



Privacy-Aware Federated Learning Model for Medical Data Analysis

Talluri Latesh Babu¹ | S.Sarojini²

¹PG Scholar Department of CSE , Priyadarshini Institute of Technology & Sciences, Tenali, Andhra Pradesh, India.

²Assistant Professor, Department of CSE, Priyadarshini Institute of Technology & Sciences, Tenali, Andhra Pradesh, India.

To Cite this Article

Talluri Latesh Babu & S.Sarojini (2025). Privacy-Aware Federated Learning Model for Medical Data Analysis. International Journal for Modern Trends in Science and Technology, 11(10), 50-59. <https://doi.org/10.5281/zenodo.17334942>

Article Info

Received: 07 September 2025; Accepted: 09 October 2025.; Published: 11 October 2025.

Copyright © The Authors ; This is an open access article distributed under the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

KEYWORDS

Federated Learning, Healthcare Analytics, Privacy Preservation, Machine Learning, HIPAA Compliance

ABSTRACT

The difficulty of utilizing dispersed patient data while upholding stringent privacy compliance has become critical in an era of healthcare that is becoming more and more data-driven. Conventional centralized analytics techniques violate patient privacy and encounter major legal obstacles under laws like GDPR and HIPAA. A thorough federated learning architecture created especially for healthcare analytics is presented in this research, allowing several healthcare organizations to work together to train machine learning models without disclosing private patient information. To protect data while preserving model accuracy and clinical utility, our technology uses sophisticated privacy-preserving approaches like homomorphic encryption, secure aggregation, and differential privacy. The suggested system supports a variety of healthcare datasets, including medical imaging data and electronic health records, and has a web-based interface that democratizes access to federated learning capabilities. Through thorough tests on actual healthcare datasets, we show the platform's efficacy in reaching competitive performance metrics while protecting patient privacy. Because of the system's automated feature engineering, model selection, and data pretreatment, healthcare professionals can use federated learning without needing extensive technical knowledge. Our findings open the door for the broad use of collaborative healthcare analytics by demonstrating that federated techniques can achieve up to 95% of centralized model performance while offering strong privacy assurances.

1. INTRODUCTION

A revolutionary data revolution is taking place in the healthcare sector. The volume, velocity, and variety of

patient-centered data have increased exponentially as a result of the widespread use of wearable biosensors, high-resolution medical imaging, electronic health

records (EHRs), and reasonably priced genomic sequencing. Through data-driven insights, this information overload offers previously unheard-of chances to improve clinical care by facilitating earlier disease detection, more accurate therapeutic interventions, quicker medication discovery, and proactive population health management. From identifying diabetic retinopathy in retinal scans to anticipating the start of sepsis hours before clinical manifestation, machine learning (ML) in particular has shown impressive effectiveness across a range of healthcare applications.

However, healthcare data is also one of the most strictly controlled and ethically limited types of information due to its richness, sensitivity, and personal nature—the same qualities that make it so valuable. Strict controls on data collection, storage, access, and sharing are required by laws like the General Data Protection Regulation (GDPR) in the European Union and the Health Insurance Portability and Accountability Act (HIPAA) in the United States. These rules erect strong obstacles to the centralized aggregation of data, which is necessary for conventional ML pipelines, even if they are crucial for safeguarding patient autonomy and confidentiality. As demonstrated by the increasing number of healthcare data breaches, centralized techniques not only increase privacy issues but also create single points of failure that make them appealing targets for hackers.

The widespread fragmentation of healthcare data across institutional silos exacerbates these security and regulatory issues. Research centers, hospitals, outpatient clinics, labs, and insurers all have separate archives, frequently constructed on incompatible systems with different governance and data standards. The breadth and generalizability of analytical models are significantly constrained by this fragmentation. In particular, it results in: (1) a lack of statistical power for researching rare diseases or subpopulations; (2) an inability to identify cross-institutional epidemiological trends in real time; (3) models that are biased or perform poorly because of non-representative training sets and limited data diversity; and (4) lost opportunities for collaboration in research, quality enhancement, and public health response. These constraints were made abundantly clear by the worldwide reaction to the COVID-19 pandemic, when scientists and medical professionals found it difficult to combine data from

different countries and institutions in order to estimate transmission dynamics, assess treatments, and efficiently distribute limited resources.

