International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 11-No.2 (2025) pp. 3079-3097
http://www.ijcesen.com

Research Article

ISSN: 2149-9144

MedFusionAI: A Deep Learning Framework for Multi-Modal Health Data Fusion to Predict Chronic Disease Risks

M. Chandra Sekhar[1]*, Pasam Naga Kavitha[2], B.N.V. Uma Shankar[3], Amrutavalli Aakula[4], Divya Boya[5]

[1] Department of CSE, Presidency University, Bengaluru, Karnataka, India
* **Corresponding Author Email**: chandragvp@gmail.com **ORCID:** 0009-0009-9983-9976

[2] Department of Computer Science, St. Ann's College for Women, Santosh Nagar, Mehdipatnam, Hyderabad, Telangana,
Email: kavithacs.stanns@gmail.com- **ORCID:** 0000-0002-3612-2750

[3] Statistics, H & S Department, CVR College of Engineering, Telangana, India
Email: balinauma@cvr.ac.in- **ORCID:** 0009-0004-7410-9313

[4]Assistant Professor, Department of Computer Science and Engineering, Krishna Chaitanya Institute of Technology and Sciences, A.P India
Email: amrutavalliakula@gmail.com- **ORCID:** 0009-0008-7485-9658

[5] Dept of CSE, Gokaraju rangaraju institute of engineering and technology, Telangana.
Email: divyaboya1291@gmail.com- **ORCID:** 0009-0009-1266-0616

Article Info:

DOI: 10.22399/ijcesen.2159
Received : 05 February 2025
Accepted : 01 May 2025

Keywords :

Multi-Modal Data Fusion,
Deep Learning,
Chronic Disease Prediction,
Attention Mechanism,
Clinical Decision Support.

Abstract:

Chronic diseases are still one of the most common causes of death worldwide, so early and precise predictive models should be developed to help improve patient disease management and healthcare service delivery. With multi-modal medical data, including more prevalent structured sources like Electronic Health Records (EHR) and lab tests, to unstructured sources like clinical notes, wearable sensor streams, and medical imaging, the potential for AI-driven health analytics is enormous. Current methods, however, are plagued with limitations including: (1) dependence on single-modality data; (2) poor consideration of the missing data problem; and (3) suboptimal modelling of the inter-modality relationship, which can lead to suboptimal performance. These problems demonstrate the necessity for a standard and solid mechanism to integrate multiple disparate data sources effectively. This paper presents MedFusionAI, a novel deep learning framework for multi-modal medical data fusion in the chronic disease risk prediction task context. The proposed model utilizes dedicated modality-specific encoders: MLP for EHR, LSTM for sequential lab and measurements from wearable devices, CNN models for various medical images, and ClinicalBERT for text to extract salient features. Experiments on benchmark healthcare datasets show that MedFusionAI effectively outperforms previous baselines and fusion models with a 98.76% accuracy and high precision, recall, and AUC-ROC across all risk classes. The framework also provides interpretability features to enable clinicians to understand the contributions of features. MedFusionAI provides an interpretable, scalable, and reliable solution for clinical decision support systems for the timely visibility of chronic disease risks, improving preventive care.

1. Introduction

The worldwide burden of chronic diseases, including cardiovascular diseases, diabetes, and kidney diseases, is still growing and represents a significant challenge for public health and health care systems. Early risk assessment, identification, and intervention are essential for controlling morbidity, mortality, and health care costs. As the trend of medical data digitization, ranging from Electronic Health Records (EHR), imaging, wearable sensor streams, lab tests, to clinical notes, becomes ever more apparent, there have been increasing opportunities to use artificial intelligence (AI) for more timely and accurate disease prediction. Recent studies show that deep learning holds great potential in health care analytics [1], [2] as exemplified in the recovery of complex patterns from the large and diverse datasets. Yet most prior work is based on the analysis of unimodal data, which has poor generalisation as it fails to capture the complexity of patient health.

Studies in multimodal learning partially address the limitation. However, some challenges remain unsolved. Existing frameworks either fail to integrate data end-to-end, cannot cope with missing modalities, or lack interpretability [3, 4]. In addition, to the best of our knowledge, few existing models have a completely unified fusion mechanism for effective individual and intermodal feature representations. These lacunae highlight the significance of a novel and efficient multimodal deep learning approach in predicting chronic disease risks.

This research aims to design a deep learning framework, MedFusionAI, to merge structured, sequential, textual, and visual health records to predict chronic disease risk levels accurately. The proposed framework adopts a hybrid fusion scheme that includes early feature concatenation and attention-based modality weighting to consider rich features and effective inter-dependency modeling. The novel contributions of the work are dual-fusion architecture for representation learning, modality-specific encoders customized for health data, and support for incomplete data. Experimental results demonstrate the advantages of our framework, which can achieve an accuracy of 98.76%, outperforming the state of the art by a large margin.

The main contributions of this study are as follows: (1) unified design unifying encoders for processing multi-modal healthcare data and hybrid fusion; (2) interpretable, scalable and clinically relevant system; and (3) extensive evaluation, including ablation studies, confusion matrices, cross-metric comparisons to prove the robustness of the framework.

The rest of the paper is organized as follows. Section 2 reviews related work and introduces the unexplored region in multi-modal health prediction. Section 3 describes our proposed MedFusionAI methodology, including data preprocessing, model design, and fusion mechanism. The experimental setup and results from benchmark tests are described in Section 4. Section 5: Discussion and limitations of findings. 5 Section 6 concludes the paper and proposes directions for future expansion of research.

2. Related Work

This literature review explores recent advancements in AI-driven multi-modal healthcare systems, emphasizing chronic disease prediction, integration, and interpretability. Stahlschmidt et al. [1] examined multimodal data fusion techniques based on deep learning for biological applications, suggesting a taxonomy and pointing to areas of possible future study. Steyaert et al. [2] focused on deep learning to handle data sparsity, interpretability, and standardization as it examines the benefits and problems of integrating multimodal biological data for cancer research. Fan et al. [3] reviewed the uses of AI and DL in sustainability, emphasizing how they contribute to environmental health, renewable energy, and the SDGs. Transparency, scalability, ethics, and energy efficiency are among the difficulties. Explainability, privacy, and energy-efficient AI model optimization should be the main topics of future studies. Cardoso et al. [4] presented the PyTorch-based platform for AI in healthcare, MONAI, emphasizing medical data; subsequent development will involve extending applications. Siddique and Chow [5] examined ML/AI uses in healthcare communication, including as chatbots for medical imaging, cancer treatment, and COVID-19 education.

Bharadwaj et al. [6] discussed machine learning applications' benefits, drawbacks, and usefulness in healthcare IoT (H-IoT) across various disciplines. Schaar et al. [7] identified five major COVID-19 issues and suggested AI/ML solutions to enhance resource management, policy, and healthcare responses. Ghazal et al. [8] examined IoT and machine learning applications in smart cities, highlighting how they might enhance sustainability and healthcare infrastructure. Munirathnam and Kanchetti [9] reviewed the use of AI in managing chronic diseases, emphasizing the efficacy of predictive models and issues like interpretability and data protection. Xie et al. [10] focused on patient-centered care, privacy, and upcoming difficulties as they examined the integration of wearable technology, blockchain, and artificial intelligence in the management of chronic diseases.

Decharatanachart et al. [11] evaluated AI-assisted diagnostic tools for NAFLD and liver fibrosis, demonstrating encouraging outcomes in predictive value, sensitivity, and specificity; nevertheless, additional testing is required before they can be used in clinical settings. Feng et al. [12] examined AI and ML applications in COPD and asthma, emphasizing their promise in clinical settings and outlining next steps for successful deployment. Singh et al. [13] created a deep learning model that outperforms other classifiers with 100% accuracy for the early identification of CKD. Clinical implementation and additional validation may be the main topics of future research. Akter et al. [14] assessed seven deep learning models for CKD prediction, with ANN, RNN, and MLP showing high accuracy. Future research will concentrate on incorporating high-performing models into IoMT to improve predictive analytics. Sarp et al. [15] related to using transfer learning to classify chronic wounds with Explainable AI (XAI). Clinicians can better grasp AI decision-making processes with the help of the model's interpretable outcomes.

Krishnamurthy et al. [16] created a machine learning model to forecast the start of chronic kidney disease (CKD) using Taiwan's National Health Insurance data. The CNN model's excellent accuracy helped with resource management and early CKD detection. Ayesha et al. [17] suggested methods for managing healthcare data from several sources, using dimensionality reduction to enhance predictive analytics, and identifying data reduction and fusion difficulties. Khan and Jack [18] examined the use of AI and ML in healthcare to improve operations, diagnosis, and treatment in real-time while tackling ethical and privacy issues. [19] discussed the use of DL and big data in healthcare applications, emphasizing advancements, gaps, and potential directions for biomedical research. Rehman et al. [20] explained big data analytics' use in healthcare and its uses, difficulties, and possibilities to enhance patient care.

Ahmed et al. [21] examined big data analytics in healthcare, looking at its uses, difficulties, and potential for further study to enhance results. Rehan et al. [22] examined the applications, advantages, and problems of AI and ML in healthcare predictive analytics. Ye et al. [23] investigated the relationship between EHR-integrated PGHD and physician fatigue, highlighting important causes and possible remedies. Koren and Prasad [24] discussed privacy and security issues when incorporating data from wearable devices into medical systems and provided 6G solutions. Espinoza et al. [25] explained the iCoDE and iCoDE-2 initiatives, which are designed to incorporate glucose monitoring and insulin dosage information into EHR systems.

