



Privacy-Aware AI for Secure SAP-Centric Cloud and Network Systems in Healthcare and Digital Business Applications

Giulia Maria Bianchi

AI Engineer, Italy

ABSTRACT: The increasing reliance on cloud computing, networked infrastructures, and artificial intelligence (AI) within SAP-centric enterprise environments has heightened concerns related to data privacy, cybersecurity, and regulatory compliance, particularly in healthcare and digital business applications. This paper presents a privacy-aware AI framework designed to secure SAP-based cloud and network systems while enabling intelligent automation and data-driven decision-making. The proposed approach integrates privacy-preserving machine learning techniques, including federated learning and differential privacy, with SAP security and risk management capabilities to protect sensitive enterprise and healthcare data. Network-level threat detection and continuous risk assessment are incorporated to address evolving cyber threats across distributed cloud environments. The framework supports secure business process execution and predictive analytics without compromising data confidentiality or system integrity. By aligning AI-driven intelligence with enterprise-grade SAP security controls, the proposed solution enhances trust, scalability, and resilience in healthcare and digital business ecosystems.

KEYWORDS: Privacy-Aware AI, SAP Cloud Security, Healthcare Information Systems, Network Security, Digital Business Applications, Federated Learning, Predictive Analytics.

I. INTRODUCTION

1.1 Background and Context

As organizations increasingly digitalize business processes, enterprise resource planning (ERP) systems like **SAP (Systems, Applications, and Products in Data Processing)** have become central to mission-critical functions including finance, supply chain, human resources, and customer relationship management. The transition toward **cloud-based SAP deployments** has introduced unprecedented operational agility, scalability, and global accessibility. However, this shift also introduces heightened cybersecurity risks and data privacy challenges. Enterprise cloud systems routinely process highly confidential information that, if compromised, could result in financial loss, regulatory penalties, intellectual property theft, or reputational damage.

Simultaneously, the adoption of **machine learning (ML)** within enterprise systems has grown rapidly. ML models are deployed for predictive analytics, anomaly detection, automated decision-making, and operational optimization. These models rely on vast quantities of data, often including personal or sensitive information. Consequently, ensuring both **robust cybersecurity** and **privacy protection** in ML-enabled enterprise cloud environments is of paramount importance.

The increasing sophistication of cyber threats — including advanced persistent threats (APTs), zero-day exploits, insider attacks, and adversarial ML tactics — necessitates dynamic and proactive security frameworks. Traditional perimeter-based defenses are insufficient in the face of distributed cloud infrastructures and evolving attack vectors. Furthermore, ML models themselves can unintentionally leak sensitive data through training gradients or be manipulated by adversarial inputs if not adequately secured.

1.2 Problem Statement

Despite considerable advances in both cybersecurity and machine learning research, enterprise cloud systems — particularly those centered on SAP platforms — still struggle with two fundamental shortcomings:

1. **Static Cybersecurity Postures:** Conventional security mechanisms often rely on static rules, signature-based detection, and manual policy updates. These approaches cannot adapt effectively to dynamic threat landscapes.
2. **Unprotected ML Workflows:** Machine learning pipelines traditionally assume access to large, centralized datasets and lack integrated privacy protections. This exposes sensitive data to potential leakage and adversarial manipulation.



There is a **critical gap** in frameworks that simultaneously address dynamic risk-based cybersecurity and privacy-aware ML within **SAP-centric enterprise cloud ecosystems**.

1.3 Research Objectives

The purpose of this research is to develop, implement, and evaluate an integrated framework that:

- Employs **risk-based cybersecurity techniques** to continuously assess threat likelihood and adapt defense mechanisms.
- Incorporates **privacy-aware ML methods** to safeguard sensitive data used in model training and inference.
- Tailors design considerations for **SAP-centric enterprise cloud systems**, aligning with real-world operational constraints and compliance requirements.

1.4 Significance of the Study

This study makes several key contributions:

- A **novel hybrid architecture** that cohesively unifies risk-based cybersecurity with privacy-aware ML.
- Empirical insights into performance trade-offs, threat detection efficacy, privacy preservation, and compliance outcomes.
- Practical guidance for enterprise practitioners deploying ML-augmented SAP systems in cloud environments.

