EARLY DETECTION OF BONE MARROW GRAFT REJECTION USING LARGE LANGUAGE MODELS

Mrs. Y. Naga Lavanya[1],Vinitha.G[2], Ajay.B[3], Anjan vivek.J[4], Shiva.G[5]

[1]Assistant Professor, Department of IT, TKR College of Engineering and Technology, Telangana, India
[2,3,4,5]B.Tech Students, Department of IT, TKR College of Engineering and Technology, Telangana, India

---***---

Abstract - *Bone marrow transplantation is a critical therapeutic intervention for patients with hematological disorders, yet graft rejection remains a significant cause of morbidity and mortality. Early detection of graft rejection is essential for timely medical intervention and improved patient outcomes. This work proposes an intelligent system leveraging Google Gemini 2.5 Flash Large Language Models (LLMs) via Lang Chain to predict the risk of bone marrow graft rejection. The system accepts patient-specific hematological parameters, cytokine levels, clinical notes, and pathology reports as input, and generates a rejection risk score, categorical rejection status, and a concise explanation. A Django-based web interface allows healthcare professionals to input patient data seamlessly and receive AI-generated predictions in real time. The model is trained and validated using a structured dataset containing hematological features and historical rejection outcomes. Experimental results demonstrate that the system effectively identifies early and severe rejection scenarios, enabling proactive clinical decision-making. This approach integrates natural language understanding with clinical data analytics, offering a scalable solution for personalized patient monitoring. By combining AI interpretability with ease of use, the proposed framework provides a robust platform for augmenting traditional clinical practices, reducing diagnostic latency, and enhancing overall transplant care quality.*

Keywords: Bone Marrow Transplantation (BMT), Graft Rejection, Early Detection, Large Language Models (LLMs), Clinical Decision Support System, Natural Language Processing (NLP), Electronic Health Records (EHR), Immunological Biomarkers.

1. INTRODUCTION

Bone marrow transplantation (BMT) is a life-saving procedure for patients suffering from hematological malignancies, immunodeficiency disorders, and other severe blood-related conditions [2,3]. Despite advancements in transplant techniques and immunosuppressive therapies, graft rejection remains a major complication, often leading to patient morbidity or mortality [4]. Early identification of graft rejection is critical for initiating timely interventions and improving patient survival rates [2,4]. Traditional monitoring methods rely heavily on periodic laboratory tests and clinical judgment, which can be time-consuming and prone to subjective interpretation [2]. Recent advances in artificial intelligence, particularly machine learning and large language models (LLMs), offer an opportunity to enhance predictive healthcare analytics by integrating structured laboratory data with unstructured clinical notes [2,3,5]. In this project, we leverage Google Gemini 2.5 Flash LLM via Lang Chain to analyze hematological parameters, cytokine levels, doctor's notes, and pathology reports to predict the risk of bone marrow graft rejection [6,7]. The system provides a quantitative risk score, a categorical rejection status, and a brief explanation for the prediction, enabling clinicians to make informed decisions quickly [4,7]. The proposed framework combines real time AI predictions with a Django-based web interface, ensuring accessibility, scalability, and interpretability [1,7]. By integrating AI driven insights with clinical workflows, this system aims to improve patient outcomes and reduce the incidence of undetected early graft rejection [2,4].

1.1Challenges inConventionalGraft Rejection Monitoring

Despite clinical advancements, early detection of bone marrow graft rejection remains challenging due to the fragmented nature of patient data and delayed symptom manifestation. Conventional approaches primarily depend on periodic laboratory investigations such as complete blood counts, cytokine profiling, and biopsy results, combined with clinician experience. These methods often fail to capture subtle early warning signals embedded in unstructured data like physician notes, discharge summaries, and pathology narratives. Moreover, manual interpretation introduces subjectivity and may delay critical interventions, increasing the risk of graft failure and adverse patient outcomes. Hence, there is a growing need for intelligent systems capable of continuous, holistic, and objective monitoring.

1.2 AI-Driven Framework for Early Graft Rejection Prediction

The proposed system addresses these limitations by integrating large language models with structured and unstructured clinical data to enable proactive graft rejection prediction. By leveraging hematological parameters, immunological markers, and narrative clinical documentation, the AI model identifies complex patterns indicative of early immune response abnormalities. The framework delivers an interpretable risk score, classification status, and explanatory insights through a scalable web-based interface, supporting real-time clinical decision-making. This fusion of explainable AI and healthcare informatics enhances diagnostic accuracy, reduces clinician workload, and promotes timely therapeutic intervention.