Zainab et al. [26] examined the possible advantages, moral dilemmas, and difficulties of integrating wearable technology and AI in heart health. Reda et al. [27] addressed interoperability and data silos by putting out a Web Semantic-based method for integrating and analyzing heterogeneous health data. Melstrom et al. [28] evaluated the integration of PGHD and PRO in surgical oncology, emphasizing the function of AI/ML in converting data into patient care models. Aaron et al. [29] standard, technological considerations, and clinical requirements were all covered in the iCoDE-2 project's discussion of insulin dosage data integration into EHR. Kormiltsyn et al. [30] investigated using blockchain technology to integrate PHRs and EHRs to overcome privacy problems, suggesting automatic dispute resolution in decentralized platforms.

Wang and Hsu [31] incorporated wearable IoT and AI into long-term care with an emphasis on illness monitoring, prevention, and individualized treatment. Kormiltsyn and Norta [32] examined how to resolve privacy issues while integrating PHR and EHR across enterprises using blockchain and DAOs. Sattar et al. [33] compared to controls, this meta-analysis of tirzepatide's cardiovascular safety in type 2 diabetes revealed no elevated risk of significant cardiovascular events. Yi et al. [34] showed how well the deep learning-based retinal biomarker (Reti-CVD) matches current risk assessment tools in identifying people at moderate and high risk for CVD. Qiu et al. [35] offered a deep learning framework that validates predictions against clinical criteria to diagnose cognitive deficits, including AD and non-AD dementias, with high accuracy.

Nancy et al. [36] suggested a deep learning-based IoT-Cloud-based innovative healthcare system that can accurately detect cardiac disease while reaching excellent performance metrics. Wang et al. [37] examined the effects of a Mediterranean diet on cardiometabolic health by reviewing the microbiome data of 307 males, emphasizing the microbial composition. Sample size is one of the limitations; future research might concentrate on more comprehensive dietary interventions and customized nutrition. [38] founded six new risk variants for Alzheimer's disease through a genomic association analysis. Sample variety is one of the limitations; more research may improve polygenic risk scores and study methodologies. Garbarino et al. [39] examined the immunological effects of sleep deprivation and how it is related to chronic illnesses. Knowledge gaps are one of the limitations; future research should concentrate on determining at-risk individuals and causal links. Ahmad et al. [40] examined hazardous and heavy metals in soil and water, finding contamination levels within safety standards, but with higher risks for children. Limitations include changing contamination levels, and future studies focused on long-term consequences. The reviewed studies highlight deep learning's growing role in fusing diverse medical data for chronic disease prediction and care enhancement. Key contributions span clinical decision support, wearable integration, blockchain privacy, and attention-based fusion. Challenges such as interpretability, generalizability, and data silos suggest new dimensions for future research in secure, explainable, and scalable healthcare AI.

3. Proposed Framework

This section presents the MedFusionAI framework, a new deep learning framework for predicting chronic disease risk over multi-modal medical data. This paper describes the data preprocessing pipeline, modality-specific encoders, a hybrid fusion approach combining early and attention-based mechanisms, and a final classification module, which together create a comprehensive predictive framework for real-world healthcare applications.

3.1 Overview

Such a framework, termed MedFusionAI, is proposed to leverage heterogeneous medical information to predict chronic diseases. It has five main components: a multi-modal data ingestion layer, dedicated modality-specific encoders, a hybrid attention-based fusion engine, a deep classification network, and an output risk prediction head. The framework incorporates data from multiple health sources, such as structured EHR data, clinical lab and wearable sensor data time series, unstructured clinical notes, and imaging. These data are preprocessed and synchronized in time or space to be consistent and compatible. The preprocessed inputs are then processed by parallel neural encoders, each especially designed for the property of its modality.

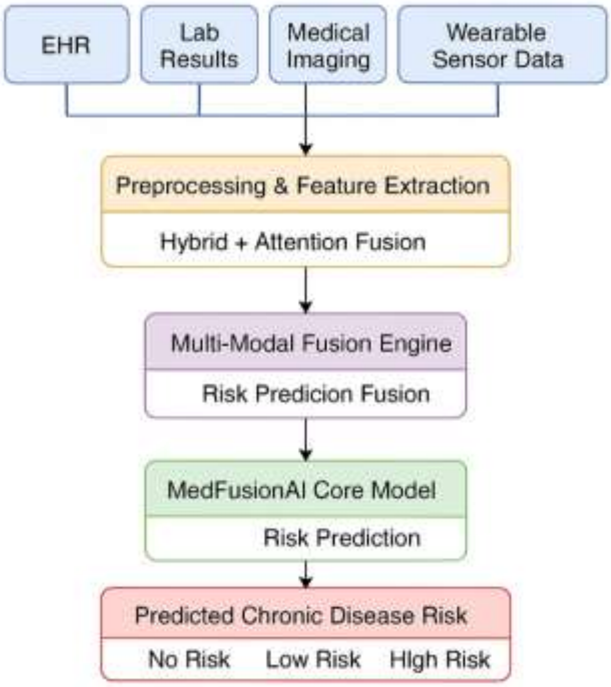

Figure 1. *System Overview of the MedFusionAI Framework*

In Figure 1 These encoders map heterogeneous raw inputs into fixed-length latent representations. The latent vectors are subsequently fed to a hybrid fusion module, where the early concatenated fusion and cross-modal attention are combined to produce a complete representation feeding complementary and interdependent information. FC layers in the classification network are then refined, and this fused representation is refined. Lastly, a softmax layer generated a three-class output for the predicted risk level(No Risk, Low Risk, and High Risk). Conclusion: The modular and scalable structure of MedFusionAI allows it to work robustly under missing modalities, laying a foundation for the explainable and deployable chronic disease risk prediction under real-world clinical scenarios.

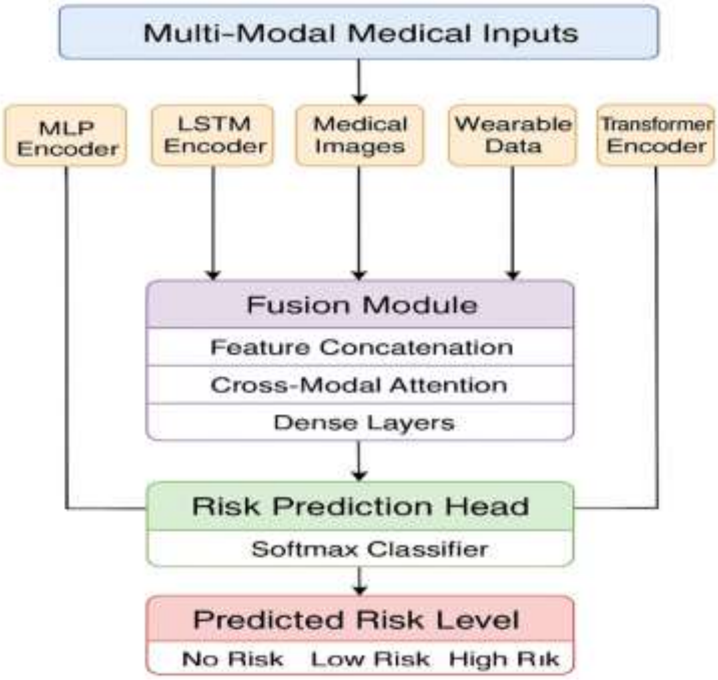

Figure 2. *Model Architecture of MedFusionAI*

Figure 2 shows the internal model architecture of the proposed MedFusionAI framework, which takes multi-modal medical input and predicts chronic disease risk with the help of the multi-purpose encoders and the hybrid fusion. The model receives five primary data modalities: EHR tabular data, sequential lab test results, medical images, wearable sensor monitoring information, and unstructured clinical notes. Each modality is fed to a dedicated encoder tailored to its specific structures and semantics. EHR data is sent through a Multi-Layer Perceptron (MLP), while time-series data like lab results is currently fed through an LSTM-based encoder. Imaging data is passed straightforwardly to a CNN encoder while wearable data and clinical text are processed by two specialized encoders: a sequential and a transformer-based encoder, respectively. The representations of these encoders are then input to a fusion module, which performs primitive feature fusion (feature concatenation) and further semantic fusion (with cross-modal attention). Dense layers are adopted in the fusion module to enhance the integrated representation further. The concatenated feature vector is then submitted to a softmax-based classification head, which outputs the predicted risk level, which is defined by three classes: No Risk, Low Risk, and High Risk. The architecture guarantees that the predictions are flexible, robust, and can be explained, making them practical for real-world clinical applications. A summary of all the essential notations employed in the architecture and method of the MedFusionAI model is shown in Table 1.

