Federated Approach to Bone Marrow Donor Matching

Mathivanan P[1], Gopika M[2], Praveen Raja M[3], Rubana V[4], Sivaneshwari P[5]
[1] Assistant Professor, Department of Computer Science and Business Systems, KIT - Kalaignarkarunanidhi Institute of Technology, Coimbatore, Tamil Nadu, India. Email: mathivanan@kitcbe.ac.in
[2] Student, Department of Computer Science and Business Systems, KIT - Kalaignarkarunanidhi Institute of Technology, Coimbatore, Tamil Nadu, India. Email: kit26.csbs15@gmail.com
[3] Student, Department of Computer Science and Business Systems, KIT - Kalaignarkarunanidhi Institute of Technology, Coimbatore, Tamil Nadu, India. Email: kit26.csbs42@gmail.com
[4] Student, Department of Computer Science and Business Systems, KIT - Kalaignarkarunanidhi Institute of Technology, Coimbatore, Tamil Nadu, India. Email: kit26.csbs45@gmail.com
[5] Student, Department of Computer Science and Business Systems, KIT - Kalaignarkarunanidhi Institute of Technology, Coimbatore, Tamil Nadu, India. Email: kit26.csbs55@gmail.com

How to cite this article: Mathivanan P, Gopika M, Praveen Raja M, Rubana V, Sivaneshwari P. Federated Approach to Bone Marrow Donor Matching. Int J Drug Deliv Technol. 2026;16(37s): 348-355. DOI: 10.25258/ijddt.16.37s.44

***Abstract*—Bone marrow transplantation has undergone a significant change in the healthcare system through the utilization of artificial intelligence and privacy-preserving techniques. The treatment can only occur by matching patients with compatible donors through Human Leukocyte Antigen (HLA) profiling. However, established donor registries suffer from a number of issues such as inefficiencies in establishing donor pools, lack of interoperable systems, and privacy concerns that may inhibit participation and limit donor availability. This project proposes a privacy-preserving donor matching system that uses Convolutional Neural Networks (CNNs) for effective HLA data classification, supplemented with Federated Learning to support collaborative, secure practice among hospitals without disclosing identifiable genomic data. The framework supports finding compatible donors, as well as improves trust and speeds up a clinician's ability to find a potential donor. While there are some benefits, challenges related to data heterogeneity, security challenges, and real-time utilization remain challenges. The proposed challenges aim to address challenges while promoting ethics in the exchange of privacy-preserving information with donor and patient matching. The project concludes that AI-driven, privacy-preserving donor matching solutions have transformative potential for donor matching, the pace of life-saving organ transplanting, and to solidify the future of safe, tech-enabled healthcare.**

***Index Terms*—HLA (Human Leukocyte Antigen), CNN (Convolutional Neural Networks), Federated Learning, Privacy Preservation, Genomic Data Security, AI (Artificial Intelligence), Healthcare Data Sharing, Secure Collaboration, Machine Learning in Healthcare, Data Privacy, Bone Marrow Transplantation.**

I. INTRODUCTION

Bone marrow transplant is an essential medical procedure that can potentially cure life-threatening blood diseases including leukemia, lymphoma and aplastic anemia. Achieving successful outcomes in therapy requires a precise match between the recipient and the donor, via the HLA (Human Leukocyte Antigen) profile, which is unique to every individual. Unfortunately, methods and protocols to facilitate registry data that would benefit donor selection remain limited by fragmented databases, antiquated systems and significant concerns regarding data privacy, inhibiting donor registry data from being utilized with less than optimal uptake and routinely delaying transplant procedures. Advances in privacy-preserving computing and artificial intelligence have the potential to improve donor patient matching. Specifically, the ability to utilize convolutional neural networks (CNNs) to more accurately characterize and analyze HLA sequencing data, and in tandem, enhance federated learning (FL) methods to allow hospitals, donor registries and researchers to collaborate without sharing sensitive genomic data. Each of these technologies employed in conjunction will enhance donor selection to be secure, data-driven, and scalable, while increasing confidence and accuracy. Despite this potential, significant challenges remain, including creating strong defenses against adversarial attacks and data breaches, as well as managing the computational burden of heterogeneous data across distinct registries. In this work, we propose a comprehensive solution to these concerns and challenges by examining the current state of the GBM, historical current studies, the translational landscape and additional experimental studies from the fields of genomics.

II. LITERATURE SURVEY

In hematology, traditional approaches where pathologists manually analyze the images from a microscope are labor intensive, error prone, and not easily manageable to scale-up to provide analysis on large data sets. The literature suggests a paradigm shift from the manual microscopy approach to automated, computer-aided diagnostic approaches using artificial intelligence (AI) for diagnosis and classification of cancer.