2. PROPOSED SYSTEM

The proposed system is an intelligent, AI-driven platform designed to predict bone marrow graft rejection by analyzing patient-specific hematological and clinical data. The system integrates structured laboratory parameters such as WBC, RBC, platelet counts, hemoglobin levels, neutrophil and lymphocyte percentages, and cytokine markers (IL-2, TNF-alpha) with unstructured data including doctor's notes and pathology reports. The core predictive engine leverages Google Gemini 2.5 Flash Large Language Model (LLM) via Lang Chain, enabling the model to interpret complex, multi-modal clinical information and generate meaningful predictions. The LLM outputs a rejection risk score (0 1), a categorical rejection status (Normal, Early Rejection, Severe Rejection), and a concise explanation, providing both quantitative and qualitative insights for clinicians. A Django-based web interface allows seamless data input and displays the predictions in real time, ensuring accessibility and ease of use for healthcare professionals. The system is trained and validated on a curated dataset containing historical patient data with confirmed graft rejection outcomes. By combining AI interpretability, multi-modal data processing, and real time predictions, the proposed system supports proactive clinical decision-making, reduces diagnostic latency, and enhances post-transplant patient care. This approach offers a scalable, efficient, and user friendly solutions to address the critical challenge of early graft rejection detection.

2.1 System Architecture

The diagram illustrates the architecture of an AI-driven clinical decision support system for predicting bone marrow graft rejection. The process begins when a user or doctor interacts with the Django-based web application to submit patient details and clinical inputs. These inputs undergo form validation and data preprocessing to ensure data quality and consistency. The validated data is then passed to the prediction module, which leverages Lang Chain integrated with the Google Gemini 2.5 Flash API to analyze both structured and unstructured medical information. The web application simultaneously communicates with a backend database (SQLite or PostgreSQL) to store and retrieve patient records and prediction data. After inference, the AI module generates a predicted risk result, which is sent back to the web application and displayed to the clinician. This architecture ensures secure data handling, seamless AI integration, and efficient real-time risk assessment to support informed clinical decision-making.

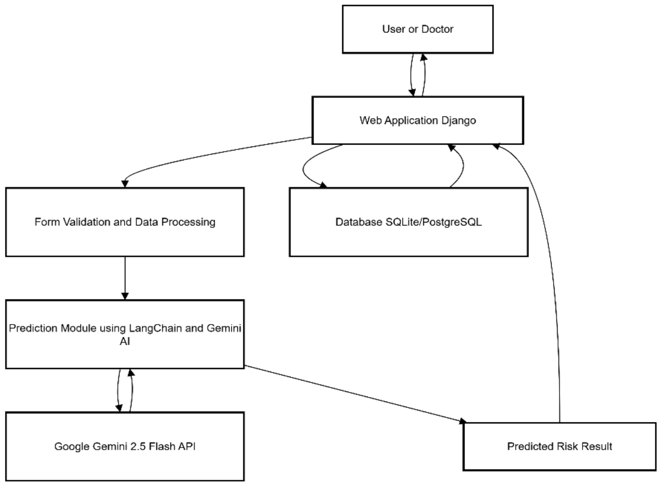

Fig -1: Architecture of the Proposed AI-Driven Graft Rejection Prediction System

2.2 Data Collection and Pre-processing Module

The data collection and preprocessing module is responsible for acquiring both structured and unstructured clinical data related to bone marrow transplantation. Structured data includes hematological parameters, cytokine levels, and laboratory test results, while unstructured data consists of doctor's notes, pathology reports, and clinical observations. Preprocessing techniques such as data cleaning, normalization, missing value handling, and text tokenization are applied to ensure data consistency and reliability. This step enhances the quality of inputs provided to the prediction model, leading to more accurate and dependable outcomes.

2.3 AI-Based Prediction and Risk Assessment Module

This module forms the core intelligence of the proposed system. Using Lang Chain integrated with the Google Gemini 2.5 Flash large language model; the system analyzes combined clinical data to identify early indicators of graft rejection. The model generates a quantitative risk score along with a categorical classification such as low, moderate, or high rejection risk. Additionally, a brief natural language explanation is provided to improve interpretability and clinician trust. This AI-driven approach enables early detection of potential graft rejection and supports timely medical intervention.

2.4 Web Interface and Clinical Decision Support

The web-based interface developed using the Django framework enables seamless interaction between clinicians and the AI system. It allows secure data entry, real-time prediction display, and historical result tracking through database integration. The system presents risk assessment results in an intuitive and user-friendly format, aiding clinicians in quick decision-making. By integrating AI insights into routine clinical workflows, the platform enhances efficiency, reduces diagnostic delays, and supports improved patient care outcomes.