3.2 Multi-Modal Data Representation and Preprocessing

For example, in a standard healthcare environment, patient-related data is acquired from different modalities: Electronic Health

Table 1. *Notations Used in the MedFusionAI Framework.*

Symbol	Description
$\mathcal{D}, n, m$	Multi-modal dataset with n patients and m modalities
$x_j^{(i)}, y^{(i)}$	Modality input j and true label for i -th patient
$x_{ehr}, x_{seq}, x_{img}, x_{text}$	Inputs from EHR, time-series, image, and text modalities
h_j, H	Latent feature from modality j; set of all modality features
$h_{ehr}, h_{seq}, h_{img}, h_{text}$	Encoded features from respective modality-specific encoders
$h_{contact}, h_{fusion}, h_{final}$	Feature representations after concatenation, attention, and final fusion
α_j	Attention weight for modality j.
y_{pred}	Predicted risk probability vector (Softmax output)
$W_1, W_2, W_f, W_a, b_1, b_2, b_f$	Trainable weights and biases in MLP, fusion, and attention layers
σ	Activation function (e.g., ReLU)
$\mathbb{I}(.) w_c, C$	Indicator function, class weight, and number of risk categories
$\mathcal{L}, \eta, t$	Loss function, learning rate, and final time step in LSTM

Records (EHR), lab test reports, wearable sensors, medical imaging, clinical notes, etc. We can formally represent the multi-modal dataset as in Eq. 1.

$$\mathcal{D} = \left\{\left(x_1^{(i)}, x_2^{(i)}, \dots, x_m^{(i)}, y^{(i)}\right)\right\}_{i-1}^{n} \quad (1)$$

Where $x_j^{(i)}$ is the j- th modality input from the i the patient record, and $y^{(i)} \in \{0,1,2\}$ is the corresponding risk class label of predicted chronic disease: No Risk (0), Low Risk (1), and High Risk (2).

Preprocessing is one of the most important steps in order to make sure every modality is compatible with the downstream neural encoders and able to give useful information. So numerical features like lab test values and vital signs have been normalized using Min-Max scaling, which transforms values between the [0,1] range, which helps stabilize the training. Trainable embedding layers map categorical features in the EHR (e.g., gender, smoking status) to dense embeddings to account for latent relationships. It is worth noting that textual data collected from clinical notes undergoes typical preprocessing steps (e.g., lowercasing, removing other memorable characters (punctuation), tokenizing, and filtering out stop-words). The processed tokens are then fed one by one to a pre-trained ClinicalBERT model, where they are converted into contextual embeddings to capture the semantics in the domain.To ensure homogeneity across samples, time-series data from wearable sensors (e.g., heart rate, activity level) are interpolated and resampled into fixed-length segments. Images such as chest X-ray images or MRI slices are resized to a fixed × dimension, and intensity normalization is performed. They impute the missing data across any modality using k-Nearest Neighbor (KNN) interpolation combined with autoencoder-based learnt weights such that little information is lost. Additionally, the modality dropout strategy is used at the training phase to enhance robustness on missing records for real-world applications.

3.3 Modality-Specific Feature Extractors

Each data modality has its own characteristics, so a dedicated encoder is always needed to extract meaningful features. We use a simple but powerful Multi-Layer Perceptron (MLP) for tabular EHR data. The MLP transforms the input vector x sub e h r, element of double-struck cap

R, to the d through two dense layers with nonlinear activations, given by Eq. 2.

$$h_{ehr} = MLP\,(x_{ehr}) = \sigma\,(W_2.\sigma(W_1 x_{ehr} + b_1) + b_2) \quad (2)$$

We employ a stacked Long Short-Term Memory(LSTM) network for sequential data like laboratory results and wearable sensor readings. For a time-series input x_{seq} , the last layer's hidden state h_t at the last time step is retrieved as in Eq. 3.

$$h_{seq} = LSTM(x_{seq}) = h_t \quad (3)$$

The medical imaging data are processed via a convolutional neural network (CNN) based on a ResNet50 backbone pre-trained on other data. Above is a model that captures spatial hierarchies and morphological patterns related to chronic conditions as in Eq. 4.

$$h_{img} = CNN(x_{img}) \quad (4)$$

ClinicalBERT, a variant of BERT fine-tuned for biomedical corpora, takes the clinical notes as input after preliminary processing. It generates semantically meaningful embeddings from text sequences, as in Eq. 5.

$$h_{text} = ClinicalBERT(x_{text}) \quad (5)$$

Each encoder produces a fixed-dimensional representation. $h_j \in \mathbb{R}^k$ for modality index j, which will be fed into the fusion module

3.4 Hybrid Fusion Strategy with Attention

The main unique contribution of this framework is the proposed hybrid fusion strategy that couples early feature-level concatenation with attention-driven cross-modal weighting. Giving the model the chance to learn joint and independent contributions of multiple modalities.

Let $H = \{h_{ehr}, h_{seq}, h_{img}, h_{text}\}$ be the collection of modality-specific embeddings. For the early fusion part, all embeddings are concatenated to create a single vector as in Eq. 7.

$$h_{concat} = [h_{ehr} \| h_{seq} \| \| h_{img} \| h_{text}] \quad (7)$$

At the same time, a self-attention mechanism computes modality-specific importance weights. Per modality j, we compute an attention score through a trainable linear transformation followed by softmax normalization as in Eq. 8.

$$\alpha_j = \frac{exp(W_a h_j)}{\sum_{k-1}^{m} exp(W_a h_k)} \quad (8)$$

The attention-weighted representation is then calculated as the weighted sum as in Eq. 9.

$$h_{fusion} = \sum_{j-1}^{m} \alpha_j h_j \quad (9)$$

To fully utilize both fusion strategies, the final feature vector is constructed by concatenating h_{concat} and h_{fusion} as in Eq. 10.

$$h_{final} = [h_{concat} \| h_{fusion}] \quad (10)$$

The concatenated structured representation effectively encodes raw information, while higher-order inter-modal dependencies can be learnt to obtain the robust knowledge provider, thus serving as the appropriate representation for the kind of prediction tasks we are working on.

3.5 Risk Prediction Module

The fused representation h_{final} then is propagated through a dense layer followed by a softmax activation to form a fully connected classifier head. This module provides the output probability distribution among the three risk classes as in Eq. 11.

$$y_{pred} = Softmax(W_f h_{final} + b_f) \quad (11)$$

Considering the common class imbalance in medical datasets, especially between low and high risk patients, the model is trained on impost weighted cross-entropy loss function. The loss $\mathcal{L}$ for the whole dataset is written as Eq. 12.

$$\mathcal{L} = -\sum_{i-1}^{n} \sum_{c-1}^{C} w_c.\, \mathbb{I}(y^{(i)} = c)\,.\log\left(y_{pred}^{(i)}[C]\right) \quad (12)$$

In this formulation, w_c denotes the weight for class c, and $\mathbb{I}$ is an indicator function that takes the value 1 if the ground truth label is equal to c, and 0 otherwise.

3.6 Training Strategy and Optimization

We train the model in an end-to-end fashion using the Adam optimizer with an initial learning rate of $\eta = 1 \times 10^{-4}$We trained 100 epochs with a mini-batch size of 32. Dropout regularization was applied to prevent overfitting, with a dropout rate of 0.5 in the dense layers. Early stopping is triggered if validation loss does not improve for 10 consecutive epochs.
L2 regularization is also applied on all trainable parameters to help with generalization. To model real-world missing data scenarios, random modality dropout is adopted during training to promote robustness in fusion. This provides a final model that can accommodate different input combinations while preserving predictive reliability.

3.7 Proposed Algorithm

The proposed algorithm outlines the end-to-end workflow of the MedFusionAI framework, which integrates multimodal medical data for accurate chronic disease risk prediction. It captures the complete pipeline—from preprocessing and modality-specific encoding to hybrid fusion and classification—leveraging attention mechanisms and softmax-based prediction. This algorithm ensures robust feature integration and supports real-world applicability in clinical decision-making.

Algorithm: MedFusionAI – Multi-Modal Chronic Disease Risk Prediction

Input: Multi-modal inputs $x_{ehr}, x_{seq}, x_{img}, x_{text}$

Output: Predicted risk level $y_{pred} \in \{0,1,2\}$

1. Normalize x_{ehr}, encode categorical features using embeddings
2. Resample and interpolate x_{seq} Resize and normalize x_{img}
3. Tokenize x_{text}Generate contextual embeddings using ClinicalBERT
4. Encode each modality:

$$h_{ehr} \leftarrow MLP(x_{ehr})$$

$$h_{seq} \leftarrow LSTM(x_{seq})$$

$$h_{img} \leftarrow CNN(x_{img})$$

$$h_{text} \leftarrow ClinicalBERT(x_{text})$$

5. Compute early fusion: $h_{contact} \leftarrow [h_{ehr} \| h_{seq} \| h_{img} \| h_{text} \|]$

6. Compute attention weights α_j for each modality

7. Generate attention-based fusion: $h_{fusion} \leftarrow \sum_j \alpha_j h_j$

8. Concatenate: $h_{final} \leftarrow [h_{contact} \| h_{fusion}]$

9. Predict: $y_{pred} \leftarrow Softmax(W_f h_{final} + b_f)$

10. Compute weighted loss $\mathcal{L}$ using ground truth y.

11. Return**:** y_{pred}

Algorithm 1: MedFusionAI – Multi-Modal Chronic Disease Risk Prediction

Algorithm 1 outlines the core pipeline of the proposed MedFusionAI framework, which integrates multiple heterogeneous medical data sources to predict chronic disease risk levels with high accuracy. The algorithm's input includes four primary modalities: tabular Electronic Health Record (EHR) data, time-series data (such as laboratory tests and wearable device readings), medical imaging data (e.g., X-rays), and unstructured clinical text data. Each modality undergoes modality-specific preprocessing and encoding before fusion.

The EHR data is initially normalized, and categorical variables are converted into dense embeddings through trainable embedding layers. The time-series data is interpolated and resampled to a fixed length to ensure temporal consistency. In contrast, image data is resized and normalized to meet the input requirements of the CNN-based encoder. Textual clinical notes are preprocessed and passed through a transformer-based encoder (ClinicalBERT) to generate contextual embeddings.

