new VM, the system must carefully consider the VM's resource requirements and the available capacity in the underlying physical machine. Similarly, when migrating a VM between physical machines, the system must carefully balance the VM's resource needs with the available capacity on the target machine. Recent research efforts [4][5] have focused on the development of techniques for rapidly forecasting virtual machine resource demand. The goal of these techniques is to achieve load balancing or to ensure that adequate resources are provisioned for each VM. In particular, proactive load balancing involves monitoring the resource demands of the VMs hosted on a physical machine (PM), and making adjustments as needed to avoid the PM becoming overwhelmed. This can be accomplished by migrating VMs with high resource requirements to other PMs, or by scaling the resources of the PM itself. While several different approaches have been proposed for resource demand forecasting, there has been little work done to compare the relative performance of these approaches [6]. This is surprising given the importance of this task for the efficient operation of cloud environments. To date, the literature has explored several approaches, including statistical methods [7], machine learning technique, and deep learning methods [6] [7]. Imam et al. (2011) introduced regression and time delay neural network (NN) techniques for predicting the workload of individual virtual machines. In a similar vein, Farahnakian et al. (2013) employed neural networks (NN) and linear regression to forecast the forthcoming demands of virtual machines (VMs) while developing resource monitoring and supply algorithms. Bankole et al. (2015) utilized three distinct machine learning models support vector regression, neural networks, and linear regression to craft a cloud client prediction model, foreseeing the resource consumption of each virtual machine. In a study by Islam et al. (2012), a methodology was presented to predict the future CPU load of a virtual machine, utilizing both Neural Network and Linear Regression techniques. The findings indicated that the Neural Network outperformed Linear Regression in terms of accuracy. Additionally, the study demonstrated that the accuracy of both algorithms is influenced by the size of the input window [3].

There is a greater need for artificial intelligence (AI) at the level of traditional robotic action and reaction due to the application of cloud computing in corporations and organizations. Therefore, the potential aim of this work is to explore the use of AI and machine learning to optimize the performance and resource utilization of cloud systems. This could include developing techniques to dynamically allocate resources based on predicted demand, or to identify and eliminate redundant or inefficient processes.

II. Methodology

Machine Learning

As a subset of artificial intelligence (AI), machine learning (ML) encompasses methods that enable computers to deduce information from data and generate AI applications (Fig. 2). Over the past few decades, machine learning has evolved significantly, diverging from its historical form due to advancements in computing technologies, as noted by SAS. The notion of computers learning without explicit programming and the concept of pattern recognition laid the foundation for artificial intelligence, prompting researchers to explore whether computers could learn from data. The iterative nature of machine learning is pivotal, as models can autonomously evolve when exposed to new data [6]. This characteristic ensures the development of reliable, reproducible judgments and outcomes, drawing on insights from prior computations. Although not a new science, machine learning has recently gained popularity. In essence, the machine learning process involves collecting datasets and utilizing them to build statistical models through algorithms, with the assumption that these models will be applied to address real-world issues. Advancements in AI technology have positioned machine learning as a critical component in various commercial applications, spanning industries such as e-commerce, banking, medicine, and research projects. Contrary to the misconception that it is exclusive to large enterprises with extensive research teams, machine learning is accessible to software engineers in the field of artificial intelligence, particularly those proficient in programming languages like Python. There are comprehensive resources available, such as a notable book focused on Python driving machine learning, offering valuable techniques for creating custom ML solutions. In the era of Big Data (BD), where vast amounts of data are readily available, the potential applications of machine learning are limitless, with creativity being the sole constraint [7]. Machine learning systems can be categorized into three types based on whether they are trained with human input, possess incremental learning capabilities, and employ Deep Learning to augment data for pattern identification in fresh data points

Machine learning (ML) is a multifaceted field that falls into four main categories: 1) Supervised; 2) Semi-Supervised; 3) Unsupervised; and 4) Reinforced, regardless of whether the models were trained on humans or not [8]. The following is a brief description of each of these types:

- *Supervised*

Machine learning algorithms that automate decision-making by extrapolating from well-known examples have shown to be the most effective. The user gives the algorithm pairs of inputs and intended outputs in this situation, which is referred to as supervised learning. The algorithm then figures out how to create the desired output given an input. Specifically, without assistance from a human, the algorithm can produce an output for an input it has never seen before.

