



Multimodal Deep Learning Fusion of Medical Imaging and Electronic Health Records for Precision Disease Prediction

Authors

Xin Nie¹, Yi Jiang²

^{1,2}School of Computer Science and Engineering, Wuhan Institute of Technology, China

Abstract

Multimodal data integration has emerged as a critical approach for improving predictive performance in clinical decision-making systems. This study proposes a comprehensive deep learning framework that integrates medical imaging and electronic health records (EHRs) for precision disease prediction. The framework combines modality-specific encoders, including convolutional neural networks or Vision Transformers for imaging and transformer-based architectures such as BEHRT and Med-BERT for EHR data, with an attention-based fusion mechanism that captures cross-modal interactions. A hybrid fusion strategy is introduced to enable effective integration of heterogeneous representations while maintaining flexibility across different clinical tasks.

The model is evaluated on large-scale benchmark datasets, including MIMIC-CXR, CheXpert, and MIMIC-IV, using a multi-task learning setup for disease classification and risk prediction. Experimental results demonstrate that the proposed multimodal framework significantly outperforms unimodal baselines, achieving superior performance across accuracy, AUC, and F1-score metrics. Ablation studies further confirm the effectiveness of the hybrid fusion approach over traditional early and late fusion methods. In addition, the framework exhibits strong robustness to missing modalities and maintains consistent performance across varying dataset sizes, highlighting its scalability and practical applicability.

To enhance clinical interpretability, the model integrates Grad-CAM and SHAP for modality-specific explanations, providing insights into both imaging features and structured clinical variables. Furthermore, the framework supports privacy-preserving deployment through compatibility with federated learning and differential privacy techniques. These features position the proposed approach as a viable solution for real-world healthcare applications.

Overall, this study demonstrates that multimodal deep learning can significantly improve predictive accuracy and reliability in clinical settings, offering a scalable and interpretable framework for next-generation precision medicine systems.

Keywords: Multimodal Deep Learning; Medical Imaging; Electronic Health Records (EHR); Precision Medicine; Disease Prediction; Attention-Based Fusion; Hybrid Fusion Strategy; Clinical Decision Support; Explainable AI; Federated Learning; Differential Privacy; Healthcare Artificial Intelligence

1. Introduction

The rapid digitization of healthcare systems has resulted in an unprecedented accumulation of heterogeneous clinical data, particularly medical imaging and electronic health records (EHRs). These modalities provide complementary perspectives on patient health. Medical images such as radiographs, computed tomography (CT), and magnetic resonance imaging (MRI) offer detailed structural and spatial representations of anatomical abnormalities, while EHRs capture longitudinal clinical information, including laboratory results, vital signs, demographic attributes, and physician notes. Despite their complementary nature, most existing predictive models in clinical practice rely on a single modality, thereby limiting their ability to capture the full complexity of patient conditions (Huang et al., 2020; Khader et al., 2023).

Deep learning has significantly advanced disease prediction from medical imaging, particularly through convolutional neural networks (CNNs) and their extensions. Architectures such as ResNet have enabled robust feature extraction from high-dimensional imaging data, achieving state-of-the-art performance in tasks such as disease classification and lesion detection (He et al., 2016). More recently, transformer-based models, including Vision Transformers (ViTs), have demonstrated superior performance in capturing long-range dependencies in medical images, further enhancing diagnostic accuracy (Dosovitskiy et al., 2021). However, imaging-only models often lack access to critical contextual and temporal patient information that is essential for accurate prognosis and decision-making in clinical settings.

In parallel, predictive modeling using EHR data has evolved rapidly with the adoption of deep learning architectures capable of modeling sequential and high-dimensional clinical data. Models such as BEHRT and Med-BERT leverage transformer-based architectures to learn contextualized patient representations from longitudinal health records, enabling improved prediction of clinical outcomes such as disease onset and mortality risk (Li et al., 2020; Rasmy et al., 2021). Additionally, recurrent neural networks and attention-based models have been widely used to capture temporal dependencies in EHR data. Nevertheless, EHR-based models are inherently limited by their inability to directly incorporate physiological and anatomical evidence that is readily available in medical imaging (Shickel et al., 2018).