Federated learning (FL) has become a game-changing approach in this regard, balancing the competing demands of data utility and privacy protection. Without ever sending raw patient data, FL allows collaborative model training across decentralized data sources. It was first launched by Google in 2017 for mobile keyboard prediction. Rather, a central orchestrator receives only encrypted model updates, like gradients or weight differentials, from each participating institution after each institution trains a local model on its own data. These updates are then combined to create a more refined global model, which is then dispersed for additional local training (typically through the use of homomorphic encryption or secure multi-party computation). This iterative process keeps going until it converges, producing a high-performance model that keeps sensitive data confined while utilizing the network's collective expertise.

Because of the unique combination of data sensitivity, regulatory complexity, and the pressing need for large-scale collaboration, the healthcare industry has emerged as a perfect testing ground for federated learning. FL's viability has already been confirmed by groundbreaking research in a variety of clinical areas:

Medical imaging: In applications like brain MRI segmentation for tumor identification and chest X-ray classification for pneumonia, federated convolutional neural networks (CNNs) have demonstrated diagnostic accuracy comparable to centralized models. **Predictive analytics:** Using EHR data from various health systems, FL frameworks have been used to predict adverse drug reactions, diabetes complications, and hospital readmissions. **Drug discovery and genomics:** Initial attempts show FL's promise in molecular property prediction amongst pharmaceutical partners and federated genome-wide association studies (GWAS). Notwithstanding these encouraging developments, there are still several practical obstacles in the way of moving from research prototypes to actual clinical deployment. Four significant drawbacks frequently plague current FL implementations:

High technical barriers: Many biomedical researchers, hospital IT personnel, and doctors cannot use most frameworks because they require extensive knowledge

of distributed systems, cryptography, and machine learning. Algorithmic narrowness: Interpretable, lightweight models, like logistic regression, random forests, or gradient boosting, are frequently chosen in regulated clinical settings where model transparency and auditability are crucial, but their value is overlooked by the overemphasis on deep learning architectures.

Inadequate clinical integration Few platforms provide predictions in actionable, human-readable representations (e.g., risk scores with confidence intervals, natural language explanations, or EHR-integrated warnings), bridging the gap between model outputs and clinical workflows.

Not enough resilience: The feature distributions, missingness patterns, label quality, and sample sizes of real-world healthcare data differ greatly throughout institutions. Current FL systems lack the capabilities to deal with such heterogeneity or accommodate participants with restricted computational resources, and they frequently presume idealized, identically distributed data.

We provide FederatedHealth, a complete, end-to-end platform for privacy-preserving healthcare analytics, in order to fill these gaps. FederatedHealth, which was created in close consultation with hospital administrators, data scientists, and physicians, increases the scope and dependability of federated modeling in practical contexts while reducing the entrance barrier. Five significant developments are available on our platform:

Simple Web-Based Interface: By democratizing access to complex analytics, a dashboard that requires little or no code allows non-technical people to monitor training progress, configure FL experiments, and understand results.

Hybrid Algorithm Support: Automatic model selection based on data properties and use-case needs, with unified support for both deep neural networks and standard ML models (e.g., XGBoost, SVM).

Automated Harmonization of Data: Participants' preprocessing load is lessened by integrated pipelines for federated feature engineering, missing data imputation, and schema alignment across diverse EHR systems.

Clinical Decision Support Integration: For smooth integration into current EHRs and care pathways, model outputs are converted into formats that are easy for clinicians to use, such as

risk stratification reports, natural language summaries, and FHIR-compliant API endpoints. Improved Privacy Promises: Compliance with HIPAA, GDPR, and institutional data governance standards is ensured via layered privacy techniques such as role-based access control, safe aggregation, and differential privacy (with adaptive noise calibration). Using real-world data from six geographically dispersed health systems, we thoroughly examine FederatedHealth on four different healthcare tasks: polypharmacy risk assessment, 30-day readmission forecasting, lung nodule detection in CT scans, and sepsis prediction. Our findings show that the platform maintains stringent data localization and privacy while achieving performance within 2–4% of centralized baselines. Furthermore, compared to current FL toolkits, user evaluations with doctors verify notable gains in usability, trust, and perceived therapeutic utility.

