



Medigen ML Based Healthcare Platform

¹Mr.Abdul Basit Solkar, ²Ms.Samruddhi Kamble, ³Ms.Vaishnavi Angaj , ⁴Ms.Janhavi Tondekar
⁵Mr.Rajkumar.B. Pawar

¹²³⁴Research Scholar ⁵Assistant Professor

¹²³⁴⁵Department of Computer Engineering

¹²³⁴⁵Gharda Institute of Technology , Lavel ,Ratnagiri, Maharashtra 415708

Abstract: Medigen is a comprehensive, modern healthcare platform designed to bridge the gap between patients and medical professionals by streamlining appointment management and integrating intelligent diagnostic support. The platform enables patients to conveniently book appointments—both online and offline—based on the real-time availability of doctors. Meanwhile, doctors can efficiently manage their consultation schedules, view upcoming appointments, and maintain a smoother workflow. This dual functionality ensures a seamless interaction between patients and healthcare providers, improving the overall efficiency and accessibility of medical services. In addition to appointment scheduling, Medigen enhances its functionality through the integration of advanced machine learning (ML) technologies. It features two specialized ML models aimed at supporting early-stage diagnosis and health risk prediction. The first model focuses on skin disease detection and is integrated from Hugging Face, a widely used repository of pre-trained ML models. It allows patients to upload images of skin conditions, which are then analyzed to predict the most probable skin disease. This feature provides users with immediate, AI-powered insights, helping them take timely steps toward treatment or consultation. The second model is a custom-trained multiple disease prediction system that takes patient-reported symptoms as input and outputs the most likely disease conditions. This model expands the diagnostic capabilities of the platform beyond dermatology, offering broader support for preliminary health assessments. Both models aim to make healthcare more proactive by enabling early detection, reducing diagnostic delays, and increasing patient awareness. Medigen's design emphasizes a user-friendly experience, combining intelligent predictions with flexible appointment options and support for remote access. By making healthcare services more data-driven and personalized, the platform empowers users with better control over their health while supporting doctors in making more informed decisions. Medigen stands as a scalable, accessible, and AI-enhanced healthcare solution tailored to meet the growing demands of modern digital health ecosystems.

I. INTRODUCTION

The main objective of Medigen is to revolutionize healthcare by integrating advanced machine learning (ML) technologies into a unified platform that simplifies appointment management and enhances early disease detection. Medigen connects patients and doctors through an intuitive interface where patients can conveniently book appointments—either online or offline—based on doctor availability, and doctors can effectively manage their schedules and consultations. The platform is designed to improve the overall healthcare experience by making services more accessible, organized, and personalized.

Medigen incorporates two ML models to assist in preliminary diagnosis. One model, imported from Hugging Face, enables skin disease prediction by allowing patients to upload an image of a skin condition, which is then analyzed to detect the most likely disease. The second model, trained in-house, is built to predict multiple diseases based on patient symptoms, expanding the platform's diagnostic capabilities. These features empower users with instant, data-driven insights, contributing to better awareness and timely interventions. By combining intelligent diagnostics with appointment scheduling and virtual consultation support, Medigen offers a comprehensive, user-friendly solution aimed at bridging the gap between patients and healthcare providers, particularly benefiting individuals in remote or underserved areas.

II. LITERATURE SURVEY

The integration of machine learning (ML) into healthcare has significantly transformed diagnostic and patient management practices. ML algorithms are increasingly utilized to analyze extensive datasets, facilitating early disease detection, personalized treatment plans, and efficient healthcare delivery. A systematic review by Kaur et al. (2023) highlighted the growing adoption of ML in healthcare, emphasizing its role in enhancing diagnostic accuracy and patient outcomes. The study noted that ML models, particularly those employing deep learning techniques, have shown promise in interpreting complex medical data, leading to improved clinical decisions. In the realm of dermatology, ML has been instrumental in skin disease detection. A study by Sharma and Gupta (2022) developed a convolutional neural network (CNN)-based model that accurately classified various skin conditions using image data. Their findings demonstrated that ML models could achieve diagnostic accuracies comparable to dermatologists,