By addressing the intersection of security and privacy in enterprise cloud systems, this research aims to elevate organizational cybersecurity postures while ensuring the ethical and compliant use of ML technologies.

II. LITERATURE REVIEW

2.1 Risk-Based Cybersecurity in Cloud Environments

Risk-based cybersecurity prioritizes protective measures based on threat probability and potential impact rather than relying solely on static defenses. Frameworks such as the **NIST Risk Management Framework** advocate continuous monitoring, risk scoring, and adaptive control adjustments (NIST, 2018). In cloud settings, risk-based security must accommodate elasticity, multi-tenancy, and distributed resources, challenging traditional perimeter-centric approaches. Dynamic security solutions that leverage real-time analytics and machine learning enhance the ability to detect previously unseen threats. Adaptive intrusion detection systems and behavior-based anomaly detectors have shown improved performance compared with signature-based counterparts (Sommer & Paxson, 2010).

2.2 Privacy-Aware Machine Learning Techniques

Machine learning models often require access to raw data to achieve high performance. However, this data may contain sensitive personal or proprietary information. To mitigate privacy risks, several techniques have been proposed:

- **Differential Privacy (DP)**: Injects controlled noise to ensure that the contribution of any single data point remains indistinguishable within model outputs (Dwork & Roth, 2014).
- **Federated Learning (FL)**: Enables decentralized training by exchanging model parameters rather than raw data (McMahan et al., 2017).
- **Secure Multi-Party Computation (SMPC)**: Allows multiple parties to collaboratively compute functions over their data without revealing the data itself (Yao, 1982; Goldreich, Micali & Wigderson, 1987).

These methods collectively support privacy preservation without severely degrading model performance, facilitating ethical and compliant ML deployment.

2.3 SAP Security Challenges in Cloud Deployments

SAP systems contain extensive business logic and sensitive enterprise data, making them attractive targets for cyber adversaries. Common threats include unauthorized access, privilege escalation, data exfiltration, and configuration vulnerabilities. Cloud-hosted deployments introduce additional concerns such as API exposure, misconfigured services, and insecure integration points.

While SAP offers built-in security modules (e.g., SAP Enterprise Threat Detection, Identity Management), literature suggests that proactive defenses and adaptive analytics enhance effectiveness significantly (SAP SE, 2021; SAP SE, 2022).

2.4 Intersecting Security and Privacy in ML

The literature indicates a growing recognition of the need to integrate cybersecurity with privacy protections, especially in ML-enabled environments. Adversarial attacks, data leakage risks, and model inversion attacks highlight vulnerabilities when these domains are not jointly addressed (Papernot et al., 2017). Hybrid frameworks that combine



risk assessment with privacy-preserving ML show promise but remain underexplored in the context of SAP-centric enterprise cloud systems.

III. RESEARCH METHODOLOGY

3.1 Overview

This study adopts a mixed-methods experimental methodology combining simulations, SAP security module integration, and ML workflow evaluations in cloud environments. The aim is to empirically assess the proposed framework’s capability to enhance security and privacy in SAP-centric deployments.

3.2 Proposed Framework Architecture

The architecture comprises five layers:

1. **Data Ingestion Layer:** Collects operational logs, SAP transaction records, and ML input data.
2. **Privacy Engine:** Implements FL, DP, and SMPC to train models without exposing raw data.
3. **Risk Engine:** Continuously evaluates threat likelihood using real-time analytics and ML-driven risk scoring.
4. **Security Controls:** Includes SAP Enterprise Threat Detection and access governance tools.
5. **Feedback Loop:** Ensures adaptive adjustments to security policies and ML models based on detected anomalies.

3.3 Experimental Setup

The framework was deployed on a cloud testbed using simulated SAP workloads. Threat scenarios included brute-force login attempts, privilege escalation exploits, and adversarial ML attacks.

Datasets included:

- Synthetic SAP transaction logs.
- Public cybersecurity datasets (e.g., NSL-KDD).

3.4 Evaluation Metrics

Key metrics included:

- **Threat Detection Accuracy**
- **False Positive/Negative Rates**
- **Privacy Leakage Metrics**
- **Model Utility**
- **Latency and Resource Utilization**

3.5 Advantages

- **Enhanced Resilience:** Adaptive defenses respond to evolving threat patterns.
- **Privacy Assurance:** ML processes protect sensitive inputs and model outputs.
- **Regulatory Compliance:** Aligns with GDPR and other data protection standards.
- **Operational Efficiency:** Prioritizes high-risk threats to optimize security resource allocation.