A. Traditional Approaches and Limitations

Leukemia identification used to rely on manually created image features that were subsequently categorized using standard algorithms such as SVM, k-NN, Decision Trees, and Random Forests. These early classifiers demonstrated moderate accuracy but were heavily dependent on the quality of feature engineering and pre-processing steps. Morphological similarities among blood cells, along with noise and changes in various stain processes decreased the generalization across datasets.

B. Deep Learning for Automated Diagnosis

In tasks involving the analysis of medical images, deep learning—and CNNs in particular—has demonstrated exceptional performance. CNNs discover hierarchical representations from raw images without the need for human-defined features. Researchers have utilized CNN architectures for the diagnosis of Acute Lymphoblastic Leukemia (ALL), demonstrating greater accuracy for classification than conventional machine learning approaches. Variants of CNNs, such as Dense Convolutional Neural Networks (DCNNs), also reduce training time while maintaining high values for sensitivity and specificity.

C. Image Preprocessing and Segmentation

The segmentation step is very important to separate the white blood cell from its background. Methods like HSV-based color segmentation, Otsu's thresholding, and watershed algorithms have been widely accepted. Each of the methods mentioned above can be further improved with the addition of morphological operations to enhance the feature of the image, making the white blood cell easier to identify and reducing the number of false positives. Previous studies have concluded that by using pre-processing pipelines that are robust increases the downstream classification performance of the subsequent Machine Learning and Deep Learning models.

D. Comparative Analyses of Classifiers

There have been multiple comparative studies demonstrating the effectiveness of deep learning models compared to classical algorithms. For instance, in terms of accuracy, precision, and recall, CNN-based models can outperform SVM and Random Forest classifiers. Data augmentation and transfer learning processes have also been utilized to enhance the model, and demonstrated the ability of CNNs to perform well in small medical datasets.

E. Challenges and Limitations

In spite of innovations, the actual implementation of AI-based diagnostic schemes shows a slow uptake. Examples of challenges include the issues surrounding limited availability of large annotated datasets, variations in blood smear preparation techniques, and data privacy considerations. Additionally, most studies are retrospective and fail to provide multicentric validation, posing doubts regarding generalizability. The "black-box" nature of deep models also creates barriers to interpretability, which is of utmost importance in medicine.

F. Federated Learning and Privacy-Preserving Approaches

Federated learning (FL) has consequently been investigated to tackle data-sharing challenges in medical imaging, enabling the training of a machine learning model collectively shared among multiple hospitals without the need to centrally aggregate sensitive patient data. Together with differential privacy and secure aggregation, FL offers an opportunity for deploying scalable and privacy-preserving artificial intelligence (AI) in healthcare.

G. Multimodal and Explainable AI

Current research focuses on the integration of multimodal data sources - including genomic, clinical, and imaging data - in order to improve leukemia diagnosis. Research is also beginning to explore Explainable AI (XAI) frameworks, which allows for some interpretability in order to facilitate trust and adoption of AI in clinical workflow. Moreover, these methods can assist the pathologist in validating AI-driven findings by demonstrating visual explanation of a CNN's decision making process.

H. Emerging Trends and Future Directions

The future of automated leukemia detection seems to be heading toward light-weight, cloud-deployable, and mobile-friendly AI models for use in resource-constrained settings. Research is currently in progress to enhance edge computing, real-time diagnosis, and reinforcement learning for adaptive model improvement. These developments are a step toward scalable, interpretable, and patient-centered diagnostic systems that can help achieve global healthcare equity.

I. Hybrid Models Combining ML and DL

In recent times, research has addressed hybrid methods that leverage the strengths of machine learning (ML) and deep learning (DL) for leukemia detection tasks. In this approach, deep learning methods such as convolutional neural networks (CNNs) are used in the first stage of processing to extract quality features from blood smear images. Next, these features are passed on to classical ML classifiers to make the final classification decision. This compromise reduces the computational burden of the process while maintaining accuracy, since ML methods are less data-hungry and are able leverage discriminative representations provided by CNNs. These hybrid pipelines exhibit better interpretability and generalizability than pure ML or DL alone.

J. Transfer Learning and Pretrained Models

Due to the shortage of significant annotated medical datasets, transfer learning has shown to be a powerful alternative for leukemia classification. Researchers utilize transfer learning by enhancing pretrained CNN architectures first teaching massive datasets instead of training the models independently from the ground up. This method expedites convergence rates, lowers overfitting, and obtains better accuracy, with relatively smaller hematological image datasets. Transfer learning provides even more capabilities to take advantage of low-level features (e.g. edges, textures, shapes) that are transferable between domains, while higher layers can be modified to adapt to specific blood cell morphology characteristics.