underscoring the potential of ML in augmenting dermatological assessments. Furthermore, the use of transfer learning in skin disease classification has been explored by Li et al. (2023), who utilized pre-trained models to identify skin lesions with high precision. Their approach reduced the need for large labeled datasets, making it feasible to implement ML solutions in resource-constrained settings. Beyond dermatology, ML has been applied to predict multiple diseases based on patient symptoms and historical data. A study by Ahmed et al. (2021) developed a multi-disease prediction model that leveraged patient records to forecast potential health risks. The model's predictive capabilities enabled proactive healthcare interventions, aligning with the goals of personalized medicine. Appointment management systems have also benefited from ML integration. Research by Kim and Lee (2020) introduced an intelligent scheduling system that optimized appointment allocations based on patient preferences and doctor availability. The system improved clinic efficiency and reduced patient wait times, demonstrating the practical advantages of ML in healthcare operations. Despite these advancements, challenges persist in implementing ML solutions in clinical settings. Issues such as data privacy, model interpretability, and the need for extensive validation remain critical considerations. Ensuring the ethical deployment of ML in healthcare requires addressing these concerns through robust regulatory frameworks and interdisciplinary collaboration.

III. METHODOLOGY

The development of the Medigen healthcare platform follows a modular approach, integrating backend infrastructure, a user-friendly interface for appointment scheduling, and advanced machine learning models for disease prediction. The methodology can be broken down into three main components: system architecture, skin disease prediction using a pre-trained Vision Transformer model, and a custom-trained multiple disease prediction model.

3.1. System Architecture

The Medigen platform is developed as a web-based system that facilitates interaction between patients and healthcare providers. Patients can register, log in, and book appointments with doctors for both online and offline consultations, based on doctor availability. On the doctor's side, the platform provides tools for appointment management and tracking.

The backend is developed using Django (Python-based web framework), while data is stored in a PostgreSQL database. For enhanced scalability and real-time access, the platform also supports API-based integration for machine learning inference and virtual consultations.