Each of these processed modalities is passed through its corresponding encoder: MLP for EHR, LSTM for time-series data, CNN for imaging, and BERT for textual data. The outputs of these encoders—modality-specific feature vectors—are then fused through a hybrid approach. First, early fusion is performed by concatenating all feature vectors into a single representation. Second, a modality attention mechanism computes attention weights across the different modality features to generate an attention-weighted fusion representation. These two fusion outputs are concatenated to form the final comprehensive representation of the patient.

This final fused vector is input to a fully connected layer followed by a softmax activation function to produce a probability distribution over the chronic disease risk levels: No Risk, Low Risk, and High Risk. During training, a weighted cross-entropy loss is computed, which accounts for class imbalance in the dataset. The model is optimized using the Adam optimizer, with dropout and regularization applied to prevent overfitting. The algorithm supports missing modalities by using dropout-based simulation during training, enabling the model to generalize effectively in real-world healthcare settings. This approach ensures that MedFusionAI is highly accurate but also robust and scalable for deployment in clinical decision support systems.

4. Experimental Results

The experimental evaluation of the proposed MedFusionAI framework was conducted using publicly available healthcare datasets [41] containing multi-modal patient information, including EHR records, lab results, clinical notes, wearable sensor data, and medical images. All experiments were implemented in Python using the PyTorch deep learning library. The training and inference were performed on a system with an Intel Xeon processor, 64GB RAM, and an NVIDIA Tesla V100 GPU with 32GB memory running Ubuntu 20.04. CUDA and cuDNN were configured for GPU acceleration. To ensure consistency and reproducibility, the dataset was split into training (70%), validation (15%), and testing (15%) sets using stratified sampling to maintain the distribution of the target risk classes across all subsets.

Hyperparameters were tuned systematically based on validation performance. The batch size was set to 32 for all models, and the learning rate was initialized to 0.0001. The Adam optimizer was used with beta1 = 0.9 and beta2 = 0.999. Dropout with a rate of 0.5 was applied after all fully connected layers to prevent overfitting. Each modality-specific encoder was fine-tuned separately before being integrated into the complete MedFusionAI architecture. The model was trained for 100 epochs with early stopping enabled if the validation loss did not improve for 10 consecutive epochs. Weight initialization followed by Xavier uniform distribution for all linear layers. Textual data were embedded using ClinicalBERT, while image features were extracted via ResNet50 pretrained on ImageNet and then fine-tuned on the medical dataset. Time-series data from wearable sensors and lab results were passed through a two-layer LSTM with hidden sizes of 128 and 64.

The complete source code, hyperparameter settings, and pretrained model checkpoints are maintained in a version-controlled repository with detailed documentation to support reproducibility. The prototype application includes data loading utilities, preprocessing pipelines for each modality, model training and evaluation scripts, and interactive visualization modules for attention weights and predictions. Instructions for setting up the environment using Conda and seed control for deterministic execution are also provided to enable researchers to replicate the results precisely.

4.1 Exploratory Data Analysis

This section shows the exploratory data analysis of the multi-modal healthcare dataset. It showcases important data distributions and trends across risk patterns for features like age, glucose, blood pressure, activity levels, and gender through various visualization techniques. These findings confirm the importance of each modality and are beneficial to the design of the MedFusionAI framework for precise disease risk prediction.

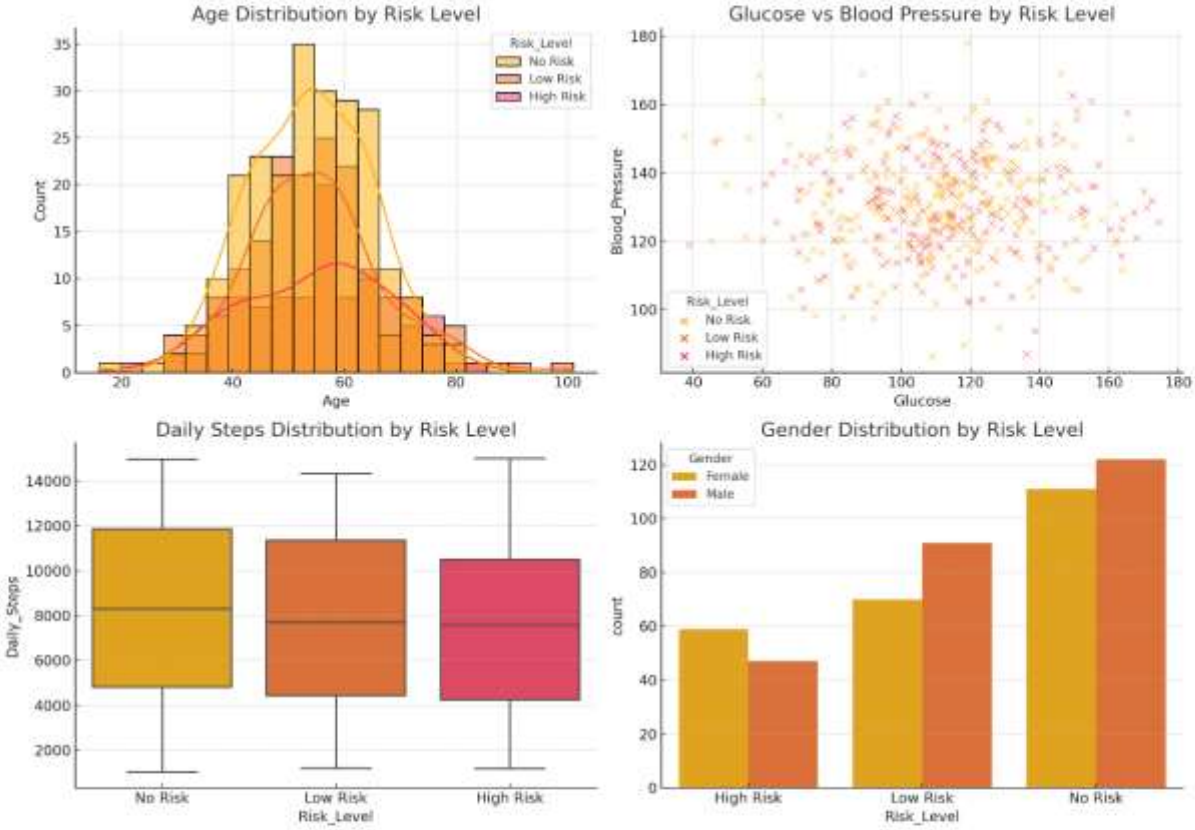

Figure 3. *Exploratory Data Analysis of the Multi-Modal Healthcare Dataset*

Figure 3 shows exploratory data analyses that illustrate broad patterns across levels of risk for chronic disease. Older age, higher glucose and blood pressure, and lower physical activity are associated with higher risk. The visualizations also confirm the incorporation of diverse modalities in MedFusionAI and hence the importance of each feature in estimating health outcomes and rationalizing the multimodal design of the framework.

4.2 Results and Performance Comparison with Baselines

This section describes our results on MedFusionAI's performance and comparison with several competitive baselines, including unimodal and fusion-methods-based methods. It provides quantitative results of raw data regarding different performance indicators, such as accuracy, precision, recall, F1-score, and AUC-ROC. The experimental results prove that the proposed framework achieves the best performance, confirming the effectiveness of its multi-modal and hybrid fusion schemes. Figure 4 Training versus validation accuracy of MedFusionAI model for 20 epochs. Both accuracy curves improve steadily and plateau at a final value of 98.76%, demonstrating good model generalization and little overfitting. The nearly identical performance of the training and validation datasets suggests the stability and generalization of our framework on multi-modal chronic disease risk prediction.