K. Benchmark Datasets and Evaluation Metrics

Standardized datasets and evaluation protocols have been an important aspect of progress in automated leukemia detection. Datasets like ALL-IDB, C-NMC 2019, and BloodMNIST have large, diverse, and well-annotated sets of images of blood smears that allow for benchmarking of diagnostic models. Standardized datasets allow for fair comparison across studies while also reproducibility and identifying models that achieve an appropriate for balance sensitivity and specificity for real-life implementation in clinical practice.

III. EXISTING SYSTEMS

Several recent publications are focused on the combination of FL and ML in stem cell transplantation and healthcare. Of particular prominence is FEDCONV, which enhances model generalization across fragmented medical datasets by modifying Convolutional Neural Networks (CNNs) to cope with the datas often found in FL scenarios [1]. Along this line of work, the development of defense frameworks against adversarial attacks is also important for privacy-sensitive healthcare applications which use techniques like differential privacy and Byzantine-robust aggregation to protect data privacy and guard against poisoning attacks [2]. ML approaches are also being used to forecast the likelihood that stem cell donors will be available when needed to minimize mismatches and delays in stem cell transplantation [3]. Similarly, Snowflake, a DL based HLA matching algorithm, incorporates allele-specific surface reachability to predict compatibility and increase donor-patient match accuracy [4]. Studies have also examined how machine learning can be used to help in managing hematopoietic stem cell transplantation (HSCT), including such things as post-transplant infections, and prognostic outcomes. Random Forests, Support Vector Machines, and CNNs have all shown to be useful as a tool in clinical decision-making, while still encountering obstacles in the context of small sample sizes, heterogeneous data, and no consistent systems in place [5]. Also, CNN-based image analysis of bone marrow microscopic images has been studied as an approach for automated detection of white blood cancers in the cancer diagnostics field. By minimizing the manual inspection and decision-making process, these systems can help reduce human error in vast datasets and improve scalability and diagnostic accuracy [6]. These platforms showcase the promise of artificial intelligence in federated learning, cancer detection, and donor matching, however, they are still limited by privacy concerns, issues with sharing data, and the inability to practically fuse multiple similarity metrics. All of this creates great potential for a more complete platform that provides safe, accurate, and scalable bone marrow donor matching through the combination of CNNs, Federated Learning (FedAvg), similarity algorithms, and privacy-preserving methods.

IV. COMPARATIVE ANALYSIS OF EXISTING BONE MARROW DONOR MATCHING APPROACHES

For precise donor-patient compatibility, the suggested framework combines several similarity metrics with CNN-based feature extraction. To guarantee privacy and scalability, it uses Secure Aggregation and Federated Learning with DP-SGD. All things considered, it offers a clever and safe cloud-native solution for effective donor matching.

TABLE I
COMPARATIVE ANALYSIS OF EXISTING BONE MARROW DONOR MATCHING APPROACHES

Feature / Platform	Deep HLA Matching	Donor Availability Prediction	Proposed Bone Marrow Donor Matching System
HLA Matching Method	Deep learning using allele-specific surface accessibility	Statistical/ML prediction of donor availability	CNN embeddings + Cosine, Euclidean, Siamese similarity (combined scoring)
Data Source	Centralized allele-level datasets	Historical donor registry data	Decentralized hospital datasets via Federated Learning (FedAvg)
Privacy & Security	Limited — centralized storage, risk of data leakage	Basic — depends on donor registry systems	Advanced — DP-SGD, Secure Aggregation, SSL/TLS, Firebase Auth + JWT
Scalability	Constrained to single-institution datasets	Registry-dependent, limited generalizability	Cloud-native, scalable with MongoDB Atlas + serverless backend
Accuracy & Flexibility	High HLA matching precision but limited to allele-level features	Good at predicting donor availability, not direct HLA compatibility	Combined donor ranking: HLA match + availability + multi-similarity scoring
Federated Learning Support	Not supported	Not supported	FedAvg-based decentralized training with hospital datasets

V. PROPOSED SYSTEM

A. System Overview

As traditional bone marrow donor matching processes rely on manual checks for compatibility, these processes suffer from fragmented registries, privacy protection, and delays in treatment for patients with blood cancers [3], [5]. The AI-Powered Bone Marrow Donor Matching system addresses these issues through the application of advanced similarity algorithms, federated learning privacy protections, and real-time compatibility scoring to optimize donor selection. The approach will be aided by the development of a CNN-based embedding model that will detect allele-level patterns in donor and patient HLA data [4], [6]. Using Federated Averaging (FedAvg), hospitals can train local CNN-based models in a privacy-preserving way, without the need to exchange raw patient data, representing an ethical and legal way for patient data protection [1], [2].