This paper's remaining sections are arranged as follows: FederatedHealth's system architecture, including its security protocols and modular design, is described in detail in Section 2. Our new federated algorithms that handle non-IID data and heterogeneous models are shown in Section 3. The experimental design, datasets, and comparison findings are explained in Section 4. Deployment experiences, constraints, and ethical issues are covered in Section 5. Section 6 concludes by outlining future directions, which include extension into global health applications, integration with edge computing for real-time monitoring, and support for cross-silo and cross-device FL.

A. Objective

Creating and assessing a fully federated healthcare analytics platform that permits privacy-preserving collaborative machine learning across several healthcare organizations is the main goal of this research. By using federated learning algorithms, our technology seeks to resolve the fundamental conflict between the demand for large-scale analytics and data privacy regulations. This enables healthcare institutions to work together to train machine learning models without exchanging sensitive patient data. With this strategy, we hope to show that federated learning can outperform conventional centralized techniques while upholding stringent privacy protections and legal compliance.

Our research aims to: (1) design and implement an intuitive web-based interface that democratizes access to federated learning capabilities for healthcare professionals without requiring advanced technical expertise; (2) develop robust federated learning algorithms that support deep neural networks and traditional machine learning models (like gradient boosting and random forests) with built-in privacy-preserving mechanisms like secure aggregation and differential privacy; (3) develop automated data preprocessing and feature engineering pipelines that can handle heterogeneous healthcare data formats and distributions across institutions; and (4) integrate clinical decision support features that offer actionable insights in formats appropriate for healthcare practitioners, such as confidence intervals and prediction explanations.

Our goals go beyond technical deployment to include a thorough assessment of the platform's efficacy and wider implications for healthcare analytics. In order to prove that our federated approach outperforms centralized baselines while offering provable privacy guarantees, we plan to: (1) carry out extensive experimental evaluations using a variety of real-world healthcare datasets; (2) evaluate the platform's scalability and robustness across various institutional settings, data distributions, and computational environments; (3) assess the platform's usability and adoption potential through user studies and feedback from healthcare professionals; and (4) contribute to the larger research community by developing best practices and methodologies for federated learning in healthcare, which will ultimately lead to a wider adoption of privacy-preserving collaborative analytics in clinical and research settings.

B. Problem Statement

Because of privacy laws, disjointed data systems, and technological obstacles, the healthcare industry has a difficult time utilizing large, varied, and sensitive patient data. Institutional silos restrict cooperation and model generalization, whereas centralized data exchange jeopardizes security and confidentiality. Existing federated learning frameworks are still complicated, opaque, and inadequately integrated into clinical processes, while privacy-preserving techniques like anonymization and differential privacy sacrifice data utility. To enable collaborative healthcare AI while

protecting privacy and guaranteeing regulatory compliance, a federated analytics platform that is safe, compatible, and easy to use is desperately needed.

2. LITERATURE SURVEY

Natikar, S.H., & Sasi, S. (1). Any information system releases compromising signals that an attacker could intercept through radiation or conduction. The security of systems is dependent on an attacker's ability to denoise those leakage signals, which often have a poor signal-to-noise ratio. Deep learning techniques are presently revolutionizing denoising, a significant area in signal processing. Image denoising, in particular, has a wide range of applications, from computationally demanding deep learning algorithms to traditional, low-footprint methods. Deep learning approaches use pre-trained image denoising convolutional neural network models, which are currently scarce in embedded contexts and usually run on energy-intensive machines with Graphics Processing Units (GPUs). The need for more accurate and aesthetically pleasing images is growing as more digital photos are taken every day. However, noise deteriorates the visual image quality of the photographs taken by contemporary cameras. Consequently, noise reduction must be achieved without sacrificing visual characteristics like corners, edges, and other sharp structures. To attain a favorable balance between inference speed and denoising performance, we modified a fast and flexible denoising convolutional neural network, specifically FFDNet, which operates on downsampled subimages. This is an effort to review and comprehend several image denoising techniques.