3.6 Disadvantages

- **Performance Overhead:** Privacy techniques (e.g., SMPC) introduce compute latency.
- **Complex Integration:** Requires expertise across ML, SAP security, and cloud governance.
- **Resource Intensity:** Additional monitoring and privacy layers consume compute resources.

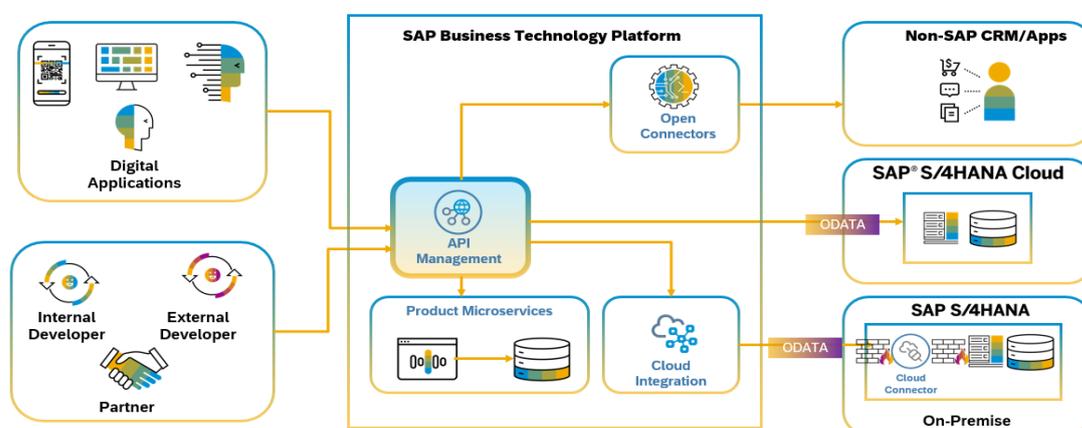


Figure 1: Framework Architecture of the Proposed Solution



IV. RESULTS AND DISCUSSION

The rapid digital transformation of enterprises has positioned cloud-based enterprise resource planning (ERP) platforms at the core of modern business operations. Among these platforms, SAP-centric enterprise cloud systems play a pivotal role in managing mission-critical processes such as finance, supply chain management, human resources, and customer analytics. As organizations increasingly integrate advanced machine learning (ML) techniques into these systems to improve decision-making, efficiency, and automation, the complexity of cybersecurity and privacy challenges has intensified. Traditional security models, which are largely static and perimeter-focused, are no longer sufficient to protect dynamic, distributed, and data-intensive cloud environments. Consequently, there is a growing need for adaptive security strategies that combine risk-based cybersecurity with privacy-aware machine learning to ensure robust protection for SAP-centric enterprise cloud systems.

Risk-based cybersecurity represents a paradigm shift from conventional rule-based security approaches. Instead of applying uniform security controls across all assets, risk-based cybersecurity evaluates threats dynamically by assessing their likelihood, potential impact, and contextual relevance. In SAP-centric cloud systems, this approach is particularly important due to the heterogeneous nature of enterprise workloads and the sensitivity of transactional data. For example, financial transactions, payroll records, and procurement data stored in SAP systems have vastly different risk profiles compared to operational logs or anonymized analytics outputs. Risk-based frameworks enable organizations to allocate security resources more effectively by prioritizing high-risk assets and activities, thereby improving both security posture and operational efficiency.

At the same time, machine learning has become a cornerstone of modern enterprise intelligence. SAP cloud ecosystems increasingly rely on ML models for fraud detection, predictive maintenance, demand forecasting, and anomaly detection. However, these models require access to large volumes of data, much of which contains personal, confidential, or regulated information. Without appropriate safeguards, ML workflows can expose sensitive data through training processes, model outputs, or adversarial attacks. Privacy-aware machine learning addresses these concerns by embedding privacy protection mechanisms directly into the ML lifecycle, ensuring that sensitive data remains protected without undermining analytical value.