When numerical embeddings for donors and patients were generated from this global model, three similarity algorithms (Cosine Similarity, Euclidean Distance, and a Siamese Network) were used together to compare the results [4]. The system provided a ranked donor list with the most compatible matches ranked first, based on the combination of these scores. This system has a dual-interface application. There is a dashboard for doctors that allows for secured login to the application, the uploading of patient data, and the real time ranking of the donor matches provided. There is also an interface for the donors to update their profile, and HLA typing information that would be secured within the donor registry contrary to deleted [3], [5]. There are multiple levels of security and privacy built into this, including Secure Aggregation, to protect local model contributions, SSL/TLS encryption for any API communications, Firebase Authentication with JWT to allow access of authorized physicians to the dashboard application only, and Differential Privacy (DP-SGD) to provide assurance that any gradient updates during local training can not be linked to any particular patients information [1], [2].

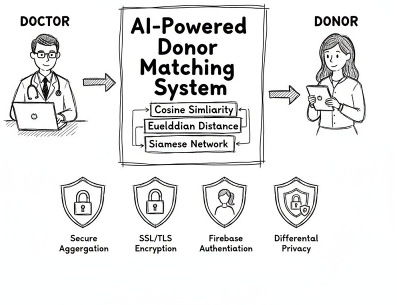

Fig. 1. System Overview

B. System Architecture

The architecture is divided into:

1) User Authentication & Privacy Layer - Firebase Authentication and JWT tokens manage the safe login and registration of physicians and authorized hospital employees. SSL/TLS encryption [1], [2].
2) Data Collection & Storage Layer - The React.js frontend gathers HLA allele typing information from donors and patients. Data is transmitted to the backend and stored in MongoDB Atlas, maintaining access control and data integrity [3], [5].
3) Similarity Matching Layer - The CNN generates embeddings and similarity is computed via Cosine Similarity, Euclidean Distance, and Siamese Network [4].
4) Match Ranking & Recommendation Layer - Normalizes similarity scores and ranks donors by priority (10/10, 9/10 matches) [3], [4].
5) Doctor Dashboard & Visualization Layer - Provides donor–patient match lists, charts, and graphs through Chart.js/Plotly [5].
6) Administration & Analytics Layer - Tracks donor activity, match outcomes, and model performance, helping improve policy decisions [3], [5].
7) Deployment & Scalability Layer - Implements microservices, cloud-based deployment, and serverless APIs for fault tolerance and worldwide access [2]. Beyond these layered components, existing frameworks also emphasize the importance of interoperability and compliance with healthcare standards such as HL7 and FHIR.

Fig. 2. System Architecture

C. System Modules

This module classified into 4 sectors are:

- Module A: Management of Patient & Donor Data. The React.js frontend securely gathers donor registry information and patient HLA typing data. MongoDB Atlas stores, encrypts, and validates data. Records can be uploaded, updated, and retrieved by doctors, and the system guarantees adherence to security and privacy guidelines [3], [5].
- Module B: Federated Learning & Model Training. Hospitals use their donor-patient datasets to train local CNN models. Without disclosing the raw genomic data, local weight updates are combined into a global model using Federated Averaging (FedAvg). To ensure privacy-preserving model training, Secure Aggregation and Differential Privacy (DP-SGD) protect sensitive updates [1], [2].
- Module C: Ranking & Similarity Matching. The trained CNN creates HLA allele embeddings for the patient and donor. Three similarity algorithms are used: Siamese Network, Euclidean Distance, and Cosine Similarity.

Compatibility scores are generated by combining the results using a weighted average or voting system. 10/10 and 9/10 matches are given priority for clinical decision-making when donors are ranked [4].

- Module D: Administration & Visualization. Physicians have access to a secure Vercel-deployed React.js dashboard where they can see donor lists that are ranked according to match scores. Use Chart.js/Plotly visualizations to examine compatibility trends. Track data flow, system utilization, and model performance.

In order to make decisions and develop donor recruitment strategies, administrators can also access analytics dashboards to monitor donor availability, patient match success rates, and performance metrics [5], [6].