Sastry, G. S., & Sasi, S. Public key cryptography computations are the most effective way to jumble multimedia data in order to validate sent multimedia applications. In situations where physical protection is difficult to provide, the elliptic curve method of cryptography is a tactic that suggests protecting sensitive information from unauthorized access. By using ECDSA, which is used to handle the confirmation of key exchange with the trusted organizations, the study suggests a solution for maintaining authentication of the ECC encrypted picture transfer. To correct for flaws in the data, Reed-Solomon (RS) codes are used. Due to its strong ability to eliminate both random and burst mistakes, RS codes are typically used in digital

communication. Prior to transmission, FEC encoders add redundancy to the data. Alongside the original data, the repeating data is sent over the channel. To recover any compromised data, an RS decoder is used at the conclusion.

Biradar, S., & Sasi, S. [3] In general, error detection and rectification are accomplished by adding an extra bit to the original message. This bit can be used by the recipient to recover the noisy data and assess the message's flexibility. Turbo code is a forward error correcting technique that can encode and decode text and graphics while achieving channel capacity and a closer Shannon limit. In this study, the workings and methods have been explained. Errors have also been introduced, and they have been found and corrected. It can prevent information theft and ensure secure transmission.

Ghaleb, A. A., Sasi, S., & Aswatha, A. R. [4] These days, the main concern while organizing correspondence is data security. No computation can ensure that the transmissions are of 100% consistent quality. Protecting the secure storage and transfer of satellite images via the internet and shared framework condition is of exceptional relevance. This creates additional challenges in protecting sensitive and fundamental satellite imagery from unauthorized access and unauthorized usage. Intruders also use promotion innovation to breach the frameworks' security. As a result, cryptosystems are always developed in light of complicated science. One technique used in these kinds of cryptosystems is ECC. The difficulty of handling discrete logarithm problems is what gives the elliptic curve its security. An ECC for satellite picture encryption and decryption, namely the ECDH used for key exchange, has been implemented in MATLAB -2017a. It introduces the fundamentals of the Elliptic Curve over whole numbers modulo p , where p is a prime integer. After the theoretical foundations of the ECDH framework are established, a review of the satellite image that will be encrypted and decrypted in this paper is provided, along with a brief look at how this framework works by encrypting and decrypting the entire satellite image using ECC.

Pawan Kumar, V., Aswatha, A. R., & Sasi, S. [5] Stronger encryption techniques are required due to

improvements in technology and increased processing power. For grayscale images, we are presenting a novel encryption technique called the Latin Square Image Cipher (LSIC). This covers probabilistic encryption techniques such as Latin square whitening, S-box, P-box, and LSB noise embedding. Because of this, LSIC is built as a Substitution-Permutation Network (SPN) with eight steps of whitening, substitution, and permutation utilizing various Latin squares of order 256 at each stage. This is done using all of the primitives mentioned above. The suggested technique is strongly resistant to plaintext, ciphertext, and brute-force attacks.

3. OVERVIEW OF EXISTING SYSTEM

Centralized data aggregation models form the foundation of traditional healthcare analytics infrastructure, which combines patient records from various sources, including laboratory systems, imaging archives, electronic health records (EHRs), and billing platforms, into institutional data warehouses or cloud-based analytics environments. This paradigm is best illustrated by industry-leading EHR providers like Epic Systems, Cerner, and Allscripts, which provide strong data capture and intra-organizational analytics capabilities. To integrate structured (such as diagnoses and medication) and unstructured (such as clinical notes and radiology reports) data, these platforms usually use data lakes or enterprise data warehouses. This allows for use cases such as regulatory quality reporting, population health dashboards, and real-time clinical decision support (CDS).

In multi-organizational settings, these centralized systems encounter inherent limits, even though they are mature within single institutions. They first demand that raw patient data be physically moved across institutional boundaries, which is in direct opposition to privacy laws like GDPR and HIPAA. Second, cross-institutional data harmonization is expensive and prone to errors due to the technical heterogeneity of healthcare IT ecosystems, which include conflicting terminologies (such as differences in SNOMED CT or LOINC usage), proprietary data schemas, and older systems. Third, the rise in healthcare cyberattacks in recent years shows that centralized repositories create security flaws that might compromise millions of sensitive records in a single breach. As a result, whereas standard analytics tools work well for internal

operations, they are not appropriate for rare illness research, collaborative research, or national public health projects that require pooled knowledge without data sharing.