- *Semi-supervised*.

In semi-supervised learning datasets, there exist both labeled and unlabeled samples, with the latter usually outnumbering the identified ones significantly. The objective of both supervised and semi-supervised learning algorithms is the same. By leveraging a substantial number of unlabeled instances, the learning algorithm aims to discover, or perhaps more precisely, "create" or "calculate," a more robust model

- *Unsupervised*

Unsupervised learning, the third category of machine learning algorithms explored here, can be conceptualized as a subset of machine learning in general. This category functions independently of teachers or trainers to guide the learning algorithm and lacks a predetermined outcome. In unsupervised learning, the input data is presented to the learning algorithm, and it is tasked with utilizing its information processing capabilities to derive knowledge, as illustrated in Figure 3 [9]

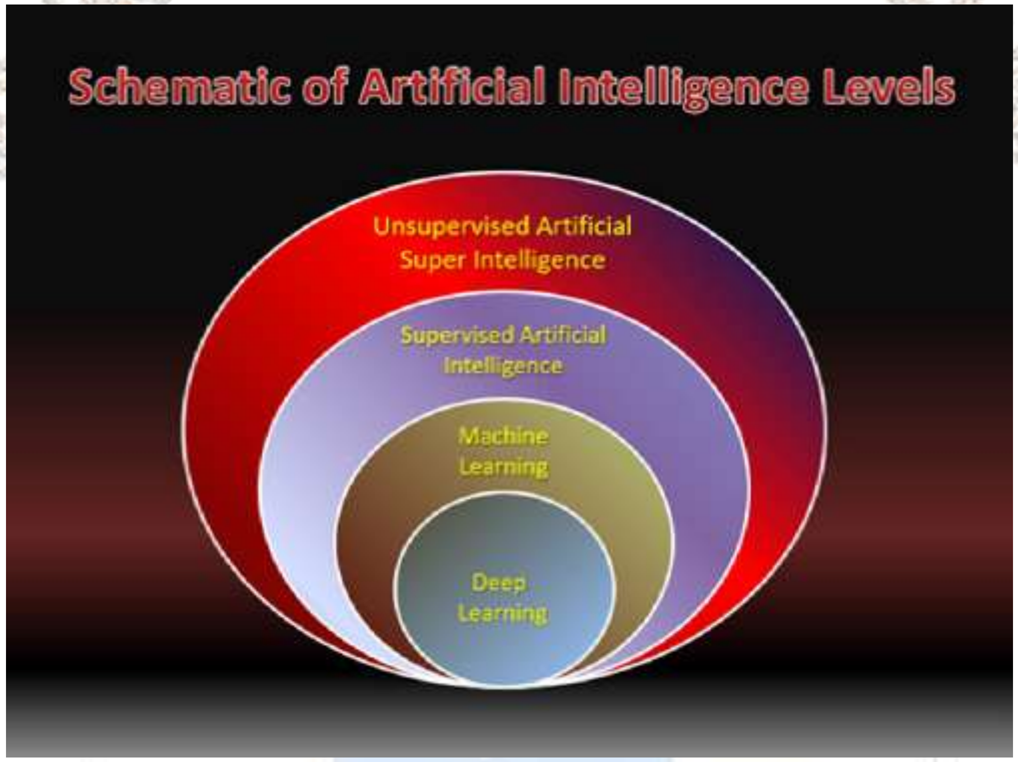

Figure 2. Relationship of Artificial Intelligence, Machine Learning, Deep learning

- Reinforcement type

Reinforcement learning, a subset of machine learning, enables a computer to navigate and interpret its environment, treating the current state as a vector of features. The system is capable of functioning across all states, with diverse actions having the potential to alter the machine's surroundings while offering varying rewards [10]. The primary objective of a reinforcement learning algorithm is to learn a policy a function that takes a state's feature vector as input and determines the optimal course of action for that state, akin to a model in supervised learning. An optimal course of action maximizes the predicted average payoff. Reinforcement learning is particularly suited for addressing specific classes of problems, such as those found in gaming, robotics, resource management, or logistics, where decisions must be made sequentially with a long-term objective in mind.