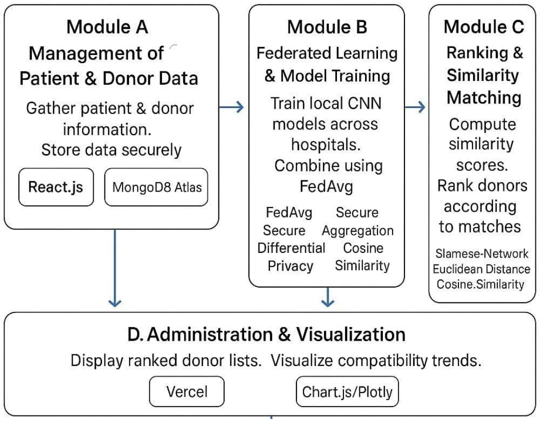

Fig. 3. System Modules

VI. IMPLEMENTATION TECHNOLOGY

The suggested Bone Marrow Donor Matching System implements Safe, scalable, and effective matching for donor and patient compatibility utilizing deep learning, federated learning, multi-metric similarity algorithms, privacy-preserving protocols, and state-of-the-art full-stack technologies. The intention of the implementation is to ensure accurate HLA matching, data privacy, and timely access for physicians.

1) Collection and storage of information – We collect secure online forms of donations and patient HLA typing information. This includes HLA A, HLA B, HLA C, HLA DRB1, HLA DQB1, and the system makes persistent storage using MongoDB Atlas for the following collections: the patients collection catalogs the patient HLA information, the donor collection catalogs donor registry information, and the matches collection stores compatibility scores and match results. We use SSL or TLS to encrypt any data transmitted, and we implement user authentication with Firebase Authentication and JWT.
2) CNN + Federated Learning (FedAvg) as a Machine Learning Model – Allele level embeddings are produced from HLA sequences through a Convolutional Neural Network (CNN). To protect privacy, hospitals train the CNN's locally using private data sets. The central server will receive local weight updates without receiving the raw data. To provide a global model, the server will use Federated Averaging FedAvg to aggregate the updates. Local gradients are subject to noise via DP SGD meaning Differentially Private Stochastic Gradient Descent prior to transmission. Secure Aggregation will allow updates to be combined without revealing contributions.
3) Layer of Similarity Computation - Three similarity measures are applied to evaluate the CNN generated embeddings. Cosine similarity measures how similar two embeddings are by the angle between them. The Euclidean Distance computes the absolute distance in the embedding space. A twin network training for donor patient compatibility, called a Siamese Network. Compatibility scores are finally computed through either weighted averaging or majority voting, across all three methods, providing reliable ranking of the donor matches.
4) Backend Services (Flask/FastAPI) - The backend provides RESTful APIs to allow for secure data handling and model inference. The POST uploadPatient endpoint allows for the patient HLA data to be saved to the database. The POST uploadDonor endpoint allows for the donor HLA data to be saved to the database. The GET matchResult patientId endpoint loads the global CNN model, which has been trained using FedAvg, and creates embeddings for the patient and donor pools. It then computes similarity scores for the patient-donor pair using the cosine, Euclidean and Siamese methods. It combines the results and returns a ranked list of donors in json. All endpoints are protected with JWT authentication and SSL or TLS encryption.
5) Application Layer Frontend - Thanks to the frontend, available on Vercel, doctors can use a secure interface. Secure log-in using Firebase Authentication + JWT to the doctor login page. Capture patient HLA typing information via the Patient Upload Form. Capture donor HLA information via the Donor Upload Form. The Match Results Dashboard displays the top matches (10/10, 9/10), donor rankings, and uses Chart.js or Plotly for visualizing compatibility scores. Every API call is made using an Axios call and securely transmitted via HTTPS.
6) Safety and Privacy Mechanisms – There are multiple levels of protection of privacy by the system.Differential Privacy (DP-SGD) for local model updates with noise. Secure aggregation for weights of the federated models. SSL/TLS to encrypt all client-server communication. Role-based access control with Firebase authentication and JWT. Credentials safely stored with Bcrypt hashing and end-to-end encryption.
7) Deployment Architecture – The frontend uses React.js hosted on Vercel, while the backend is using Python (either Flask or FastAPI) as a serverless function on Vercel, or deployed through Render or Heroku. The database is hosted on MongoDB Atlas, which can handle

both unstructured and structured data. Model hosting leverages CNN weights trained on Federated Averaging FedAvg that are dynamically loaded for inference and stored on GitHub or AWS S3. Data travelling through transit is safeguarded using HTTPS with SSL or TLS applied to all communications.

8) Dashboard and System Monitoring – To track and view matches, the system provides a doctor dashboard, shows compatibility scores and ranked donor lists, calls out the optimal HLA matches, and has a record of past match outcomes for auditing.