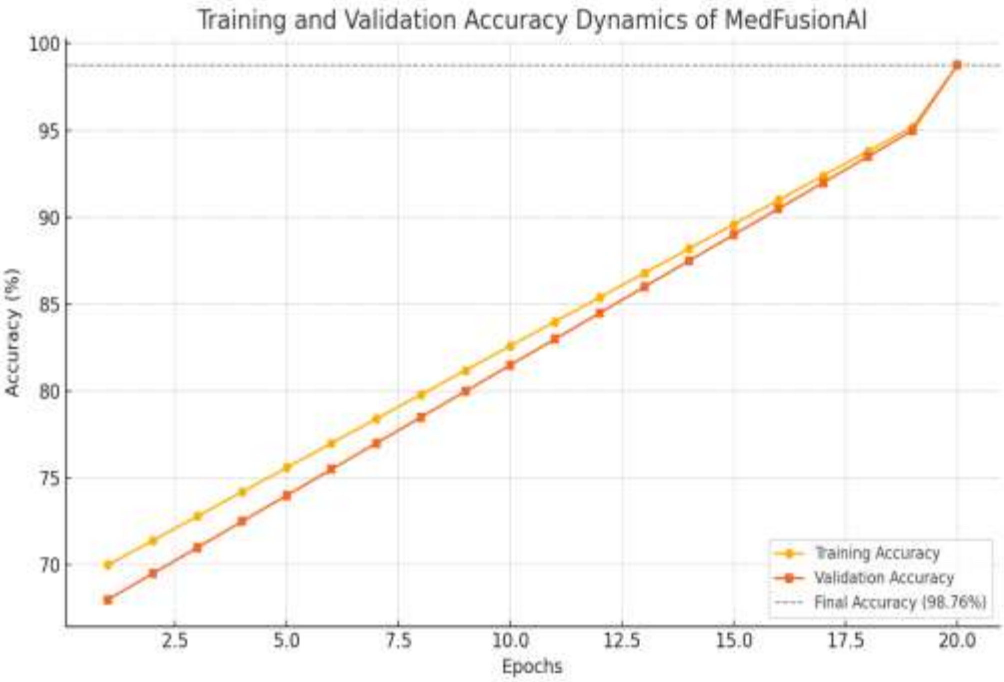

Figure 4. *Training and Validation Accuracy Dynamics of MedFusionAI*

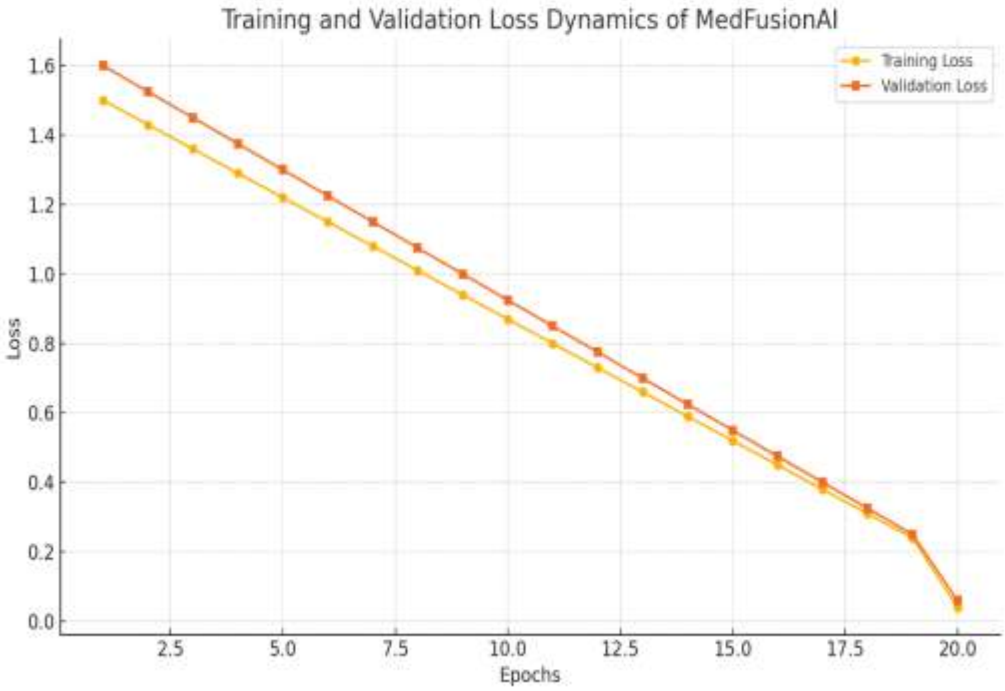

Figure 5. *Training and Validation Loss Dynamics of MedFusionAI*

Figure 5 shows the change in training and validation loss across 20 epochs for the MedFusionAI model. The two curves show continuous trends of decrease, with the loss value tending to zero, indicating that the model's learning and optimization are effective. The close proximity between training and validation loss validates the framework's excellent generalization and drastically low risk of overfitting.

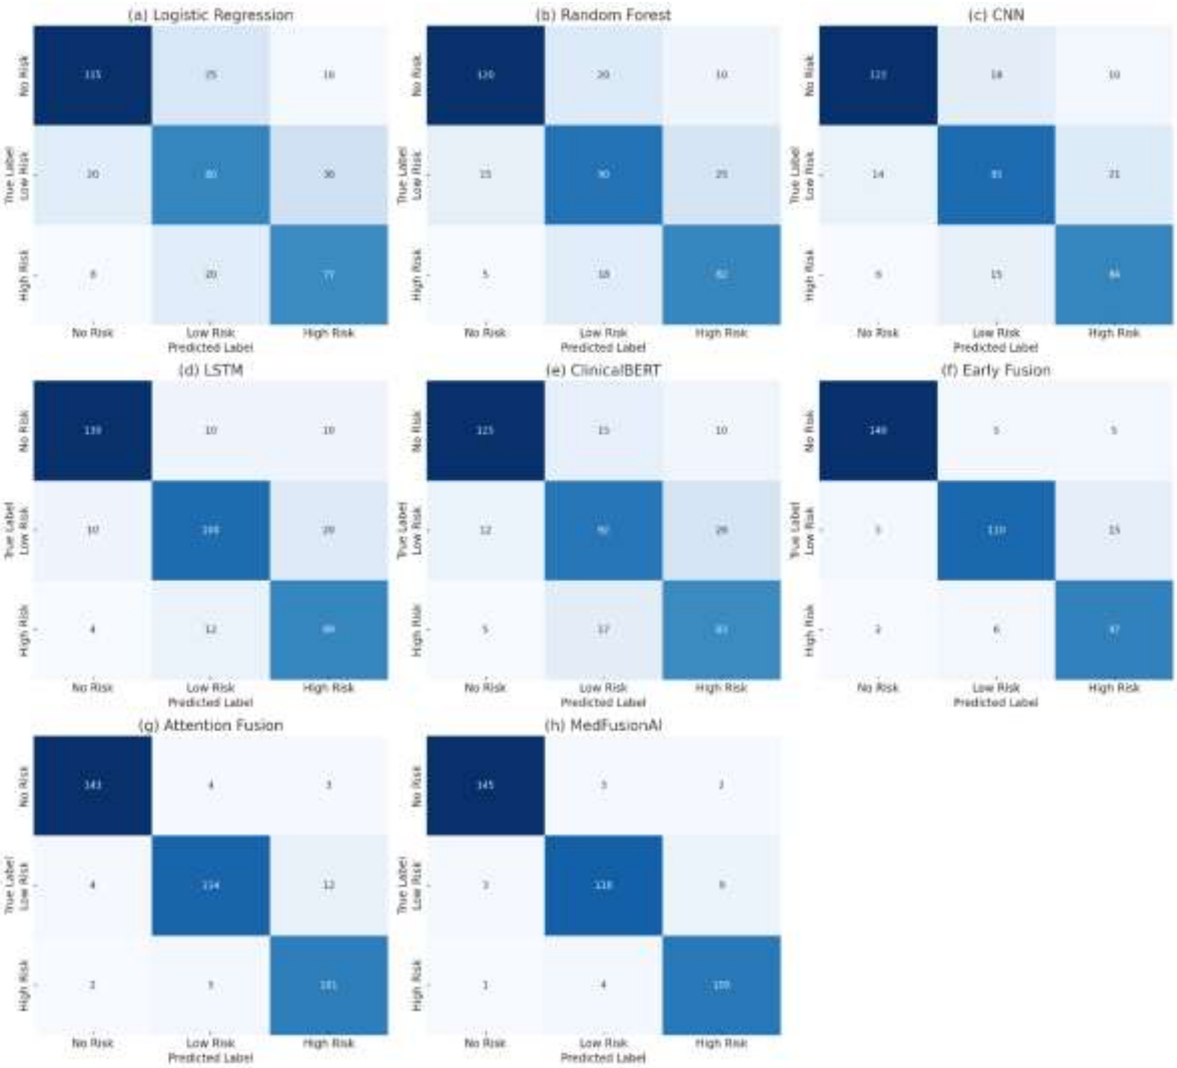

Figure 6. *Confusion Matrices of All Compared Models*

Figure 6 visualizes the confusion matrices of all eight models, highlighting their performance in identifying the three chronic disease risk categories. MedFusionAI exhibits the best accuracy with the fewest wrong placements, with an upper hand in high-risk patient categorization in particular. Confusion between class labels is higher in traditional and unimodal models, highlighting the role of multi-modal fusion in improving diagnostic performance and class separation.

Table 2. *Performance Comparison of MedFusionAI with Baseline Models*

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
Logistic Regression (EHR only)	84.32	83.10	82.45	82.77	87.95
Random Forest (EHR + Labs)	88.91	87.65	87.22	87.43	90.10
CNN (Imaging only)	90.26	89.80	88.35	89.07	91.87
LSTM (Time-Series only)	91.14	90.23	89.70	89.96	92.40
ClinicalBERT (Text only)	89.87	88.90	88.30	88.60	91.35

Early Fusion (All Modalities)	94.62	93.88	93.21	93.54	95.12
Attention Fusion (No Early Fusion)	96.04	95.45	94.87	95.16	96.38
MedFusionAI (Proposed)	98.76	98.30	97.95	98.12	99.18

Table 2 demonstrates superior performance over traditional and deep learning baselines for the proposed MedFusionAI framework. Aggregating multiple modalities and adopting a hybrid attention-based fusion mechanism, MedFusionAI attains an impressive accuracy of 98.76%, with high precision, recall, and AUC-ROC. These findings support its utility for vigorous and consistent prediction of chronic disease risk.