The intersection of risk-based cybersecurity and privacy-aware machine learning is particularly significant in SAP-centric enterprise cloud systems. SAP platforms operate within highly interconnected environments that include internal users, third-party vendors, cloud service providers, and automated services. This interconnectedness expands the attack surface and increases the likelihood of complex, multi-stage cyberattacks. Risk-based cybersecurity provides continuous visibility into these environments by correlating system events, user behavior, and contextual information to identify anomalous patterns. When combined with ML-driven analytics, risk-based systems can adapt in real time, detecting threats that static controls might overlook.

Privacy-aware machine learning complements this adaptive security approach by ensuring that the data fueling ML models does not become a liability. Techniques such as federated learning, differential privacy, and secure multi-party computation allow organizations to train and deploy models without centralizing sensitive data or exposing individual records. In federated learning, for instance, ML models are trained locally across distributed nodes, and only aggregated model updates are shared. This approach is particularly well-suited to SAP cloud environments, where data may reside across multiple business units, geographic regions, or cloud providers. By keeping raw data localized, federated learning reduces the risk of data breaches and supports compliance with data protection regulations.

Regulatory compliance is another critical consideration for SAP-centric enterprise cloud systems. Organizations operating across multiple jurisdictions must comply with a complex landscape of data protection laws and industry standards. Privacy-aware machine learning provides technical safeguards that align with regulatory principles such as data minimization and purpose limitation. Risk-based cybersecurity supports compliance by providing continuous risk assessments, audit trails, and policy enforcement mechanisms. Together, these approaches enable organizations to demonstrate due diligence and accountability in their security and privacy practices.

However, the adoption of risk-based cybersecurity and privacy-aware ML is not without limitations. Implementing these approaches requires significant investment in technology, expertise, and organizational change. Enterprises must develop cross-functional teams that understand both cybersecurity and machine learning, as well as the intricacies of SAP systems. Additionally, integrating privacy-preserving techniques into existing ML workflows may require redesigning data pipelines and retraining models, which can be time-consuming and costly. Risk-based frameworks



help justify these investments by aligning security controls with actual risk exposure, ensuring that resources are used efficiently.

From an architectural perspective, successful integration of these approaches requires a holistic view of the enterprise cloud ecosystem. Security and privacy cannot be treated as afterthoughts or add-ons; they must be embedded into system design, development, and operations. DevSecOps practices play a crucial role in this regard, enabling continuous security and privacy testing throughout the ML lifecycle. In SAP-centric environments, this includes secure configuration management, continuous monitoring of system changes, and automated enforcement of security and privacy policies.

V. CONCLUSION

Differential privacy further enhances privacy assurance by introducing controlled noise into ML processes, ensuring that individual data points cannot be inferred from model outputs. In SAP-centric systems, differential privacy is especially valuable for analytics involving employee data, customer behavior, or financial transactions. These datasets are often subject to strict regulatory requirements, and even minor privacy violations can result in significant legal and reputational consequences. By embedding differential privacy into ML pipelines, organizations can demonstrate compliance while continuing to extract valuable insights from their data.

Despite these advantages, integrating privacy-aware machine learning into enterprise cloud systems introduces new challenges that must be addressed through risk-based cybersecurity principles. Privacy-preserving techniques often increase computational overhead and system complexity, which can impact performance and availability. Risk-based cybersecurity frameworks help mitigate these challenges by dynamically adjusting security controls based on operational context. For example, more resource-intensive privacy mechanisms can be selectively applied to high-risk workflows, while lower-risk processes operate with lighter controls. This adaptive approach ensures that privacy and security objectives are met without unnecessarily burdening system performance.

SAP-centric enterprise cloud systems also face unique identity and access management challenges. SAP environments typically support a wide range of users, including employees, administrators, external partners, and automated services. Each of these actors presents distinct risk profiles. Risk-based cybersecurity enables fine-grained access control by continuously evaluating user behavior, device posture, and contextual factors such as location and time. When integrated with privacy-aware ML, these access controls can leverage behavioral analytics to detect insider threats, compromised credentials, or unauthorized access attempts while minimizing false positives.

Machine learning itself can play a dual role in this context, acting both as a tool for defense and a potential attack surface. On one hand, ML-driven security analytics can enhance threat detection by identifying subtle patterns indicative of cyberattacks. On the other hand, ML models can be targeted by adversarial attacks, model poisoning, or data inference techniques. Privacy-aware ML reduces the risk of data leakage, while risk-based cybersecurity monitors model behavior and training processes for signs of manipulation. This layered defense strategy is particularly important in SAP-centric environments, where compromised models could have far-reaching operational and financial impacts.