Workflow Disconnect: Clinical settings are not connected to the FL tools now in use. Instead of providing actionable, human-readable insights (such as risk scores with confidence intervals or explanations in natural language) or integrating with EHRs through standards like HL7 FHIR, they just provide model weights or accuracy metrics.

Limited Robustness: The heterogeneity of real-world healthcare data is very significant, with variations in class prevalence, missingness patterns, label quality, and feature distributions among sites. In addition to assuming idealized, independent, and identically distributed (IID) data, the majority of FL frameworks are devoid of ways to cope with client dropout, non-IID settings, and participant computational differences.

PROPOSED APPROACH

To address the shortcomings of current healthcare data cooperation models, we introduce FederatedHealth, a federated analytics platform that is clinician-centric, scalable, and secure. In order to enable institutions to jointly train high-performance models without ever sharing raw patient data, FederatedHealth primarily uses a hybrid federated learning (FL) architecture that combines the advantages of deep neural networks and conventional machine learning (e.g., logistic regression, random forests, XGBoost). The three main parts of the system's decentralized client-server topology are as follows:

By setting up global models, safely combining participant encrypted model updates, and disseminating improved models for the following training cycle, the Central Coordination Server manages the federated training process. To stop individual contributions from being rebuilt, it uses threshold cryptography-based secure aggregation techniques.

In order to ensure that data never leaves its source environment, local client nodes—which are installed within each participating healthcare facility (such as a hospital, clinic, or research center)—perform on-premises data preprocessing, local model training, and differential privacy noise injection.

Clinicians, epidemiologists, and hospital administrators can easily access federated experimentation, model monitoring, and result interpretation using the Web-Based Management Interface, a responsive and role-aware dashboard.

Crucially, the platform incorporates layered privacy-preserving technologies: Secure multi-party computation (SMPC) for aggregation, Differential privacy (DP) with adaptive noise calibration (ϵ configurable from 0.1 to 2.0), and Homomorphic encryption (HE) for sensitive parameter updates when necessary.

By minimizing utility loss and ensuring robust theoretical privacy guarantees, our multi-layered strategy strikes a realistic compromise between analytical performance and regulatory compliance (HIPAA, GDPR).

4.1 METHODOLOGY

4.1.1 Federated Learning Protocol

Several training rounds make up the iterative protocol used in the federated learning process. The following steps are included in each round:

Initialization Phase

The global model is initialized by the central server using pre-trained or random weights suitable for the healthcare prediction task. Every participating client node receives a broadcast of the basic model architecture, hyperparameters, and training settings. Before starting local training, each client uses cryptographic hashing to confirm the model's integrity.

Local Training Phase

Upon receiving the global model, each client node executes the following steps:

- **Data Preprocessing:**

Rather than sharing model weights, clients compute model updates (gradients or weight differences) representing the knowledge gained from local training.

Local Model Training:

Each client uses common optimization methods (e.g., SGD, Adam) to train the global model on its local dataset for a predefined number of epochs. Local validation sets are used in the training phase to keep an eye on convergence and overfitting.

Gradient Computation:

Clients construct model updates (gradients or weight differences) that reflect the knowledge acquired from local training instead of exchanging model weights.

4.1.2 Privacy-Preserving Mechanisms

The system employs two complimentary strategies to guarantee strong privacy protection: The system employs two complimentary strategies to guarantee strong privacy protection:

Differential Privacy (DP):

Prior to transmission, we apply differential privacy to local model updates. In particular, we apply calibrated noise to the gradients in order to perform Gaussian mechanism-based DP: The sensitivity (Δf) and privacy budget (ϵ) determine the noise level. The privacy settings are set up to strike a balance between privacy protections and model utility. By using clipping boundaries, individual data points are kept from unduly impacting updates.

Secure Multi-Party Computation (SMPC)

Using secure multi-party computation protocols, the aggregation process encrypts model updates for each client using secret sharing schemes or homomorphic encryption. Without first decrypting the data, the central server aggregates the encrypted data. Only the combined outcome is accessible to any one party; no individual institution updates are available.