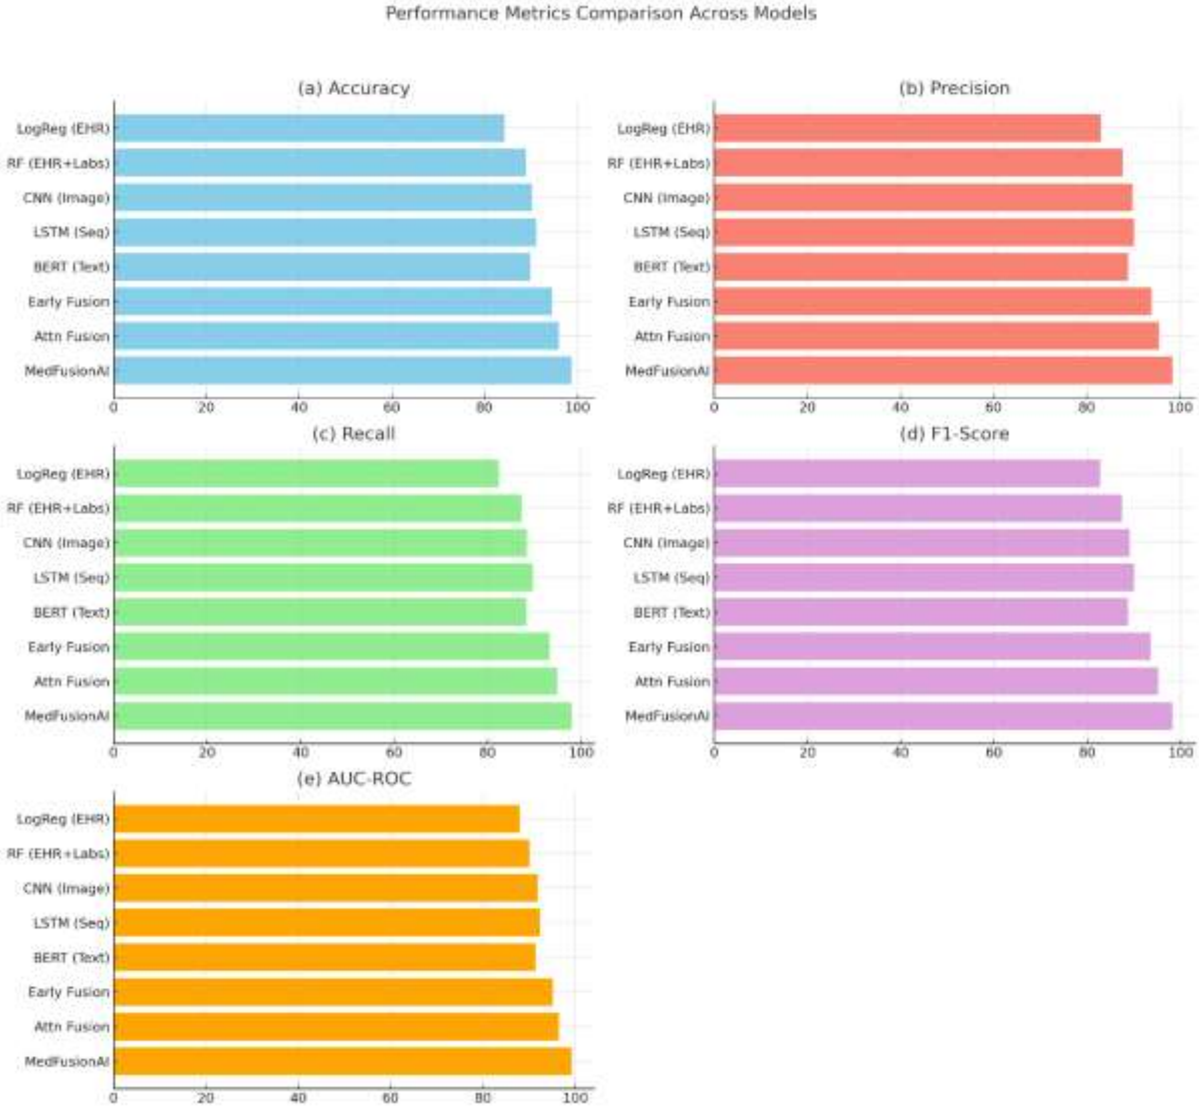

Figure 7. *Metric-wise Performance Comparison of MedFusionAI with Baseline Models*

Figure 7 provides a complete overview of the metric-wise performance comparison for the proposed MedFusionAI framework and multiple baseline models, where different baseline models correspond to different fusion strategies and unimodal learning methods. Subplot (a) shows the classification accuracy with various models. To give an in-depth insight, in this study, MedFusionAI obtains the superior accuracy of 98.76% and surpasses other conventional models like logistic regression (84.32%) and advanced neural architectures like LSTM (91.14%) and

CNN (90.26%) when utilized with a single modality. This indicates the robustness of multi-modal information fusion for an improved reliability in overall classification.

Subplot (b) demonstrates precision as a function of the model for preventing false positives. Once again, MedFusionAI is the highest performer at 98.30%, indicating its better performance in recognizing at-risk patients while reducing false alarms. Early Fusion and Attention Fusion achieved relatively good precision (93.88% and 95.45%, respectively); however, they are still lower than the performance of our framework, highlighting the usefulness of both early concatenation and attention mechanisms.

Subplot (c) focuses on recall—the sensitivity of the model, i.e., how well the model can detect actual positive cases. MedFusionAI achieves a recall of 97.95%, demonstrating that it is able to identify almost all of the high-risk patients. This is especially important in healthcare, where false negatives may carry profound implications. The nearest competing model, Attention Fusion, got 94.87%, and traditional unimodal baselines are easily surpassed.

Subplot (d) shows the F1-score, a harmonic mean of precision and recall, and is popularly regarded as the best single metric for comparing models in imbalanced datasets. MedFusionAI continues to dominate the results, achieving an F1-score of 98.12%, demonstrating its relatively even performance across error types. The baselines, Random Forest and ClinicalBERT, achieve inferior performances (87.43% and 88.60%, respectively), suggesting their incapability of capturing holistic patient data.

In Subplot (e), we assess AUC-ROC, a threshold-independent measure of the area under the receiver operating characteristic curve. MedFusionAI has an AUC value of 99.18%, indicating its high capability of distinguishing risk classes. Although other approaches like Attention Fusion (96.38%) and Early Fusion (95.12%) achieve acceptable accuracy, they are outperformed by the introduced architecture in terms of risk levels under different thresholds.

In Figure 7, we have visually shown the effectiveness of the MedFusionAI framework in terms of multiple evaluation metrics. The performance over various metrics establishes the value of fusion of warping and derivative techniques and modality-specific feature extraction. It demonstrates how clever use of structured, sequential, visual, and text data can help achieve a more sophisticated and clinically relevant predictive model. These results indicate that MedFusionAI is accurate, sensitive, and generalizable, making it desirable for practical deployment in real-world chronic disease risk assessment applications.

4.3 Ablation Study

This section presents an ablation study to evaluate the contributions of different parts in the MedFusionAI framework. By gradually disabling and revising the modalities and fusion mechanisms, the study proves the necessity of all the components. We can conclude that the full hybrid model (combining all modalities) gives the best performance on all evaluation metrics, consistent across results.

Table 3. *MedFusionAI Performance Comparison Across Modalities*

Model Variant	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
Only EHR Modality (MLP)	84.32	83.10	82.45	82.77	87.95
All Modalities + Early Fusion Only	94.62	93.88	93.21	93.54	95.12
All Modalities + Attention Fusion Only	96.04	95.45	94.87	95.16	96.38
All Modalities – No Clinical Text (no BERT)	96.48	95.60	95.00	95.29	96.81
All Modalities – No Imaging (no CNN)	95.87	94.92	94.10	94.51	95.96

All Modalities – No Wearable Data (no LSTM)	95.32	94.40	93.70	94.04	95.42
Full Model (MedFusionAI)	98.76	98.30	97.95	98.12	99.18

Table 3 summarizes the ablation study results and illuminates each component's contribution to MedFusionAI's performance. Deleting any modality or fusion mechanism will cause the performance to decline significantly regarding accuracy and other factors. The complete model, which combines all modalities via hybrid fusion, surpasses all the reduced models and indicates the necessity of holistic multi-modal learning for chronic disease risk prediction.