Deep Learning

As illustrated in Fig. 2, deep learning (DL) is a broader field encompassing machine learning, designed to enable computers to tackle more complex tasks. It achieves this by learning from experience and perceiving the world as a hierarchy of concepts, each defined by simpler notions [10]. Another perspective on deep learning involves comprehending multiple levels of abstraction and representation, facilitating the interpretation of various types of data, including text, sound, and images. In essence, deep learning involves the application of neural networks characterized by higher numbers of neurons, layers, and interconnections. While we are still distant from fully replicating the intricacies of the human brain, notable progress has been made in that direction. Macroeconomic digital transformation, or DX, is becoming more widespread. Artificial intelligence (AI), machine learning (ML), and continuous deep learning (DL)-based intelligent applications are the next big thing in technology that are revolutionizing the way businesses and consumers work, learn, and play. The new digital economy is centered around data, but it also involves managing data from the edge to the core to the cloud, sensing the environment, analyzing it in almost real-time, taking action based on what you learn, and influencing outcomes [11].

Mobile devices, big data, artificial intelligence (AI), machine learning (ML), and deep learning (DL) collaborate seamlessly to perceive and continuously adapt to the environment. The capacity to leverage this integration for delivering insightful, value-added predictions and actions is a distinguishing factor for successful organizations [12]. The business objectives of AI strike a balance between tactical and strategic goals, encompassing activities such as optimizing revenue from existing products, introducing innovative digital revenue streams, and enhancing operational efficiencies to bolster competitive distinctiveness.

Propelled by Artificial Neural Networks (ANNs), Super Artificial Intelligences (SAIs) may eventually forge a future landscape that enhances their human counterparts, possibly surpassing them to become Homo-Fabians and Homo-Sapiens not too distant from the present [13]. A pivotal tool in machine learning is the artificial neural network (ANN), a computational model inspired by the structure and functions of biological neural networks. ANNs represent artificial intelligence (AI) systems meticulously designed to emulate human learning processes. Within the realm of information processing, ANNs are part of a suite of methodologies employed to extract knowledge, patterns, or models from extensive datasets. Broadly speaking, attention should be directed towards the following three categories of artificial neural networks:
1) Multilayer Perceptron (MLP).
2) Convolutional Neural Network (CNN).
3) Recurrent Neural Network (RNN).

Super Artificial Intelligences (SAIs) are entrusted with addressing a diverse array of challenges, and these three types of neural networks have demonstrated their reliability and flexibility over decades of use. They significantly enhance the adaptability of SAIs, catering to the intricacies of various forecasting and prediction problems, spanning both structured and unstructured datasets at the Big Data level. Moreover, they encompass numerous subtypes to better suit specific needs. As this emerging technology rapidly expands, particularly among startups in Silicon Valley, questions arise regarding its impact on humanity whether it genuinely benefits or harms. The increasing automation of traditional labor tasks, exemplified by fully automated assembly lines, raises concerns about the evolving landscape [14]. In the modern world, there is a sobering reality regarding economic growth, with the effectiveness of traditional production levers like labor and capital investment diminishing. Artificial intelligence (AI) emerges as a novel production component capable of fostering new sources of growth, reshaping work processes, and underscoring the pivotal role of people in driving corporate growth.