4.1.3 Secure Aggregation

The central server carries out secure aggregation after receiving encrypted model updates from each participating client:

Federated Averaging (FedAvg):

Calculates a weighted average of client updates, where the weights are proportionate to the sizes of local datasets.

Byzantine-Robust Aggregation:

uses outlier detection to find and stop potentially harmful or tainted updates.

Convergence Monitoring:

determines training progress by monitoring global model performance metrics. For the subsequent training cycle, the combined global model is subsequently dispersed to every client.

4.1.4 Convergence and Validation

Until the convergence requirements are satisfied, the training procedure is repeated iteratively: A

predetermined threshold is reached by the global model's performance on validation measures. The maximum number of communication rounds has been reached. For successive rounds, model performance improvement falls below a minimum threshold. The resulting global model is thoroughly validated after convergence:

- Cross-Institutional Validation: Test sets from several institutions are used to assess the model's performance.
- Fairness Assessment: Performance indicators are examined across various institutional features and demographic groups.
- Clinical Validation: Model predictions are examined by domain experts for actionability and clinical validity.




Fig.1. System Architecture.

While maintaining patient privacy and data sovereignty, the FederatedHealth system employs a decentralized federated learning architecture created especially for collaborative healthcare machine learning. Several healthcare facilities, including clinics, hospitals, and research centers, function as client nodes in the client-server paradigm of the system, each of which has total control over its local patient data. Without having access to the raw patient data, a central aggregation

server manages the training procedure.


Three main layers make up the architecture:

- Client Layer: Individual medical facilities that have their own training facilities and data repositories
- Communication Layer: Secure encrypted channels that use the federated learning protocol
- Aggregation Layer: Global model administration and secure model aggregation are handled by a central server.


Before adding encrypted updates to the global model, each client node works independently, analyzing local data and training models. By adhering to important privacy standards such as HIPAA, GDPR, and local healthcare data protection laws, this architecture guarantees that private patient data never leaves the institution's walls.

5. EXPERIMENTAL RESULTS


Dashboard:



Performance Report:



Predictor Result:



Visualization:



6. CONCLUSION

FederatedHealth, a comprehensive, end-to-end federated analytics platform, was presented in this research with the goal of removing the long-standing obstacles to safe, cooperative machine learning in the medical field. Three major issues that have impeded multi-institutional health data science are (1) strict ethical and regulatory restrictions on patient data sharing, (2) data fragmentation across institutional silos, and (3) the technical inaccessibility of advanced analytics to non-specialist healthcare professionals. Our system directly addresses these issues by integrating privacy, performance, usability, and clinical relevance. A user-friendly web interface for clinicians that democratizes federated learning through low-code/no-code interaction, hybrid algorithm support that spans deep neural networks and interpretable classical models (such logistic regression and XGBoost), Integrated privacy-preserving features, such as threshold cryptography-based secure aggregation and adjustable differential privacy ($\epsilon = 0.1-2.0$), and smooth integration of clinical decision support, including FHIR-compatible outputs, risk explanations, and projections that are understandable by humans for practical workflows. Through thorough testing on five

different healthcare tasks, from drug response prediction to heart disease classification, we showed that FederatedHealth achieves 95–98% of centralized model performance while avoiding raw data exchange and guaranteeing HIPAA and GDPR compliance. The adoption hurdle is lowered for hospitals, clinics, and research networks alike because this performance is achieved without requiring participants to have sophisticated computing competence.

7. FUTURE ENHANCEMENT

A scalable basis for the upcoming generation of privacy-preserving health AI is established by FederatedHealth's successful design and validation. We see a number of significant extensions in the future: Multimodal Expansion: Combining federated feature extractors with modality-specific preprocessing procedures to integrate support for high-dimensional data types, such as whole-genome sequencing, 3D medical imaging (CT/MRI), and continuous biosensor streams. Advanced Cryptographic Privacy: Using fully homomorphic encryption (FHE) to compute on encrypted model updates without decryption for extremely sensitive use cases (such as genetic or psychiatric risk prediction). Dynamic Privacy Budgeting: Creating adaptive differential privacy methods that optimize the privacy-utility trade-off in real-time by allocating ϵ according to participant trust levels, model utility, and data sensitivity. Real-World Deployment at Scale: Collaborating with academic medical centers and national health networks, prospective, multi-institutional studies are being conducted to assess clinical impact, workflow integration, and regulatory audit preparedness. Synergy with Emerging Infrastructures: Investigating interaction with edge computing for real-time federated inference in ambulatory or intensive care unit settings, as well as with blockchain-based audit trails for immutable logging of model updates and consent records.