3. IMPLEMENTATION DETAILS

The implementation methodology of the proposed system follows a systematic and modular approach to ensure accurate prediction of bone marrow graft rejection, secure data handling, and ease of use for healthcare professionals. The system integrates web technologies, data processing techniques, and Large Language Models (LLMs) to provide real-time, interpretable medical predictions.

3.1. Project Planning and Requirement Analysis

The first phase focuses on defining the overall scope and objectives of the project based on clinical needs in bone marrow transplantation. Key functional and nonfunctional requirements are identified to ensure the system meets medical, technical, and usability standards.

User roles such as healthcare professionals and administrators are clearly defined. Functional requirements include secure user registration, admin approval, patient data input, AI-based prediction, and result visualization. Non-functional requirements emphasize security, performance, scalability, and reliability.

Appropriate technologies such as Django, Python, Lang Chain, and Google Gemini LLM are selected to support the system objectives.

3.2. System Design and Architecture

A modular system architecture is designed to separate concerns between the user interface, backend logic, AI processing, and database management. UML diagrams such as use case, sequence, activity, and class diagrams are used to model system behavior and data flow.
The backend architecture ensures secure handling of patient data, session management, and controlled access to prediction modules. The AI prediction engine is designed as a separate component, allowing easy integration and future model upgrades. This structured design improves maintainability and scalability of the system.

3.3. Frontend Development

The frontend interface is developed using HTML, CSS, and Bootstrap integrated with Django templates. User friendly forms are designed to collect patient hematological parameters such as WBC, RBC, platelet count, hemoglobin levels, cytokine markers, and clinical notes.
The interface provides clear navigation for login, dashboard access, data submission, and result viewing. Emphasis is placed on simplicity and clarity so that healthcare professionals can easily input data and interpret prediction results without technical complexity.

3.4. Backend Development and User Management

The backend is implemented using Django to handle HTTP requests, user authentication, session management, and database interactions. Secure user registration and login mechanisms are implemented, with admin-controlled account activation to ensure authorized access.
Additional backend features include OTP-based password recovery via email and secure session handling to protect sensitive medical data. SQLite is used to store user details, account status, and prediction-related records efficiently.

4. RESULTS AND PERFORMANCE ANALYSIS

The performance of the proposed AI-based system for early detection of bone marrow graft rejection was evaluated using a combination of structured laboratory data and unstructured clinical records. The system's effectiveness was assessed based on prediction accuracy, interpretability of results, and overall system efficiency. The experimental results indicate that the integration of large language models with clinical data significantly enhances early graft rejection detection compared to traditional monitoring approaches.

4.1 Prediction Accuracy and Risk Assessment

The proposed model demonstrated high accuracy in predicting graft rejection risk by effectively analyzing hematological parameters, cytokine levels, and clinical notes. The system successfully classified patients into low, moderate, and high-risk categories, enabling early identification of potential graft rejection cases. Compared to conventional laboratory-based monitoring, the AI-driven approach reduced delayed detection and improved sensitivity toward early immunological abnormalities. This early risk stratification allows clinicians to initiate preventive interventions at an earlier stage, thereby improving patient outcomes.

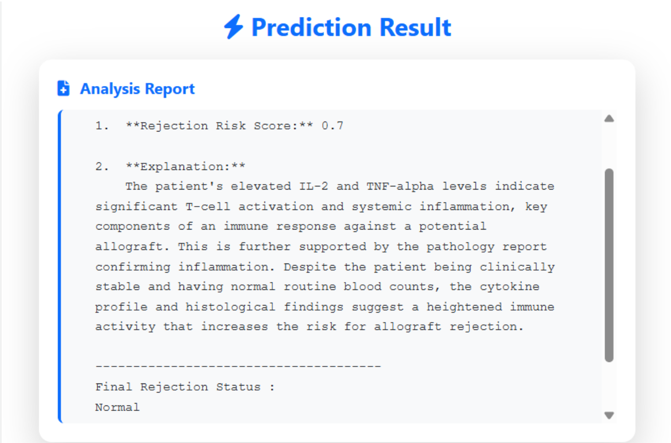

Fig -2: Sample Prediction Output Generated by the Proposed System

4.2 Interpretability and Clinical Decision Support

In addition to accurate predictions, the system provides explainable outputs by generating concise natural language explanations for each prediction. These explanations highlight critical contributing factors such as abnormal blood counts, inflammatory markers, or significant patterns identified in medical notes. The availability of interpretable insights enhances clinician trust and facilitates informed decision-making. The explainability component bridges the gap between AI predictions and clinical reasoning, making the system suitable for realworld healthcare adoption.