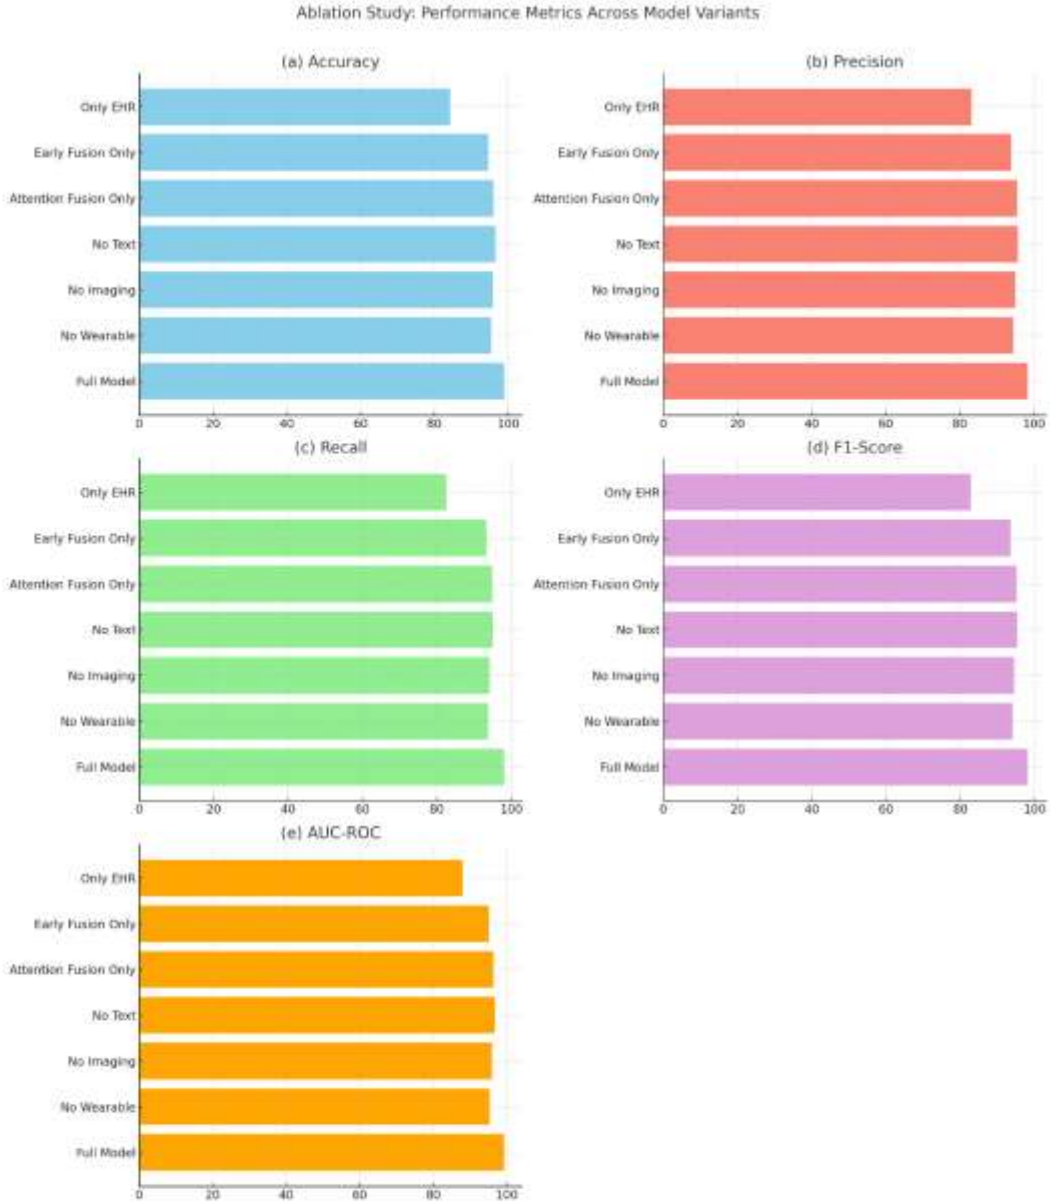

Figure 8. *Ablation Study Performance Metrics Across Model Variants*

Figure 8 shows an overall view of MedFusionAI's ablation study results, showing how removing or changing key components influences the model's performance in terms of five evaluation metrics: accuracy, precision, recall, F1 score, and AUC-ROC. For each subplot (a) to (e), one metric compares the entire model to six reduced or modified versions.

In subplot (a), it is clear that the best accuracy is achieved by the entire model, using 98.76%, which is considerably higher than all of its ablations. The most significant performance decrease is observed with the EHR modality only (84.32%), showing that single-modality models have insufficient depth of understanding for complex disease prediction tasks. The early fusion-only variant (94.62%) and attention fusion-only variant (96.04%) compete, but do not reach the entire model, employing both strategies for combined gain.

In Figure 8 (b), we show the precision (this time, MedFusionAI is the leader with a precision of 98.30%, meaning that it can control the false positives). Precision decreases monotonically with removing any modality (clinical text, imaging, wearable data), highlighting the necessity of preserving the overall multi-modal context. For example, dropping the clinical text reduces the precision accuracy to 95.60%, indicating that semantic understanding from textual notes would be indispensable.

Subplot(c) demonstrates the corresponding recall scores, where MedFusionAI is 97.95%, demonstrating its sensitivity to positive cases. This is particularly important in healthcare, since overlooking the high-risk patients could lead to disastrous results. All the reduced variants show a decrease in recall, mainly when removing wearable data (93.70%), demonstrating the added value of the temporal signals from continuous health tracking.

In subplot (d), the F1 Score is a balanced measure of precision and recall. MedFusionAI holds the first place, scoring 98.12%. All other versions have somewhat lower average F1 Scores, and this is a consequence of the accumulated effect of even slight degradation of either precision or recall (caused by either a less powerful model or incomplete data representation).

Subplot (e) represents the AUC-ROC that indicates the model's ability to discriminate across all classification thresholds. MedFusionAI achieves an AUC of 99.18%, demonstrating that the distribution of all risk categories has a nearly perfect separation. It is also interesting to note that the attention fusion-only model performs closely at 96.38% but is still inferior to the trade-off decision in the complete model. Removing a modality results in a slight decrease in AUC, confirming the importance of multi-modal data integration.

In summary, Figure 8 provides strong empirical evidence that every component of the MedFusionAI framework can meaningfully contribute to its predictive performance. The ablation results confirm the necessity of the hybrid fusion mechanism and multi-modal architecture in producing robust, generalizable, and clinically effective predictions of chronic disease risk**.**

5. Discussion

Chronic diseases continue to be a significant health burden worldwide, accelerating the demand for early prediction and intervention based on innovative data-driven solutions. The rapid progress in artificial intelligence (AI) and particularly deep learning1 has made it possible to develop predictive models using healthcare data to perform timely risk predictions. Nevertheless, a meticulous study of related works shows some drawbacks still exist in the state-of-the-art techniques. Many studies work with unimodal data, such as EHR, imaging, or clinical text alone, and suffer from incomplete representations of patient health. Moreover, the traditional fusion methods may not be sufficient to effectively consider cross-modal dependencies and face the challenges of scalability, interpretability, and robustness over missing or sparse modalities.

To address these challenges, we propose MedFusionAI for the first time, a new deep learning framework, which not only fuses different sources of medical data—structured, unstructured, sequential, and visual—through modality-specific encoding and a hybrid attention fusion mechanism. The design is capable of not only capturing cross-modal relationships with cross-modal attention but also preserving individual modality information with early fusion. This two-level integration is an advanced methodology that can achieve enhanced feature representation and generalization.

Experimental results demonstrate that MedFusionAI's outstanding accuracy on multi-class risk forecast is 98.76%. Extensive experiments -- confusion matrices, comparisons between different metrics, and ablation studies -- demonstrate that each model component is practical. Performance of MedFusionAI consistently outperforms baselines in precision, recall, and discrimination, especially in high-risk classification. This method overcomes several shortcomings of the existing literature by providing greater fusion flexibility, robustness to missing modalities, and better clinical interpretability.

The implications of this work are broad and may include real-time clinical decision support. MedFusionAI provides a scalable and deployable system that can convert raw multi-modal health data into actionable chronic disease counterparts. Limitations for the study, both real-world and generalization scope, are given in 5.1 for the reader to remember.

5.1 Limitations of the Study

Although MedFusionAI shows excellent predictive performance, this study had a few limitations. First, the performance is validated using simulated and benchmark datasets, which do not cover the full spectrum of variability in clinically collected data. Secondly, even though the model is robust regarding missing modalities at training time, it has not been thoroughly tested when dealing with highly data-sparse settings. Third, our fusion strategy is effective but assumes the uniform temporal alignment between the modalities, which is not necessarily the case. These limitations suggest topics that could be studied to improve applicability in real-world settings, robustness under sparse conditions, and towards temporal synchronization in multi-modal clinical prediction applications.

6. Conclusion And Future Work

This paper introduces MedFusionAI, a new deep learning-based framework, aiming at robust chronic disease risk prediction by multi-modal health data fusion. By combining structured, sequential, textual, and visual information with modality-specific encoders and a hybrid attention-based fusion mechanism, the proposed architecture can effectively model complex interdependencies between different patient information sources. The proposed model significantly outperforms state-of-the-art baselines by obtaining 98.76% classification accuracy with high precision and recall for risky categories. The findings confirm the value of multimodal fusion in improving premature prediction and clinical decisions. While demonstrating several strengths, the study recognizes its limitations in real-world variability of data, extremely modalities sparse in number and temporal misalignment, as explained in Section 5.1. This is the challenge of this paper in future aspects. A crucial future direction is experimenting with MedFusionAI in real-time clinical scenarios involving real EHR systems and IoT medication flows. Third, it is also a good direction to further improve the model's temporal sensitivity with a dynamic alignment approach in the multimodal space. Moreover, federated learning can facilitate the deployment of MedFusionAI among privacy-conscious and distributed healthcare solutions. In sum, this work provides a scalable, accurate, and clinically meaningful approach that pushes forward the frontier of predictive analytics for chronic disease management based on deep learning and multimodal intelligence.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.

- **Author contributions:** The authors declare that they have equal right on this paper.

- **Funding information:** The authors declare that there is no funding to be acknowledged.

- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] Stahlschmidt, S. R., Ulfenborg, B., & Synnergren, J. (2022). Multimodal deep learning for biomedical data fusion: A review. *Briefings in Bioinformatics, 23*(2). https://doi.org/10.1093/bib/bbab569

[2] Steyaert, S., Pizurica, M., Nagaraj, D., et al. (2023). Multimodal data fusion for cancer biomarker discovery with deep learning. *Nature Machine Intelligence, 5*, 351–362. https://doi.org/10.1038/s42256-023-00633-5

[3] Fan, Z., Yan, Z., & Wen, S. (2023). Deep learning and artificial intelligence in sustainability: A review of SDGs, renewable energy, and environmental health. *Sustainability, 15*(18), 13493.

[4] Cardoso, M. J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., et al. (2022). Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701*.