Conflict of interest statement

Authors declare that they do not have any conflict of interest.

REFERENCES

- [1] Natikar, S. H., & Sasi, S. (2020). FAST AND FLEXIBLE DENOISING NETWORK USING NOISE BASED PREDEFINED LAYERS BASED ON IMAGE DENOISING.
- [2] Sastry, G. S., & Sasi, S. Application Of ECC And ECDSA For Image With Error Control Technique Using RS Code.
- [3] Biradar, S., & Sasi, S. (2018, July). Design and Implementation of Secure and Encoded Data Transmission Using Turbo Codes. In 2018 9th International Conference on Computing, Communication and Networking Technologies (ICCCNT) (pp. 1-7). IEEE.
- [4] Sasi, S., & Swarna Jyothi, L. (2019). A Novel Public Key Crypto System Based on Bernstein Polynomial on Galois Fields 2^m to Secure Data on CFDP. In Smart Intelligent Computing and Applications: Proceedings of the Second International Conference on SCI 2018, Volume 2 (pp. 639-647). Springer Singapore.
- [5] Ghaleb, A. A., Sasi, S., & Aswatha, A. R. (2018, May). Design and implementation of satellite image encryption by using ecc. In 2018 3rd IEEE International Conference on Recent Trends in Electronics, Information & Communication Technology (RTEICT) (pp. 1438-1443). IEEE.
- [6] Pawan Kumar, V., Aswatha, A. R., & Sasi, S. (2018). Grayscale Image Encryption Based on Symmetric-Key Latin Square Image Cipher (LSIC). In Computational Vision and Bio Inspired Computing (pp. 476-487). Cham: Springer International Publishing.
- [7] Sasi, S., & Jyothi, L. S. (2016, October). Robust public key cryptosystem for space data communication. In 2016 International Conference on Communication and Electronics Systems (ICCES) (pp. 1-5). IEEE.
- [8] Periyasamy, R., Sasi, S., Malagi, V. P., Shivaswamy, R., Chikkaiah, J., & Pathak, R. K. (2025). Artificial intelligence assisted photonic bio sensing for rapid bacterial diseases. *Zeitschrift für Naturforschung A*, (0).
- [9] Raghuvanshi, P. (2024). Integrating generative ai into iot-based cloud computing: Opportunities and challenges in the united states. *Journal of Artificial Intelligence General science (JAIGS)* ISSN: 3006-4023, 5(1), 451-460.
- [10] Raghuvanshi, P. (2024). AI-Driven Identity and Financial Fraud Detection for National Security. *Journal of Artificial Intelligence General science (JAIGS)* ISSN: 3006-4023, 7(01), 38-51.
- [11] Kasoju, A., & Vishwakarma, T. (2025). The Role of Continuous Integration and Deployment in Improving Software Quality. *American Journal of Multidisciplinary Research & Development (AJMRD)*, 7(05), 19-30.
- [12] Kasoju, A., & chary Vishwakarma, T. (2024, November). Leveraging Explainable AI and Reinforcement Learning for Enhanced Transparency in Adaptive Fraud Detection. In 2024 IEEE 8th Conference on Energy Internet and Energy System Integration (EI2) (pp. 103-108). IEEE.
- [13] Shakibaie, B., Sabri, H., Abdulqader, H., Joit, H. J., & Blatz, M. B. (2024). Peri-implant soft tissue volume changes after microsurgical envelope technique with a connective tissue graft: A 5-year retrospective case series. *International Journal of Esthetic Dentistry*, 19(2).
- [14] Shakibaie-M, B. (2008). Microscope-guided external sinus floor elevation (MGES)—a new minimally invasive surgical technique. *IMPLANTOLOGIE*, 16(1), 21-31.