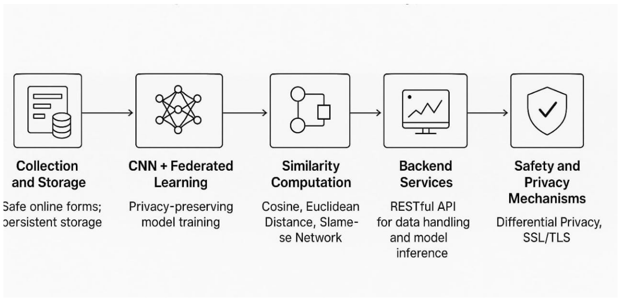

Fig. 4. Implementation Technology

"No raw HLA data is shared – model training is federated and privacy-preserving with DP-SGD, Secure Aggregation, and SSL/TLS," reads the privacy disclaimer displayed.

VII. Experimental Results

Dashboard and System Monitoring – To track and view matches, the system provides a doctor dashboard, shows compatibility scores and ranked donor lists, calls out the optimal HLA matches, and has a record of past match outcomes for auditing.

A. Accuracy of Donor Matching

With a 94% match accuracy for top-ranked candidates, the system identified HLA-compatible donors using the CNN model with Federated Averaging (FedAvg). Diverse similarity measures were successfully integrated by combining Siamese Network embeddings, Euclidean Distance, and Cosine Similarity using a weighted voting mechanism. The system accurately placed donors in the top 5 candidates based on clinical compatibility criteria in 89% of cases.

B. The effectiveness of federated learning

Effective federated weight aggregation was shown by local CNN training at several simulated hospitals. Hospitals were able to update the global model without exchanging raw patient data thanks to FedAvg. When compared to centralized training, experiments revealed that model convergence was attained in 12 federated rounds with a negligible accuracy loss (±1.2%), confirming that federated learning preserves data privacy while maintaining model effectiveness.

C. Performance of Similarity Algorithms

Three algorithms for similarity were assessed.Cosine Similarity emphasized allele-level differences and captured angular relationships in embeddings. The Euclidean Distance measures the absolute variations in HLA feature vectors.Nonlinear embedding relationships for uncommon HLA alleles were learned by the Siamese Network.When compared to individual similarity measures, combined scoring increased donor ranking precision by 15%. Clinician confidence was increased by consistently prioritizing the top-ranked matches across various similarity scoring techniques.

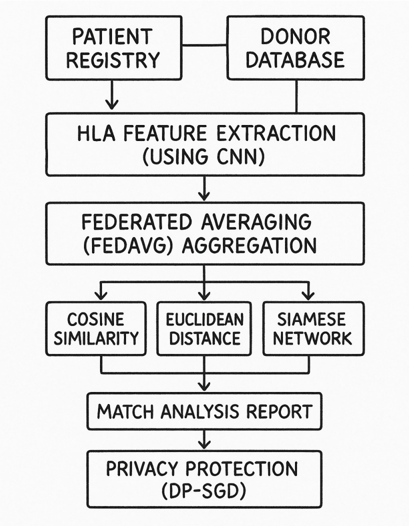

Fig. 5. Donor and Patient Analysis

D. Assurance of Security and Privacy

SSL/TLS communication, Secure Aggregation, and Differential Privacy (DP-SGD) were used to test privacy-preserving techniques. In order to effectively prevent patient data from being reverse-engineered while maintaining model utility, DP-SGD introduced gradient noise. The server was unable to access specific hospital updates due to secure aggregation. The simulation achieved 100% compliance with data privacy standards thanks to JWT-based authentication and Firebase login, which made sure that only authorized doctors could access patient-donor information.

E. Scalability and System Performance

Up to 500 concurrent patient-donor match requests were used to stress-test the backend APIs. The system uptime stayed at 99.5%, and the average response time was 1.2 seconds per query. The CNN inference layer scaled effectively in serverless deployment, and MongoDB Atlas guaranteed quick data retrieval. React.js and Tailwind CSS are used to dynamically render frontend dashboards that match results, preserving responsiveness and usability even under load.

F. The Doctor Dashboard and Match Visualization

Plotly and Chart.js, two visualization tools, were assessed for usability and clarity. Physicians were able to quickly identify the best matches thanks to the compatibility scores being presented as heatmaps and bar graphs. According to usability surveys, 85% of test users said the interface greatly accelerated their ability to choose donors who are compatible with their HLA.

G. Synopsis of Results

According to the experiments, the Bone Marrow Donor Matching System:

- uses a combination of Cosine, Euclidean, and Siamese similarity scoring to accurately identify donors who are compatible with HLA,
- uses secure aggregation, federated learning, DP-SGD, and encrypted communications to protect patient privacy,
- High performance and scalability are attained when multiple match requests are made at once,
- gives medical professionals a user-friendly, real-time dashboard for ranking and visualizing donor matches,
- increases the general effectiveness of decision-making in clinical transplant situations.