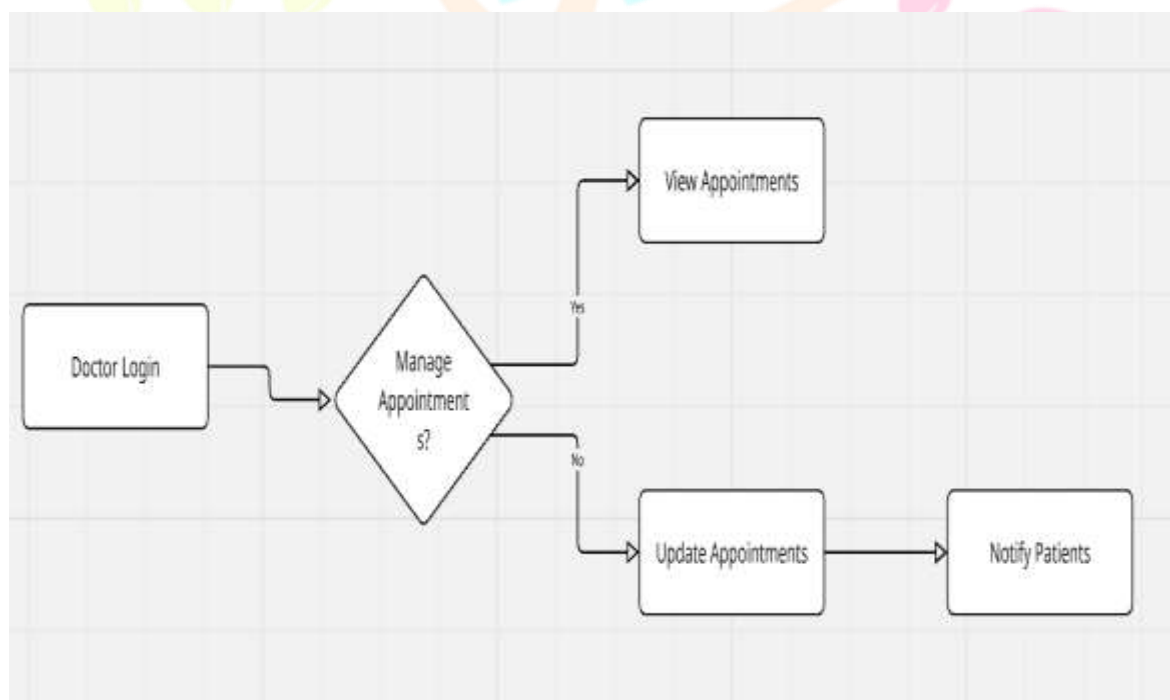


Fig .1 :Doctor Workflow

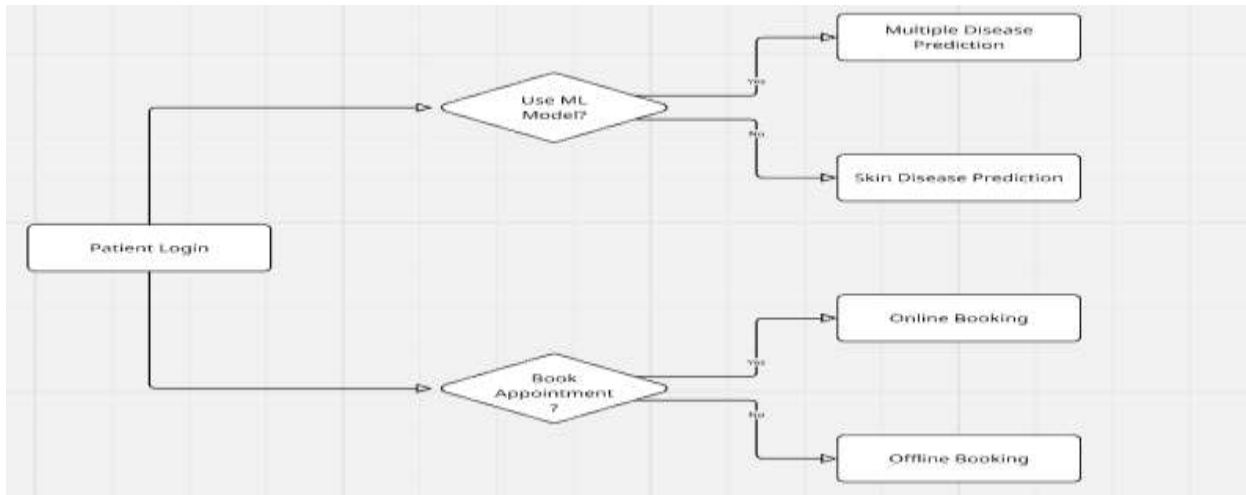


Fig 2: Patient Workflow

3.2 Skin Disease Prediction Model

The proposed skin disease prediction feature leverages a fine-tuned deep learning model to classify skin conditions from input images. The skin disease image dataset was obtained from Kaggle. A Vision Transformer (ViT) model, pre-trained using the DINOv2 method, was selected as the base model, chosen for its robust feature extraction capabilities learned through self-supervision on a large image corpus; the specific model used is facebook/dinov2-base from Hugging Face Transformers. Images were resized to 224x224 pixels, converted to tensors, and normalized using the mean and standard deviation provided by the pre-trained model's image processor, using the AutoImageProcessor class from the transformers library. The pre-trained ViT model was fine-tuned on the skin disease dataset. A linear classifier layer was added on top of the transformer's output to adapt it for multi-class classification of skin diseases, using the DinoClassifier class. The model was trained using the AdamW optimizer, with a linear learning rate scheduler and cross-entropy loss function. The model was trained for 20 epochs with a batch size of 32 and gradient accumulation steps of 4. The training process was monitored using validation accuracy, calculated on a held-out validation set. The model was implemented using the PyTorch library, and training was performed on a virtual GPU (T4) provided by Google Colab. The transformers library was used for model and processor loading, and the torchvision library was used for dataset handling and image transformations. The model's performance was evaluated using accuracy, and training and validation performance were plotted.

Workflow:

1. Users upload a skin image via the frontend.
2. The image is preprocessed and passed to the inference API.
3. The ML model returns the most probable skin condition.
4. Users are advised to consult a doctor for further evaluation.