III. USE CASES FOR AI/ML IN IMPROVING CLOUD RESILIENCE

The use of artificial intelligence (AI) and machine learning (ML) technologies has emerged as a promising strategy for improving cloud resilience. In recent years, there has been a growing interest in applying AI/ML techniques to address the challenges of resilience in cloud computing environments. AI/ML can be used to identify patterns and anomalies in large volumes of data, monitor system performance, and detect and predict potential failures or outages. It can also be used to optimize resource utilization and dynamically adjust resource provisioning to meet changing demand. This combination of AI/ML techniques has the potential to improve the reliability and efficiency of cloud systems [15]. The use of AI/ML for cloud resilience is based on the ability of these technologies to process large volumes of data, analyze patterns, and make predictions. AI/ML models can be trained on historical data to detect patterns and anomalies in cloud system performance. These models can then be used to monitor current system performance and identify potential issues in real-time. Additionally, AI/ML models can be used to forecast future system performance and identify potential outages or failures. This information can then be used to take proactive actions to improve system resilience, such as load balancing or resource reallocation. In terms of specific use cases, AI/ML can be used to:

I. **Detect and diagnose issues in real-time**, such as anomalies in data center power consumption, hardware failure, or network congestion; Predict potential outages and failures based on historical data and trends; Optimize resource utilization by dynamically adjusting resource allocation based on changing demand; Optimize workload placement by identifying the best location for workloads based on resource availability and performance requirements; Automate system responses to failures or issues, such as triggering failover.

II. **Dynamic resource allocation** is an approach to improving cloud resilience by adjusting the allocation of resources in real-time based on changing demand and system performance. Traditional cloud resource management methods rely on static resource allocation, where resources are allocated based on static policies and predetermined thresholds [15]. However, this approach may not be able to effectively handle unexpected changes in demand or system performance. Dynamic resource allocation addresses this issue by allowing for the automatic reallocation of resources based on real-time data. This approach can help improve cloud resilience by ensuring that resources are available when and where they are needed.

III. **Anomaly detection** is an important area of research in cloud resilience, as it involves identifying unexpected and abnormal events or patterns that can lead to service disruptions or failures. There are various AI/ML-based approaches for anomaly detection, including unsupervised learning methods such as clustering, density-based methods, and neural networks. These methods can be used to detect anomalies in system performance metrics such as resource utilization, network traffic, and CPU utilization. The detected anomalies can then be analyzed to understand the root cause and take appropriate actions to prevent service disruptions [9]. The application of anomaly detection in cloud resilience can help to improve system availability, reduce downtime, and improve system performance. It can also help to improve system security by detecting and preventing malicious activity [16]. Additionally, anomaly detection can be used to optimize system resources by identifying suboptimal resource utilization and taking appropriate actions. Overall, the application of anomaly detection can help to improve the reliability, security, and efficiency of cloud systems.

IV. **Automated healing** is a key aspect of cloud resilience, as it involves the use of intelligent systems to automatically detect and respond to system failures or performance issues. With automated healing, the cloud system can automatically recover from failures or degradations without human intervention. This can help to reduce downtime and improve system availability. Some examples of automated healing techniques include auto-scaling, self-healing containers, and intelligent fault detection and isolation. These techniques can be used to ensure that the system is always operating at optimal performance and can quickly recover from failures [18]. One important aspect of automated healing is the use of machine learning algorithms to detect and classify system failures. By using historical data and real-time monitoring data, machine learning algorithms can learn the normal behavior of the system and identify anomalies [17] [19] [20]. Once an anomaly is detected, the system can take appropriate actions to restore the system to a healthy state [21]. The combination of machine learning and automation allows for faster and more effective recovery from system failures. This in turn improves system availability and resilience. Additionally, automated healing can also be used to optimize system resources and reduce costs.