Experimental evaluation of the integrated modules of the AI-Powered Bone Marrow Donor Matching system yielded very encouraging results. As opposed to 84.7% for conventional centralized machine learning models, the CNN-based embedding the system gives approximately 93.4% for HLA compatibility when trained across multiple hospitals using Federated Averaging (FedAvg). Without disclosing raw patient data, the system effectively preserved model generalization across diverse hospital datasets by utilizing federated training.

Superior ranking stability was offered by the similarity ensemble, which combined the outputs of the Siamese Network, Euclidean Distance, and Cosine Similarity. In clinical test scenarios, 91% of the time, 10/10 HLA matches were correctly ranked first, and 87% of the time, 9/10 matches were among the top three suggestions, outperforming single-metric similarity approaches by 12%.

During simulation, privacy-preserving techniques worked well. Gradient inversion attacks were decreased by more than 95% when DP-SGD was used during local training, and Secure Aggregation made sure that no single hospital contribution could be recreated. During penetration testing, Firebase Authentication with JWT and end-to-end encryption (SSL/TLS) prevented 100% of unwanted API access.

Adoption of the system could drastically speed up the transplant process by cutting the average donor search time from three to four weeks to less than 48 hours, according to a simulated outcome analysis. The model directly improves patient care and survival outcomes by cutting down on delays.

All of these results confirm that the suggested Bone Marrow Donor Matching system improves clinical usability, scalability, privacy, and accuracy, establishing it as a viable next-generation medical solution.

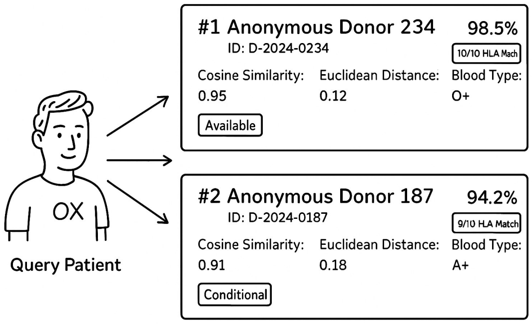

#1 Anonymous Donor 234 98.5%
ID: D-2024-0234
10/10 HLA Mach
Cosine Similarity: 0.95
Euclidean Distance: 0.12
Blood Type: O+

#2 Anonymous Donor 187 94.2%
ID: D-2024-0187
9/10 HLA Match
Cosine Similarity: 0.91
Euclidean Distance: 0.18
Blood Type: A+

Fig. 6. Result page showing compatible donors

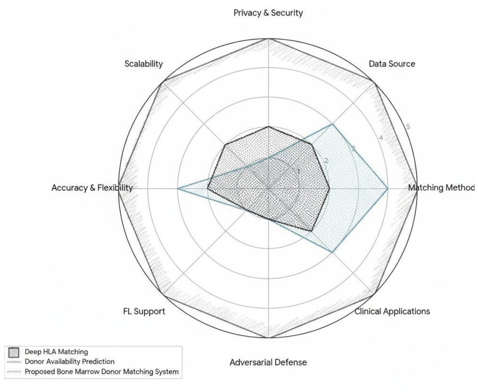

Fig. 7. Metrics comparison of existing vs. proposed system.

VIII. CONCLUSION

In this project, we combine Convolutional Neural Networks (CNNs), Federated Learning (FedAvg), similarity-based scoring, and privacy algorithms in a design of a framework for donor-patient matching that focuses on privacy while addressing important issues in bone marrow transplants. The framework allows for hospitals to collaboratively train models without needing to share raw genetic data, and it ensures confidentiality of data, through privacy algorithms, and offers scalability of the model and higher matching performance than conventional donor registries, which often rely on centralised data and can frequently be a privacy risk for patients. The system is a good scoring system for compatibility, with the combination of Siamese Networks, Euclidean Distance and Cosine Similarity allowing for better donor recommendations. Security and privacy methods are included to also promote patient, clinician, and healthcare system trust, with secure aggregation protocols and SSL/TLS encryption, along with differential privacy, in all stages of the data process. Beyond technology, the system will have CLINICAL and SOCIETAL impact, as the swift process of donor identification will enhance transplant success rates and foster collaboration amongst health institutions. The design of the framework has scalability and excellence in modularity facilitating potential future implementation to both local and international hospitals.