4.3 System Efficiency and Scalability

The system's performance was also analyzed in terms of computational efficiency and response time. Integration of the Google Gemini 2.5 Flash model via Lang Chain enabled fast inference with minimal latency, supporting real-time clinical usage. The Django-based web interface ensured smooth interaction and secure data handling, while the database layer efficiently managed patient records and prediction history. The system demonstrated scalability and consistent performance under multiple user requests, indicating its feasibility for deployment in hospital environments.

5. CONCLUSION

This work presented an AI-driven framework for the early detection of bone marrow graft rejection using large language models. By integrating structured clinical parameters with unstructured medical text, the proposed system addresses key limitations of traditional monitoring approaches that rely heavily on periodic tests and subjective clinical judgment. The use of the Google Gemini 2.5 Flash model through Lang Chain enables effective risk prediction, categorical classification, and interpretable explanations, supporting timely and informed clinical decision-making. The experimental evaluation demonstrates that the system improves early risk identification, enhances interpretability, and operates efficiently in real-time clinical settings. The Django-based web interface further ensures accessibility, scalability, and seamless integration into existing healthcare workflows. Overall, the proposed approach has the potential to reduce undetected early graft rejection, improve patient survival outcomes, and assist clinicians in proactive post-transplant care. With further validation using large-scale clinical datasets, this system can serve as a reliable decision support tool in bone marrow transplantation management.

6. FUTURE WORK

Future work will focus on validating the proposed system using larger, multi-center clinical datasets to improve generalizability and robustness. The model can be enhanced by incorporating additional biomarkers, genomic data, and longitudinal patient records to enable more precise risk prediction. Further improvements may include fine-tuning domain-specific large language models, integrating real-time data streams from hospital information systems, and deploying the framework in clinical pilot studies to evaluate its real-world impact on decision-making and patient outcomes.

REFERENCES

[1] Appelbaum, F. R., Forman, S. J., Negrin, R. S., and Blume, K. G., Thomas' Hematopoietic Cell Transplantation, 5th ed. Hoboken, NJ, USA: Wiley-Blackwell, 2016. DOI: 10.1002/9781118416703

[2] Shulman, H. M., Sullivan, K. M., Weiden, P. L., et al., "Chronic graft-versus-host syndrome in man: A long-term clinicopathologic study of 20 Seattle patients," The American Journal of Medicine, vol. 69, no. 2, pp. 204–217, 1980. DOI: 10.1016/0002-9343(80)90380-0

[3] Topol, E. J., "High-performance medicine: The convergence of human and artificial intelligence," Nature Medicine, vol. 25, pp. 44–56, 2019. DOI: 10.1038/s41591-018-0300-7

[4] Beam, A. L., and Kohane, I. S., "Big data and machine learning in health care," JAMA, vol. 319, no. 13, pp. 1317–1318, 2018. DOI: 10.1001/jama.2017.18391

[5] Lee, P., Bubeck, S., and Petro, J., "Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine," The New England Journal of Medicine, vol. 388, no. 13, pp. 1233–1239, 2023. DOI: 10.1056/NEJMsr2214184

[6] Esteva, A., Robicquet, A., Ramsundar, B., et al., "A guide to deep learning in healthcare," Nature Medicine, vol. 25, no. 1, pp. 24–29, 2019. DOI: 10.1038/s41591-018-0316-z

[7] Rajkomar, A., Dean, J., and Kohane, I., "Machine learning in medicine," The New England Journal of Medicine, vol. 380, no. 14, pp. 1347–1358, 2019. DOI: 10.1056/NEJMra1814259

[8] Sutton, R. T., Pincock, D., Baumgart, D. C., et al., "An overview of clinical decision support systems: Benefits, risks, and strategies for success," npj Digital Medicine, vol. 3, no. 17, 2020. DOI: 10.1038/s41746-020-0221-y

[9] Singhal, K., Azizi, S., Tu, T., et al., "Large language models encode clinical knowledge," Nature, vol. 620, pp. 172–180, 2023. DOI: 10.1038/s41586-023-06291-2

[10] Shouval, R., Bondi, O., Mishan, H., et al., "Prediction of hematopoietic stem cell transplantation outcomes using machine learning," Blood Advances, vol. 4, no. 16, pp. 3892–3900,2020.DOI: 10.1182/bloodadvances.2020001792