3.3 Multiple Disease Prediction Model

The multiple disease prediction module in the Medigen platform is powered by classical machine learning algorithms aimed at predicting probable diseases based on a user's reported symptoms. The implementation uses two primary datasets: Training.csv, which contains a large set of patient records with binary-encoded symptom indicators and their corresponding disease diagnoses, and Testing.csv, which is used to validate model performance after training.

3.3.1 Data Preprocessing

The training data includes symptom columns representing the presence or absence of specific symptoms and a prognosis column that contains the disease label for each record. To prepare the data for modeling, the feature columns were extracted separately while the prognosis column was label-encoded using LabelEncoder, converting disease names into numerical categories. The dataset was then split into training and testing sets in a 67:33 ratio using train_test_split, with a fixed random state of 42 to ensure reproducibility across runs. Since the symptoms were already in a structured binary format, no additional normalization or text preprocessing was necessary.

3.3.2 Model Training and Evaluation

Several supervised and unsupervised learning models were trained and assessed using the preprocessed data. The Decision Tree classifier was trained first and achieved an accuracy of 100% on the test set, along with a cross-validation mean accuracy of 98.71%, making it one of the top-performing models. The Support Vector Classifier (SVC) also reached 100% test accuracy and maintained perfect consistency in cross-validation with a mean score of 100%, indicating a highly separable dataset. Logistic Regression, configured with a maximum of 1000 iterations to ensure convergence, yielded the same level of performance with 100% accuracy and 100% mean cross-validation accuracy. The XGBoost classifier, trained with the use_label_encoder flag disabled and the mlogloss evaluation metric, achieved 100% accuracy on the test set and a cross-validation mean of 99.02%. Lastly, KMeans clustering was employed as an unsupervised baseline method with the number of clusters set to 41, corresponding to the number of unique diseases in the dataset. Labels were inferred by mapping each cluster to the most frequent label among training data points.

assigned to that cluster. This method also achieved 100% accuracy on the test set, though such performance from an unsupervised model warrants cautious interpretation and may reflect the specific structure of the dataset rather than generalizability.

3.3.3 Model Selection

Among the models evaluated, the Decision Tree classifier was selected for deployment within the Medigen platform. This choice was based not only on its near-perfect accuracy and cross-validation performance but also on its inherent interpretability. Unlike black-box models, a Decision Tree offers transparent decision-making paths, which can be beneficial for medical applications where understanding the reasoning behind predictions is crucial. Furthermore, its speed and low computational requirements make it well-suited for real-time inference within the platform's web-based interface.