To overcome the limitations of unimodal systems, multimodal deep learning has emerged as a promising paradigm for integrating heterogeneous healthcare data. By combining imaging and EHR modalities, multimodal systems can construct a more comprehensive and holistic representation of patient health. Prior studies have demonstrated that multimodal fusion can significantly improve predictive performance in tasks such as mortality prediction, disease diagnosis, and clinical risk stratification (Huang et al., 2020; Jabbour et al., 2022; Hayat et al., 2022). In particular, hybrid models that jointly learn from radiological images and clinical variables have consistently outperformed unimodal baselines, highlighting the importance of cross-modal integration in precision medicine.

Despite these promising developments, several critical challenges remain. First, many existing multimodal approaches rely on simplistic fusion techniques, such as direct feature concatenation, which fail to capture complex interactions between modalities (Baltrušaitis et al., 2019). Second, real-world clinical datasets frequently exhibit missing, incomplete, or asynchronous modalities, which complicates the design of robust multimodal systems. Third, interpretability remains a major concern, as clinicians require transparent and explainable models to trust AI-driven predictions in high-stakes environments (Tonekaboni et al., 2019). Furthermore, issues related to data heterogeneity, scalability, and privacy constraints pose additional barriers to real-world deployment (Johnson et al., 2023).

In this study, we propose a unified multimodal deep learning framework for precision disease prediction that integrates medical imaging and EHR data through advanced fusion mechanisms. The proposed framework systematically explores multiple fusion strategies, including early, intermediate, and hybrid fusion, and evaluates their effectiveness across clinically relevant prediction tasks. In addition, the framework incorporates attention-based fusion and modular architecture design to enhance both predictive performance and interpretability.

The main contributions of this work are fourfold. First, we present a comprehensive multimodal fusion architecture that integrates imaging and EHR data within a unified deep learning framework. Second, we provide a systematic comparative analysis of fusion strategies to identify optimal design choices for precision disease prediction. Third, we demonstrate the effectiveness of multimodal integration through performance comparisons with unimodal baselines. Finally, we discuss practical considerations for clinical deployment, including interpretability, scalability, and privacy-aware learning.

2. Related Work

2.1 Deep Learning in Medical Imaging

Deep learning has transformed medical imaging analysis by enabling automated extraction of high-level features from complex visual data. Convolutional neural networks (CNNs) have been the dominant architecture in this domain, achieving state-of-the-art performance in tasks such as disease classification, lesion detection, and organ segmentation. The introduction of deep residual networks (ResNet) addressed the vanishing gradient problem and enabled the training of significantly deeper architectures, resulting in improved performance across various medical imaging benchmarks (He et al., 2016).

In addition to CNN-based models, encoder–decoder architectures such as U-Net have been widely adopted for medical image segmentation tasks due to their ability to preserve spatial resolution while capturing contextual information (Ronneberger et al., 2015). More recently, transformer-based models, particularly Vision Transformers (ViTs), have gained attention for their ability to model long-range dependencies in image data. These models have demonstrated competitive and, in some cases, superior performance compared to CNNs in medical imaging tasks, especially when trained on large-scale datasets (Dosovitskiy et al., 2021).

Despite these advances, imaging-based models remain limited in their ability to incorporate patient-specific contextual information, such as clinical history and laboratory findings, which are often critical for accurate diagnosis and prognosis. This limitation motivates the integration of imaging data with other clinical modalities.

2.2 Predictive Modeling with Electronic Health Records

Electronic health records provide a rich source of longitudinal patient data, enabling the modeling of disease progression and clinical outcomes over time. Traditional machine learning approaches have been largely replaced by deep learning models capable of capturing temporal dependencies and complex feature interactions.