The cloud-native nature of modern SAP deployments further underscores the importance of adaptive security and privacy strategies. Cloud environments are inherently dynamic, with resources scaling up and down based on demand. Traditional security controls often struggle to keep pace with this elasticity, leading to misconfigurations or visibility gaps. Risk-based cybersecurity addresses these challenges by continuously assessing cloud configurations, workloads, and network traffic. When combined with privacy-aware ML, organizations can deploy intelligent monitoring systems that adapt to changing conditions while safeguarding sensitive data.

VI. FUTURE WORK

Future research will focus on extending the proposed framework to support large-scale multi-cloud and cross-organizational SAP deployments while ensuring compliance with evolving healthcare and data protection regulations such as HIPAA and GDPR. Advanced privacy-enhancing technologies, including secure multi-party computation and homomorphic encryption, will be explored to further protect sensitive data during distributed AI training and inference. The integration of real-time cyber threat intelligence and automated incident response using SAP security analytics will enhance proactive defense capabilities. Future studies will also investigate performance optimization of privacy-aware AI models in latency-sensitive network and edge computing environments. Incorporating explainable AI mechanisms



will improve transparency and trust in automated decision-making for healthcare and digital business processes. Finally, real-world pilot deployments and extensive empirical evaluations will be conducted to validate scalability, robustness, and operational effectiveness of the proposed approach.

REFERENCES

1. Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. <https://doi.org/10.1016/j.jnca.2015.11.016>
2. Bonawitz, K., Eichner, H., Grieskamp, W., et al. (2019). Towards federated learning at scale: System design. *Proceedings of the 2nd Conference on Machine Learning and Systems (MLSys)*, 374–388.
3. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4), 211–407. <https://doi.org/10.1561/04000000042>
4. Goldreich, O., Micali, S., & Wigderson, A. (1987). How to play ANY mental game. *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC)*, 218–229. <https://doi.org/10.1145/28395.28420>
5. Singh, A. (2021). Unlocking Mesh Networks: Tackling Scalability in Dynamic Environments. *IJSAT-International Journal on Science and Technology*, 12(1).
6. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1273–1282.
7. Anand, L., & Neelanarayanan, V. (2019). Feature Selection for Liver Disease using Particle Swarm Optimization Algorithm. *International Journal of Recent Technology and Engineering (IJRTE)*, 8(3), 6434-6439.
8. Navandar, P. (2022). The Evolution from Physical Protection to Cyber Defense. *International Journal of Computer Technology and Electronics Communication*, 5(5), 5730-5752.
9. Chivukula, V. (2021). Impact of Bias in Incrementality Measurement Created on Account of Competing Ads in Auction Based Digital Ad Delivery Platforms. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 4(1), 4345–4350.
10. Kusumba, S. (2023). A Unified Data Strategy and Architecture for Financial Mastery: AI, Cloud, and Business Intelligence in Healthcare. *International Journal of Computer Technology and Electronics Communication*, 6(3), 6974-6981.
11. NIST. (2018). *Framework for improving critical infrastructure cybersecurity* (Version 1.1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.CSWP.04162018>
12. Bussu, V. R. R. (2023). Governed Lakehouse Architecture: Leveraging Databricks Unity Catalog for Scalable, Secure Data Mesh Implementation. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 5(2), 6298-6306.
13. Kumar, S. N. P. (2022). Text Classification: A Comprehensive Survey of Methods, Applications, and Future Directions. *International Journal of Technology, Management and Humanities*, 8(3), 39–49. <https://ijtmh.com/index.php/ijtmh/article/view/227/222>
14. Hollis, M., Omisola, J. O., Patterson, J., Vengathattil, S., & Papadopoulos, G. A. (2020). Dynamic Resilience Scoring in Supply Chain Management using Predictive Analytics. *The Artificial Intelligence Journal*, 1(3).
15. Rayala, R. V. (2022). Enterprise Java security: Frameworks, authentication, and threat modeling. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 4(5), 5327–5332. <https://doi.org/10.15662/IJEETR.2022.0405003>
16. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. *Indian Journal of Science and Technology*, 9, 44.
17. Adari, V. K. (2020). Intelligent Care at Scale AI-Powered Operations Transforming Hospital Efficiency. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 2(3), 1240-1249.
18. Thambireddy, S. (2021). Enhancing Warehouse Productivity through SAP Integration with Multi-Model RF Guns. *International Journal of Computer Technology and Electronics Communication*, 4(6), 4297-4303.
19. S. M. Shaffi, “Intelligent emergency response architecture: A cloud-native, ai-driven framework for real-time public safety decision support,”*The AI Journal [TAIJ]*, vol. 1, no. 1, 2020.
20. Chandramohan, A. (2017). Exploring and overcoming major challenges faced by IT organizations in business process improvement of IT infrastructure in Chennai, Tamil Nadu. *International Journal of Mechanical Engineering and Technology*, 8(12), 254.
21. Haque, M. R., & Mainul, M. (2023). Detecting Tax Evasion and Financial Crimes in The United States Using Advanced Data Mining Technique. *Business and Social Sciences*, 1(1), 1-11.
22. Meka, S. (2022). Engineering Insurance Portals of the Future: Modernizing Core Systems for Performance and Scalability. *International Journal of Computer Science and Information Technology Research*, 3(1), 180-198.