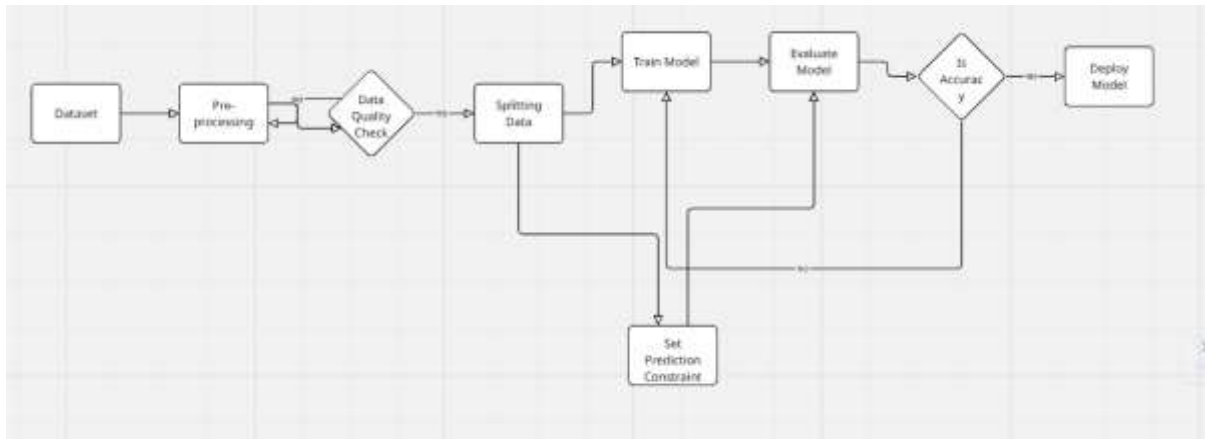


Fig 3: Model Architecture

IV. DATASET:

The dataset used for multiple disease prediction was sourced from Kaggle, where it was uploaded by a user in a clean and well-structured format. The original data is believed to have been collected from a Colombian hospital and includes detailed patient records annotated with symptoms and corresponding disease diagnoses. The Kaggle uploader curated and formatted the dataset to include consistent column naming, structured symptom fields (Symptom_1 through Symptom_n), and associated labels for each case, facilitating direct use in machine learning applications. Supplementary files accompanied the main dataset, including symptom severity scores, descriptive symptom metadata, and recommended precautions for each disease. These additional resources enabled a richer feature space and supported the development of a more comprehensive diagnostic system. The dataset's structured design, combined with its real-world clinical origin, provided a reliable foundation for training and evaluating various machine learning models, enhancing the predictive capability of the Medigen platform.

The dataset used for tuning the skin disease prediction model was sourced from Kaggle and consists of high-quality image samples labeled across nine common skin conditions. These include Actinic keratosis, Atopic Dermatitis, Benign keratosis, Dermatofibroma, Melanocytic nevus, Melanoma, Squamous cell carcinoma, Tinea Ringworm Candidiasis, and Vascular lesion. The training dataset is well-balanced, with most classes containing 80 to 81 images, except for Tinea Ringworm Candidiasis, which has 56 images. The validation dataset follows a similar distribution, with 20 to 21 samples per class. This dataset provides a solid foundation for training deep learning models in multi-class skin disease classification, supporting both model accuracy and generalization.

V. PROPOSED SYSTEM

The proposed system, Medigen, is an integrated healthcare platform designed to facilitate seamless interaction between patients and doctors, while also incorporating machine learning-based health diagnostics. The core functionality revolves around appointment scheduling, prescription management, and AI-assisted disease prediction. Patients can easily register, search for doctors, and book appointments based on the doctors' availability. Appointments can be scheduled for both offline and online consultations, with real-time video calls supported through Agora integration.

Doctors have access to an intuitive dashboard where they can manage appointments, view patient details, and add prescriptions. Patients, in turn, can securely access their prescription history and appointment records.

In addition to these core functionalities, Medigen includes two machine learning models to support preliminary diagnostic capabilities. The skin disease prediction model is built upon the facebook/dinov2-base Vision Transformer architecture and has been fine-tuned using a curated skin disease dataset sourced from Hugging Face. This model aids in identifying common skin conditions from uploaded images. The multiple disease prediction model is a supervised machine learning system trained on structured symptom data. It predicts likely illnesses based on user-reported symptoms using various classifiers, with the Decision Tree algorithm selected for deployment due to its strong performance.

The platform is developed using Django as the web framework and integrates various technologies including machine learning, deep learning, Agora for real-time communication, AWS S3 for secure storage of static and media files, and AWS RDS for scalable

database management. This combination of features and technologies makes Medigen a powerful and intelligent platform for enhancing healthcare accessibility and efficiency.