IV. CONCLUSIONS

The use of AI and machine learning has the potential to improve the resilience of cloud systems in several ways. By developing models that can detect and prevent failures, or optimize resource allocation, cloud systems can become more reliable and efficient, leading to improved uptime and reduced costs. With the increasing complexity and scale of cloud systems, the need for tools that can improve their resilience will continue to grow. The use of AI and machine learning represents a promising approach for addressing this need and improving the overall performance of cloud systems. In summary, artificial intelligence has grabbed the center stage of business intelligence, despite having been around for decades, due to the growing pervasiveness of data, the scalability of cloud computing, the availability of AI accelerators, and the sophistication of the ML and DL algorithms. Technology of AI in its today's growth, companies like IDC technology, predict 40% of Digital Transformation (DX) and big data initiatives will use AI services; 75% of commercial enterprise apps will use AI, over 90% of consumers will interact with customers support bots, and over 50% of new industrial robots will leverage AI. In conclusion, Artificial Intelligence (AI) offers a pathway to augment growth by functioning as a hybrid of capital and labor, thereby amplifying and surpassing the existing potential of these factors to propel economic growth. Our research unveils unprecedented potential for value development. Businesses and organizations embracing this trend are advised to do so promptly, considering the swift advancement of artificial intelligence and its evolution into super artificial intelligence.

Conflicts of Interest

The authors declare that they have no conflicts of interest concerning the publication of this paper.

V. REFERENCES

[1] B. Temitope A and A. Oluwaseyi O, "Solving Key Business Challenges For A Client's E-Commerce Using Ml Techniques," 2023, [Online]. Available: airfoiltools.com

[2] J. Jung, M. Maeda, A. Chang, M. Bhandari, A. Ashapure, and J. Landivar-Bowles, "The potential of remote sensing and artificial intelligence as tools to improve the resilience of agriculture production systems," *Curr. Opin. Biotechnol.*, vol. 70, pp. 15–22, 2021, doi: 10.1016/j.copbio.2020.09.003.

[3] J. Gao, H. Wang, and H. Shen, "Machine Learning Based Workload Prediction in Cloud Computing," *Proc. - Int. Conf. Comput. Commun. Networks, ICCCN*, vol. 2020-Augus, 2020, doi: 10.1109/ICCCN49398.2020.9209730.

[4] Bahman Zohuri and Farhang Mossavar Rahmani, "Artificial Intelligence Driven Resiliency with Machine Learning and Deep Learning Components," *J. Commun. Comput.*, vol. 15, no. 1, pp. 1–13, 2019, doi: 10.17265/1548-7709/2019.01.001.

[5] J. Wang, L. Zhao, J. Liu, and N. Kato, "Smart Resource Allocation for Mobile Edge Computing: A Deep Reinforcement Learning Approach," *IEEE Trans. Emerg. Top. Comput.*, vol. 9, no. 3, pp. 1529–1541, 2021, doi: 10.1109/TETC.2019.2902661.

[6] P. Muralidhara, "the Evolution of Cloud Computing Security: Addressing Emerging Threats," *Int. J. Comput. Sci. Technol.*, vol. 1, no. 4, pp. 1–33, 2017, [Online]. Available: http://ijcst.com.pk/IJCST/article/view/237%0Ahttps://ijcst.com.pk/IJCST/article/download/237/191

[7] A. Reddy Kunduru, "Central Asian Journal of Mathematical Theory and Computer Sciences Artificial Intelligence Usage in Cloud Application Performance Improvement," pp. 42–47, 2023, [Online]. Available: https://cajmtcs.centralasianstudies.org

[8] S. Saad, W. Briguglio, and H. Elmiligi, "The Curious Case of Machine Learning in Malware Detection," *Int. Conf. Inf. Syst. Secur. Priv.*, pp. 528–535, 2019, doi: 10.5220/0007470705280535.

[9] G. Qu, H. Wu, R. Li, and P. Jiao, "DMRO: A Deep Meta Reinforcement Learning-Based Task Offloading Framework for Edge-Cloud Computing," *IEEE Trans. Netw. Serv. Manag.*, vol. 18, no. 3, pp. 3448–3459, 2021, doi: 10.1109/TNSM.2021.3087258.