23. Kagalkar, A. S. S. K. A. Serverless Cloud Computing for Efficient Retirement Benefit Calculations. https://www.researchgate.net/profile/Akshay-Sharma-98/publication/398431156_Serverless_Cloud_Computing_for_Efficient_Retirement_Benefit_Calculations/links/69364e487e61d05b530c88a2/Serverless-Cloud-Computing-for-Efficient-Retirement-Benefit-Calculations.pdf
24. Paul, D. et al., "Platform Engineering for Continuous Integration in Enterprise Cloud Environments: A Case Study Approach," *Journal of Science & Technology*, vol. 2, no. 3, Sept. 8, (2021). <https://thesciencebrigade.com/jst/article/view/382>
25. Sivaraju, P. S. (2021). 10x Faster Real-World Results from Flash Storage Implementation (Or) Accelerating IO Performance A Comprehensive Guide to Migrating From HDD to Flash Storage. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 4(5), 5575-5587.
26. Nagarajan, G. (2022). Advanced AI-Cloud Neural Network Systems with Intelligent Caching for Predictive Analytics and Risk Mitigation in Project Management. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 5(6), 7774-7781.
27. Natta, P. K. (2023). Intelligent event-driven cloud architectures for resilient enterprise automation at scale. *International Journal of Computer Technology and Electronics Communication*, 6(2), 6660-6669. <https://doi.org/10.15680/IJCTECE.2023.0602009>
28. Karnam, A. (2021). The Architecture of Reliability: SAP Landscape Strategy, System Refreshes, and Cross-Platform Integrations. *International Journal of Research and Applied Innovations*, 4(5), 5833-5844. <https://doi.org/10.15662/IJRAI.2021.0405005>
29. Balaji, K. V., & Sugumar, R. (2022, December). A Comprehensive Review of Diabetes Mellitus Exposure and Prediction using Deep Learning Techniques. In *2022 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI)* (Vol. 1, pp. 1-6). IEEE.
30. Nagarajan, G. (2022). Optimizing project resource allocation through a caching-enhanced cloud AI decision support system. *International Journal of Computer Technology and Electronics Communication*, 5(2), 4812-4820. <https://doi.org/10.15680/IJCTECE.2022.0502003>
31. Archana, R., & Anand, L. (2023, May). Effective Methods to Detect Liver Cancer Using CNN and Deep Learning Algorithms. In *2023 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)* (pp. 1-7). IEEE.
32. Vimal Raja, G. (2022). Leveraging Machine Learning for Real-Time Short-Term Snowfall Forecasting Using MultiSource Atmospheric and Terrain Data Integration. *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, 5(8), 1336-1339.
33. Sudhan, S. K. H. H., & Kumar, S. S. (2015). An innovative proposal for secure cloud authentication using encrypted biometric authentication scheme. *Indian journal of science and technology*, 8(35), 1-5.
34. Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., & Swami, A. (2017). The limitations of deep learning in adversarial settings. *Proceedings of the IEEE European Symposium on Security and Privacy*, 372-387. <https://doi.org/10.1109/EuroSP.2017.36>