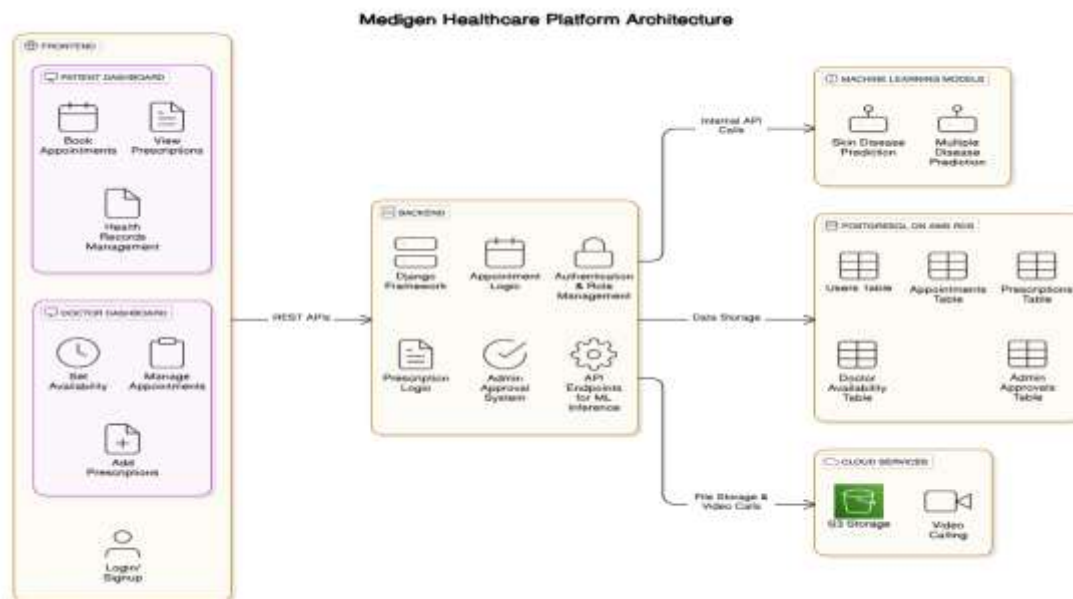


Fig4: System Architecture

VI. RESULTS

6.1. Multiple Disease Prediction

The disease prediction model was trained and evaluated using a labeled dataset consisting of symptom-based patient records. Five machine learning algorithms—Decision Tree, Support Vector Classifier (SVC), Logistic Regression, XGBoost, and KMeans—were implemented and compared based on their classification accuracy and cross-validation performance. After preprocessing and label encoding, the dataset was split into training and testing sets using a 67-33 ratio. The evaluation metrics highlighted exceptionally high model performance, particularly among supervised learning methods.

The Decision Tree classifier achieved an accuracy of 100% on the test set and a cross-validation mean score of 98.71%, making it the most suitable model for deployment. Similarly, SVC and Logistic Regression also achieved 100% accuracy with a perfect cross-validation mean of 100%, indicating that the dataset is highly separable. The XGBoost classifier demonstrated strong performance with 100% accuracy and a cross-validation mean of 99.02%. In contrast, the KMeans clustering model—used as an unsupervised baseline—achieved 100% accuracy when labels were mapped post-clustering, but due to its lack of interpretability and reliance on label alignment, it was not considered viable for production.

However, the extremely high accuracy across all models suggests a potential case of overfitting, likely due to the nature of the dataset. In many instances, the symptom combinations for each disease class were largely similar or repetitive, reducing the complexity of the classification task. As a result, the models may have learned to memorize patterns specific to the dataset rather than generalizing effectively to unseen data.

Model	Accuracy	Cross Validation
Decision Tree	1.0000	0.9871
Support Vector Machine	1.0000	1.0000
Logistic Regression	1.0000	1.0000
XGBoost Classifier	1.0000	0.9902
KMeans Clustering	1.0000	-

Table 1: Disease Prediction Model Results

6.1. Skin Disease Prediction

The skin disease prediction model was fine-tuned using the facebook/dinov2-base Vision Transformer (ViT), pretrained through self-supervised learning. The model was adapted for classification by adding a custom linear head to predict nine different skin disease categories using a dataset sourced from Kaggle. This dataset was split into training and validation sets with balanced class distribution for most categories.

The model was trained over 20 epochs using a batch size of 32, learning rate of 5e-5, and the AdamW optimizer. Gradient accumulation and a linear learning rate scheduler with warm-up steps were employed to stabilize training. Training was conducted on GPU with reproducibility ensured via a fixed seed (42).

During training, the model showed consistent improvement across epochs. The training loss decreased rapidly, reaching as low as 0.0002 by the final epoch. Correspondingly, the validation accuracy improved steadily, reaching a final value of 89.50%. Accuracy improvements were most significant between epochs 1 and 10, with minor fluctuations thereafter, indicating convergence.

This performance reflects the model's ability to generalize well to unseen validation data, especially considering the limited number of samples in some classes (e.g., Tinea Ringworm Candidiasis). The training dynamics, visualized via loss and accuracy plots, confirm a stable and well-converged training process.