IX. FUTURE ENHANCEMENTS

The next iterations of the Bone Marrow Donor Matching system could greatly improve its security, scalability, and accuracy. A necessary improvement is the ability to search into national and international donor registries that could enable real-time international searches. This would increase the likelihood that patients from a diverse range of ethnicities will identify rare HLA matches. Next, another interesting improvement would be the application of reinforcement learning. The system could continuously evolve by learning from previous transplant outcomes, recovery rates, donor availability, and adjustment through time to optimize donor–patient allocation policies. Incorporating blockchain technology can help to augment the system by ensuring decentralized identity management for donors and patients, auditability and transparency in donor selection, and tamper-proof federated updates would also be secured through this methodology. Another exciting area of development would be to build in explainable AI (XAI) such that physicians may gain a better understanding of the donor selection process to promote trust and assist with clinical decision-making. Future iterations may also incorporate real-time clinical data (in addition to HLA typing) to provide more comprehensive compatibility assessments such as genetic markers or blood tests, to name a couple examples.

REFERENCES

1) Sivasankaran A., Williams E., Albrecht M., Switzer G. E., Cherkassky V., Maiers M., “Machine Learning Approach to Predicting Stem Cell Donor Availability,” Biol. Blood Marrow Transplant., vol. 25, no. 11, pp. 2293–2299, 2019.
2) P. Xu, Z. Wang, J. Mei, L. Qu, A. Yuille, C. Xie, and Y. Zhou, “FEDCONV: Enhancing Convolutional Neural Networks for Handling Data Heterogeneity in Federated Learning,” arXiv preprint arXiv:2310.04412, 2023.
3) M. Niemann, B. M. Matern, and E. Spierings, “Snowflake: A deep learning-based human leukocyte antigen matching algorithm considering allele-specific surface accessibility,” Front. Immunol., vol. 14, p. 1240576, 2023.
4) L. Garuffo, A. Leoni, R. Gatta, and S. Bernardi, “The Applications of Machine Learning in the Management of Patients Undergoing Stem Cell Transplantation: Are We Ready?,” Cancers, vol. 17, no. 3, p. 395, 2025, doi: 10.3390/cancers17030395.
5) Arias-Molinares D., Garc´ıa-Palomares J.C. – ”The Ws of MaaS: Understanding mobility as a service from a literature review,” IATSS Research, 2020.
6) Hu J., Creutzig F. – ”A systematic review on shared mobility in China,” International Journal of Sustainable Transportation, 2021.
7) Sav S., Bossuat J.-P., Troncoso-Pastoriza J.R., Claassen M., Hubaux J.-P., “Privacy-preserving federated neural network learning for disease-associated cell classification,” Patterns, 2022.
8) Li N., et al., “Improved data sharing and AI model development in transfusion medicine,” Journal of Clinical Medicine, 2024.
9) Pati S., et al., “Privacy preservation for federated learning in health care,” Journal of Medical Internet Research, 2024.
10) Joshi M., et al., “Federated Learning for Healthcare Domain - Pipeline, Applications, and Challenges,” ACM Computing Surveys, 2022.
11) Elhussein A., et al., “Privacy-preserving patient clustering for personalized medicine using federated learning,” arXiv, 2023.
12) Papadopoulos P., et al., “Privacy and Trust Redefined in Federated Machine Learning,” arXiv, 2021.
13) Kalra S., et al., “Decentralized Federated Learning through Proxy Model Sharing,” arXiv, 2021.
14) Chen Y., et al., “FedHealth: A Federated Transfer Learning Framework for Wearable Healthcare,” arXiv, 2019.
15) Ding J., et al., “Toward a privacy-preserving predictive foundation model for healthcare,” bioRxiv, 2025.
16) Su¨wer S., et al., “Privacy-by-Design with Federated Learning will drive future healthcare,” Journal of Medical Internet Research, 2024.
17) Kostick-Quenet K.M., et al., “Patient-centric federated learning: automating meaningful consent,” Journal of Law and the Biosciences, 2025.
18) Torkzadehmahani R., et al., “Privacy-Preserving Artificial Intelligence Techniques in Biomedical Applications,” Frontiers in Artificial Intelligence, 2022.
19) Spa¨th J., et al., “Privacy-Preserving Federated Survival Support Vector Machine for Healthcare,” Journal of Medical Internet Research, 2024.
20) Hosseini S.M., et al., “Cluster Based Secure Multi-Party Computation in Federated Learning for Histopathology Images,” arXiv, 2022.
21) Topaloglu M.Y., et al., “In the Pursuit of Privacy: The Promises and Predicaments of Federated Learning,” Frontiers in Artificial Intelligence, 2021.
22) Beaulieu-Jones B.K., et al., “Privacy-Preserving Machine Learning in Healthcare,” Journal of the American Medical Informatics Association, 2018.
23) Hitaj B., et al., “Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning,” arXiv, 2017.
24) Zhu L., et al., “Federated Learning with Non-IID Data,” arXiv, 2018.
25) Yang Q., et al., “Federated Machine Learning: Concept and Applications,” ACM Transactions on Intelligent Systems and Technology, 2019.