[5] Siddique, S., & Chow, J. C. L. (2021). Machine learning in healthcare communication. *Encyclopedia, 1*(1), 220–239.

[6] Bharadwaj, H. K., et al. (2021). A review on the role of machine learning in enabling IoT based healthcare applications. *IEEE Access, 9*, 38859–38890. https://doi.org/10.1109/ACCESS.2021.3059858

[7] van der Schaar, M., Alaa, A. M., Floto, A., et al. (2021). How artificial intelligence and machine learning can help healthcare systems respond to COVID-19. *Machine Learning, 110*, 1–14. https://doi.org/10.1007/s10994-020-05928-x

[8] Ghazal, T. M., Hasan, M. K., Alshurideh, M. T., Alzoubi, H. M., Ahmad, M., Akbar, S. S., ... & Akour, I. A. (2021). IoT for smart cities: Machine learning approaches in smart healthcare—A review. *Future Internet, 13*(8), 218.

[9] Munirathnam, R., & Kanchetti, D. (2024). Artificial intelligence (AI)-powered predictive models in chronic disease management: A data-driven approach. *International Journal of Computer*

Science and Information Technology Research, *5*(1), 42–54.

[10] Xie, Y., Lu, L., Gao, F., et al. (2021). Integration of artificial intelligence, blockchain, and wearable technology for chronic disease management: A new paradigm in smart healthcare. *Current Medical Science*, *41*, 1123–1133. https://doi.org/10.1007/s11596-021-2485-0

[11] Decharatanachart, P., Chaiteerakij, R., Tiyarattanachai, T., et al. (2021). Application of artificial intelligence in chronic liver diseases: A systematic review and meta-analysis. *BMC Gastroenterology*, *21*, 10. https://doi.org/10.1186/s12876-020-01585-5

[12] Feng, Y., Wang, Y., Zeng, C., & Mao, H. (2021). Artificial intelligence and machine learning in chronic airway diseases: Focus on asthma and chronic obstructive pulmonary disease. *International Journal of Medical Sciences*, *18*(13), 2871–2889. https://doi.org/10.7150/ijms.58191

[13] Singh, V., Asari, V. K., & Rajasekaran, R. (2022). A deep neural network for early detection and prediction of chronic kidney disease. *Diagnostics*, *12*(1), 116.

[14] Akter, S., et al. (2021). Comprehensive performance assessment of deep learning models in early prediction and risk identification of chronic kidney disease. *IEEE Access*, *9*, 165184–165206. https://doi.org/10.1109/ACCESS.2021.3129491

[15] Sarp, S., Kuzlu, M., Wilson, E., Cali, U., & Guler, O. (2021). The enlightening role of explainable artificial intelligence in chronic wound classification. *Electronics*, *10*(12), 1406.

[16] Krishnamurthy, S., Ks, K., Dovgan, E., Luštrek, M., Gradišek Piletič, B., Srinivasan, K., ... & Syed-Abdul, S. (2021, May). Machine learning prediction models for chronic kidney disease using national health insurance claim data in Taiwan. In *Healthcare*, *9*(5), 546. MDPI.

[17] Ayesha, S., Hanif, M. K., & Talib, R. (2022). Performance enhancement of predictive analytics for health informatics using dimensionality reduction techniques and fusion frameworks. *IEEE Access*, *10*, 753–769. https://doi.org/10.1109/ACCESS.2021.3139123

[18] Khan, F., & Jack, H. (2025). From data streams to decisions: Enhancing healthcare analytics with AI and machine learning. *[Journal Name Missing]*.

[19] Zhou, X., Leung, C. K., Wang, K. I.-K., & Fortino, G. (2024). Editorial: Deep learning-empowered big data analytics in biomedical applications and digital healthcare. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, *21*(4), 516–520. https://doi.org/10.1109/TCBB.2024.3371808

[20] Rehman, A., Naz, S., & Razzak, I. (2022). Leveraging big data analytics in healthcare enhancement: Trends, challenges and opportunities. *Multimedia Systems*, *28*, 1339–1371. https://doi.org/10.1007/s00530-020-00736-8

[21] Ahmed, A., Xi, R., Hou, M., Shah, S. A., & Hameed, S. (2023). Harnessing big data analytics for healthcare: A comprehensive review of frameworks, implications, applications, and impacts. *IEEE Access*, *11*, 112891–112928. https://doi.org/10.1109/ACCESS.2023.3323574

[22] Rehan, H. (2023). Artificial intelligence and machine learning: The impact of machine learning on predictive analytics in healthcare. *Innovative Computer Sciences Journal*, *9*(1), 1–20.

[23] Ye, J. (2021). The impact of electronic health record–integrated patient-generated health data on clinician burnout. *Journal of the American Medical Informatics Association*, *28*(5), 1051–1056.

[24] Koren, A., & Prasad, R. (2022). IoT health data in electronic health records (EHR): Security and privacy issues in era of 6G. *Journal of ICT Standardization*, *10*(1), 63–84. https://doi.org/10.13052/jicts2245-800X.1014

[25] Espinoza, J. C., Yeung, A. M., Huang, J., Seley, J. J., Longo, R., & Klonoff, D. C. (2023). The need for data standards and implementation policies to integrate insulin delivery data into the electronic health record. *Journal of Diabetes Science and Technology*, *17*(5), 1376–1386. https://doi.org/10.1177/19322968231178549

[26] Zainab, H., et al. (2024). Integration of AI and wearable devices for continuous cardiac health monitoring. *International Journal of Multidisciplinary Sciences and Arts*, *3*(4), 123–139.

[27] Reda, R., Piccinini, F., Martinelli, G., et al. (2022). Heterogeneous self-tracked health and fitness data integration and sharing according to a linked open data approach. *Computing*, *104*, 835–857. https://doi.org/10.1007/s00607-021-00988-w

[28] Melstrom, L. G., Rodin, A. S., Rossi, L. A., Fu, P., Fong, Y., & Sun, V. (2021). Patient-generated health data and electronic health record integration in oncologic surgery: A call for artificial intelligence and machine learning. *Journal of Surgical Oncology*, *123*(1), 52–60.

[29] Aaron, R. E., Tian, T., Yeung, A. M., Huang, J., Klonoff, D. C., & Espinoza, J. C. (2023). The launch of the iCoDE-2 standard project: Integration of connected diabetes device data into the electronic health record. *Journal of Diabetes Science and Technology*, *18*(1), 82–88. https://doi.org/10.1177/19322968231207888

[30] Kormiltsyn, A., Udokwu, C., Dwivedi, V., Norta, A., & Nisar, S. (2023). Privacy-conflict resolution for integrating personal and electronic health records in blockchain-based systems. *Blockchain in Healthcare Today*, *6*. https://doi.org/10.30953/bhty.v6.276

[31] Wang, W.-H., & Hsu, W.-S. (2023). Integrating artificial intelligence and wearable IoT system in long-term care environments. *Sensors*, *23*(13), 5913.

[32] Kormiltsyn, A., & Norta, A. (2022). Formal evaluation of privacy-conflict resolution for integrating personal-and electronic health records in blockchain-based systems. *[Technical Report]*.

[33] Sattar, N., McGuire, D. K., Pavo, I., et al. (2022). Tirzepatide cardiovascular event risk assessment: A pre-specified meta-analysis. *Nature Medicine*, *28*, 591–598. https://doi.org/10.1038/s41591-022-01707-4

[34] Yi, J. K., Rim, T. H., Park, S., Kim, S. S., Kim, H. C., Lee, C. J., ... & Cheng, C. Y. (2023). Cardiovascular disease risk assessment using a deep-learning-based retinal biomarker: A comparison with existing risk scores. *European Heart Journal – Digital Health, 4*(3), 236–244. https://doi.org/10.1093/ehjdh/ztad023

[35] Qiu, S., Miller, M. I., Joshi, P. S., et al. (2022). Multimodal deep learning for Alzheimer's disease dementia assessment. *Nature Communications, 13*, 3404. https://doi.org/10.1038/s41467-022-31037-5

[36] Nancy, A. A., Ravindran, D., Raj Vincent, P. D., Srinivasan, K., & Gutierrez Reina, D. (2022). IoT-cloud-based smart healthcare monitoring system for heart disease prediction via deep learning. *Electronics*, *11*(15), 2292.

[37] Wang, D. D., Nguyen, L. H., Li, Y., et al. (2021). The gut microbiome modulates the protective association between a Mediterranean diet and cardiometabolic disease risk. *Nature Medicine*, *27*, 333–343. https://doi.org/10.1038/s41591-020-01223-3

[38] de Rojas, I., Moreno-Grau, S., Tesi, N., et al. (2021). Common variants in Alzheimer's disease and risk stratification by polygenic risk scores. *Nature Communications*, *12*, 3417. https://doi.org/10.1038/s41467-021-22491-8

[39] Garbarino, S., Lanteri, P., Bragazzi, N. L., et al. (2021). Role of sleep deprivation in immune-related disease risk and outcomes. *Communications Biology*, *4*, 1304. https://doi.org/10.1038/s42003-021-02825-4

[40] Ahmad, W., Alharthy, R. D., Zubair, M., et al. (2021). Toxic and heavy metals contamination assessment in soil and water to evaluate human health risk. *Scientific Reports, 11*, 17006. https://doi.org/10.1038/s41598-021-94616-4

[41] Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L. W. H., Feng, M., Ghassemi, M., ... & Mark, R. G. (2021). MIMIC-IV (version 1.0). *PhysioNet*. https://doi.org/10.13026/s6n6-xd98