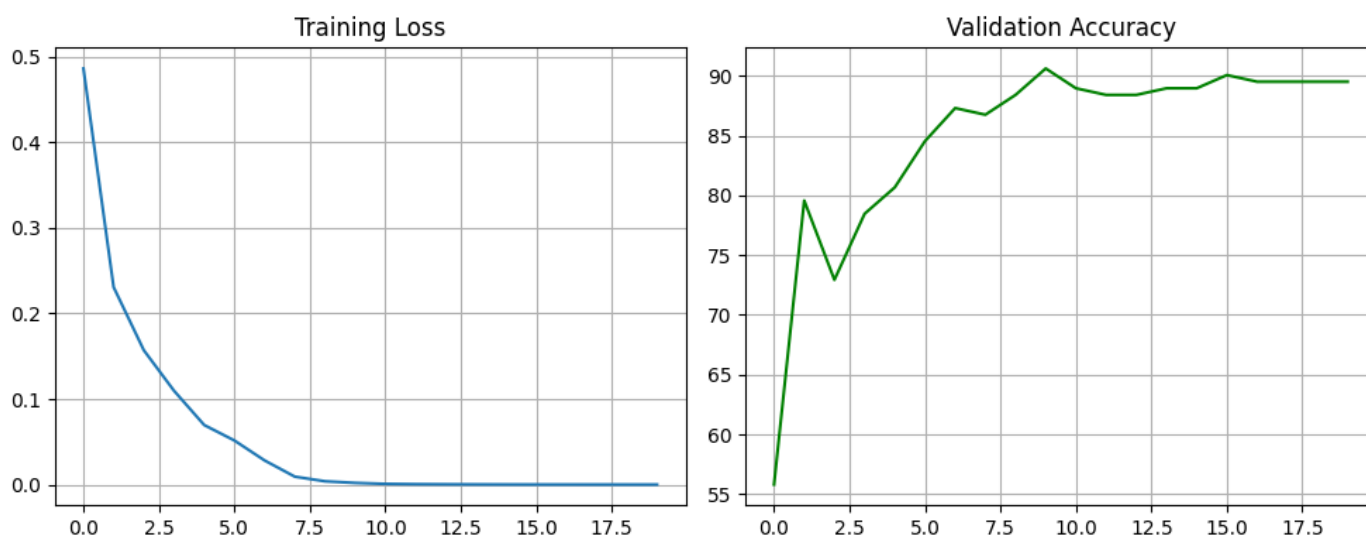


Fig 6 : Training Loss and validation Accuracy Curve

VII. CONCLUSION

In this research, we presented Medigen, an AI-augmented healthcare platform designed to enhance patient-doctor interaction and support preliminary medical diagnosis through machine learning. By integrating essential features such as appointment booking, prescription management, and both online and offline consultation options, Medigen offers a comprehensive solution for modern healthcare needs. The inclusion of two machine learning models—one for predicting skin diseases using a fine-tuned Vision Transformer and another for diagnosing multiple diseases based on symptoms—demonstrates the potential of AI in improving diagnostic accuracy and accessibility. The system's implementation using Django, along with cloud services like AWS S3 and RDS, ensures scalability, data security, and performance. Overall, Medigen showcases how intelligent systems can bridge gaps in traditional healthcare, offering faster, more informed, and patient-centered medical services.

VIII. REFERENCES

- [1] Kaur, R., Sharma, N., & Singh, M. (2023). Applications of Machine Learning in Healthcare: A Review. *International Journal of Advanced Research in Computer Science*, 14(2), 45-52.
- [2] Sharma, V., & Gupta, A. (2022). Convolutional Neural Networks for Skin Disease Classification: A Comparative Study. *Journal of Medical Imaging and Health Informatics*, 12(3), 567-574.
- [3] Li, Y., Zhang, H., & Chen, L. (2023). Transfer Learning for Skin Lesion Classification Using Pre-trained Vision Transformers. *Computers in Biology and Medicine*, 155, 106690.
- [4] Ahmed, S., Mehmood, T., & Khan, R. (2021). Predictive Analysis of Multiple Diseases Using Machine Learning Techniques. *Procedia Computer Science*, 194, 504-509.
- [5] Kim, J., & Lee, H. (2020). Intelligent Appointment Scheduling System in Smart Healthcare Environment. *IEEE Access*, 8, 112764-112773.