Recurrent neural networks (RNNs), including long short-term memory (LSTM) networks, have been widely used for modeling sequential EHR data, particularly for tasks such as mortality prediction and disease onset forecasting (Shickel et al., 2018). However, these models often struggle with long-term dependencies and scalability.

To address these limitations, transformer-based architectures have been introduced for EHR representation learning. Models such as BEHRT and Med-BERT leverage self-attention mechanisms to capture contextual

relationships within longitudinal patient records, enabling improved predictive performance across a range of clinical tasks (Li et al., 2020; Rasmy et al., 2021). These approaches treat EHR data analogously to natural language, where medical codes are embedded into high-dimensional representations and processed using attention mechanisms.

Nevertheless, EHR-based models lack direct access to physiological and anatomical evidence captured in medical imaging, which limits their diagnostic capability in certain clinical scenarios.

2.3 Multimodal Learning in Healthcare

Multimodal deep learning aims to integrate heterogeneous data sources to improve predictive performance and provide a more comprehensive representation of patient health. In healthcare, this typically involves combining medical imaging with structured and unstructured EHR data.

Early studies have demonstrated that integrating imaging and clinical data leads to significant improvements in predictive accuracy compared to unimodal approaches. For example, Huang et al. (2020) proposed a multimodal fusion framework for clinical outcome prediction, showing that joint modeling of imaging and EHR data enhances performance in mortality prediction tasks. Similarly, Jabbour et al. (2022) and Hayat et al. (2022) highlighted the effectiveness of multimodal learning in improving diagnostic accuracy and robustness across various medical applications.

Multimodal fusion strategies are generally categorized into three types: early fusion, intermediate fusion, and late fusion. Early fusion involves combining raw input features before model training, allowing the model to learn cross-modal interactions from the outset. Intermediate fusion combines latent representations learned independently from each modality, enabling more flexible integration of heterogeneous data. Late fusion aggregates predictions from modality-specific models, often improving robustness but limiting deep interaction between modalities (Baltrušaitis et al., 2019).

While these approaches have shown promise, many existing studies rely on relatively simple fusion mechanisms, such as feature concatenation, which may not fully capture complex relationships between modalities. More advanced techniques, including attention-based fusion and cross-modal transformers, have been proposed to address these limitations, but their application in clinical settings remains limited.

2.4 Datasets and Benchmark Resources

The development of multimodal models in healthcare has been facilitated by the availability of large-scale public datasets that include both imaging and EHR data. Among the most widely used datasets are MIMIC-IV and MIMIC-CXR, which provide comprehensive clinical and radiological data for critical care patients (Johnson et al., 2023). These datasets enable the development and evaluation of multimodal models for tasks such as mortality prediction, disease diagnosis, and risk stratification.

CheXpert is another widely used dataset for chest radiograph analysis, providing labeled imaging data for multiple thoracic diseases (Irvin et al., 2019). When combined with EHR datasets, such resources offer a valuable foundation for multimodal learning research.

Despite the availability of these datasets, challenges remain in aligning imaging and EHR modalities due to differences in data structure, temporal resolution, and missing values. These challenges highlight the need for robust preprocessing and data integration techniques in multimodal frameworks.

2.5 Limitations of Existing Approaches

Although multimodal deep learning has demonstrated promising results, several limitations hinder its widespread adoption in clinical practice. First, many studies rely on datasets with limited sample sizes or incomplete modality coverage, which reduces the generalizability of the models. Second, the handling of

missing modalities remains an open challenge, as real-world clinical data are often incomplete or asynchronously recorded.

Third, interpretability remains a critical issue, as deep learning models are often perceived as “black boxes,” limiting clinician trust and adoption (Tonekaboni et al., 2019). Finally, issues related to data privacy, scalability, and computational cost present additional barriers to deployment in real-world healthcare systems (Johnson et al., 2023).