[10] G. Zhou, W. Tian, and R. Buyya, *Deep Reinforcement Learning-based Methods for Resource Scheduling in Cloud Computing: A Review and Future Directions*, vol. 1, no. 1. Association for Computing Machinery, 2021. [Online]. Available: http://arxiv.org/abs/2105.04086

[11] W. Wu and A. Plakhtii, "E-Learning Based on Cloud Computing," *Int. J. Emerg. Technol. Learn.*, vol. 16, no. 10, pp. 4–17, 2021, doi: 10.3991/ijet.v16i10.18579.

[12] O. O. Alabi, S. A. Ajagbe, O. A. Adeaga, and M. O. Adigun, "Investigating Fuel Adulteration Using Arduino as an Engine Protection Device (EPD)," *J. Hunan Univ. (Natural Sci.*, vol. 50, no. 9, 2023, doi: https://doi.org/10.55463/issn.1674-2974.50.9.11.

[13] O. O. Alabi, O. A. Adeaga, S. A. Ajagbe, T. S. Adekunle, and M. O. Adigun, "Performance of the Solar PV Module of the Dual Solar Axis Tracker of a Smart Home Monitoring System," *J. Hunan Univ. (Natural Sci.*, vol. 50, no. 9, 2023, doi: https://doi.org/10.55463/issn.1674-2974.50.9.3.

[14] F. O. Olowononi, D. B. Rawat, and C. Liu, "Resilient Machine Learning for Networked Cyber Physical Systems: A Survey for Machine Learning Security to Securing Machine Learning for CPS," *IEEE Commun. Surv. Tutorials*, vol. 23, no. 1, pp. 524–552, 2021, doi: 10.1109/COMST.2020.3036778.

[15] B. Zohuri, M. Moghaddam, and F. Mossavar-rahmani, "International Journal of Theoretical & Computational Physics Business Resilience System Integrated Artificial Intelligence System," vol. 3, no. 1, pp. 1–7, 2022.

[16] U. A. Butt *et al.*, "A review of machine learning algorithms for cloud computing security," *Electron.*, vol. 9, no. 9, pp. 1–25, 2020, doi: 10.3390/electronics9091379.

[17] T. S. Adekunle, T. A. Adeleke, O. Sunday, and O. O. Alabi, "A Framework for Robust Attack Detection and Classification using Rap-Densenet," *Paradig. plus J. Informatics Artif. Intell.*, vol. 4, no. 2, pp. 1–17, 2023, doi:

10.55969/paradigmplus.v4n2a1.

[18] O. S. Ojo, M. O. Oyediran, B. J. Bamgbade, A. E. Adeniyi, G. N. Ebong, and S. A. Ajagbe, "Development of an Improved Convolutional Neural Network for an Automated Face Based University Attendance System," *ParadigmPlus*, vol. 4, no. 1, pp. 18–28, 2023, doi: 10.55969/paradigmplus.v4n1a2.

[19] M. A. Mukwevho and T. Celik, "Toward a Smart Cloud: A Review of Fault-Tolerance Methods in Cloud Systems," *IEEE Trans. Serv. Comput.*, vol. 14, no. 2, pp. 589–605, 2021, doi: 10.1109/TSC.2018.2816644.

[20] S. Mishra and A. K. Tyagi, "The Role of Machine Learning Techniques in Internet of Things-Based Cloud Applications," *Internet of Things*, no. January, pp. 105–135, 2022, doi: 10.1007/978-3-030-87059-1_4.

[21] Temitope S. Adekunle; Morolake O. Lawrence; Oluwaseyi O. Alabi; Adenrele A. Afolorunso; Godwin N. Ebong; Matthew A. Oladipupo, "Deep Learning for Plant Disease Detection," Comput. Sci. Inf. Technol., vol. 5, no. 1, pp. 49–56, 2024, https://doi.org/10.11591/csit.v5i1.pp49-56