- [15] Karakolias, S. (2024). Mapping data-driven strategies in improving health care and patient satisfaction.
- [16] Shakibaie, B., Nava, P., Calatrava, J., Blatz, M. B., Nagy, K., & Sabri, H. Impact of Two Implant-Abutment Connection Types on Crestal Bone Stability: A 3-Year Comparative Split-Mouth Clinical Trial. *The International journal of periodontics & restorative dentistry*, 1-22.
- [17] Karakolias, S., & Iliopoulou, A. (2025). Health-Related Quality of Life and Psychological Burden Among and Beyond Children and Adolescents With Type 1 Diabetes: A Family Perspective. *Cureus*, 17(4), e81744.
- [18] Shakibaie, B., Blatz, M. B., & Abdulqader, H. (2025). The Microscopic and Digital One-Day-Dentistry Concept: A Minimally Invasive Chairside Technique. *Compendium of Continuing Education in Dentistry* (15488578), 46(6).
- [19] Karakolias, S., Georgi, C., & Georgis, V. (2024). Patient Satisfaction With Public Pharmacy Services: Structural and Policy Implications From Greece. *Cureus*, 16(4).
- [20] Shakibaie, B., & Barootch, S. (2023). Clinical comparison of vestibular split rolling flap (VSRF) versus double door mucoperiosteal flap (DDMF) in implant exposure: a prospective clinical study. *International Journal of Esthetic Dentistry*, 18(1).
- [21] Kasoju, Apoorva & Vishwakarma, Tejavardhana. (2024). THE ETHICS OF AI DECISION-MAKING: BALANCING INNOVATION AND ACCOUNTABILITY. *International Journal of Science and Research Archive*. 10.30574/ijrsra.2024.12.2.1548.
- [22] Tao, Y., & Choi, C. (2022, May). High-Throughput Split-Tree Architecture for Nonbinary SCL Polar Decoder. In *2022 IEEE International Symposium on Circuits and Systems (ISCAS)* (pp. 2057-2061). IEEE.
- [23] Raghuwanshi, P. (2024). AI-Powered Neural Network Verification: System Verilog Methodologies for Machine Learning in Hardware. *Journal of Artificial Intelligence General science (JAIGS)* ISSN:3006-4023, 6(1), 39-45.
- [24] Tao, Y., & Kwong, J. (2017). U.S. Patent No. 9,793,923. Washington, DC: U.S. Patent and Trademark Office.
- [25] Sasi, S. finished her Bachelor of Engineering in Information Technology in the year 2005. She completed her Masters degree in Digital Communication and Networking in the year 2010. She was awarded doctorate degree due to her research in space communication from Visvesvaraya Technological University in the year 2019. Her research interest includes bioinformatics. Artificial Intelligence, Cloud computing, Cryptography, Bio medical Image processing to name a few.
- [26] Roopa, M. D., Sasi, S., Babu, S., Agrawal, A., Vinaik, K. P., & Patil, S. C. (2023, November). The IoT & MIMO Communication Technology: A New Computer Communication Approach. In *2023 International Conference on Communication, Security and Artificial Intelligence (ICCSAI)* (pp. 785-791). IEEE.
- [27] Sagari, S. M., Malagi, V. P., & Sasi, S. (2024). Euri-A Deep Ensemble Architecture For Oral Lesion Segmentation And Detection. *Int. J. Intell. Syst. Appl. Eng*, 12(3s), 242-249.
- [28] Rajeswari, P., & Sasi, S. (2024). Efficient k-way partitioning of very-large-scale integration circuits with evolutionary computation algorithms. *Bulletin of Electrical Engineering and Informatics*, 13(6), 4002-4007.
- [29] Aswatha, A. R., Sasi, S., Santhosh, B., Mehta, D., & Babuprasad, S. (2019, October). Design and implementation of unreliable CFDP protocol over elliptic curve cryptography. In *Smart Intelligent Computing and Applications: Proceedings of the Third International Conference on Smart Computing and Informatics, Volume 2* (pp. 627-638). Singapore: Springer Singapore