These limitations underscore the need for more robust, scalable, and interpretable multimodal frameworks, which this study aims to address.

3. Proposed Multimodal Fusion Framework

3.1 System Overview

The proposed framework is designed to enable robust and scalable integration of heterogeneous healthcare data by combining medical imaging and electronic health records (EHRs) within a unified deep learning architecture. The motivation for this design stems from the recognition that different data modalities encode complementary clinical information. Medical images provide high-resolution structural and morphological insights into disease pathology, whereas EHRs capture longitudinal patient trajectories, including physiological measurements, treatment histories, and clinician observations.

Traditional unimodal models fail to exploit these complementary signals, often resulting in suboptimal predictive performance. To address this limitation, the proposed system adopts a multimodal learning paradigm in which modality-specific feature representations are learned independently and subsequently integrated through an advanced fusion mechanism. This design enables the model to capture both intra-modal patterns and inter-modal dependencies, leading to more comprehensive patient representations (Huang et al., 2020; Khader et al., 2023).

The architecture consists of four major components:

- ❖ Imaging feature extraction module
- ❖ EHR representation learning module
- ❖ Multimodal fusion module
- ❖ Task-specific prediction module

These components are designed to operate in a coordinated pipeline, with bidirectional information flow and optional feedback mechanisms for iterative refinement.

3.2 Architecture Design




Figure 1. Multimodal Fusion Architecture for Imaging and EHR Integration.

Figure 1. Proposed multimodal deep learning architecture for precision disease prediction integrating medical imaging and electronic health records. The framework consists of an image encoder (ResNet or Vision Transformer), an EHR encoder (transformer-based models such as BEHRT or Med-BERT), and an optional clinical text encoder. Multimodal fusion is performed using early, intermediate (cross-attention), or late fusion strategies. The architecture supports multi-task prediction, including diagnosis and risk estimation. **Interpretability is achieved using Grad-CAM for imaging and SHAP for EHR features.** The model is designed to operate under both centralized and federated learning settings with differential privacy mechanisms, and includes a missing modality handling component for real-world clinical robustness.

3.3 Imaging Representation Learning

The imaging encoder is responsible for extracting discriminative spatial features from high-dimensional medical images. In this framework, two categories of architectures are considered: convolutional neural networks (CNNs) and Vision Transformers (ViTs).

CNN-based models, such as ResNet, are widely used due to their ability to capture hierarchical spatial features through convolutional operations and pooling layers. These models are particularly effective in identifying local patterns such as lesions, nodules, or structural abnormalities (He et al., 2016). However, CNNs are inherently limited in modeling long-range dependencies due to their localized receptive fields.

To address this limitation, transformer-based models such as ViTs are incorporated. These models utilize self-attention mechanisms to capture global contextual relationships across the entire image, enabling improved performance in complex imaging tasks (Dosovitskiy et al., 2021). The output of the imaging encoder is a high-dimensional embedding vector:

$$F_{img} \in \mathbb{R}^{d_1}$$

This representation encodes both local and global visual features relevant to disease prediction.

3.4 EHR Representation Learning

The EHR encoder is designed to process heterogeneous clinical data, including structured variables (e.g., laboratory results, vital signs, demographic attributes) and unstructured clinical narratives. Due to the temporal nature of EHR data, the encoder must effectively capture sequential dependencies and contextual relationships.

Traditional approaches utilize recurrent neural networks (RNNs) and long short-term memory (LSTM) networks to model temporal dynamics. While effective, these models suffer from limitations in capturing long-term dependencies and handling irregular time intervals.

To overcome these challenges, the proposed framework incorporates transformer-based architectures such as BEHRT and Med-BERT. These models treat EHR data as sequences of medical events and apply self-attention mechanisms to learn contextualized representations (Li et al., 2020; Rasmy et al., 2021). This approach enables the model to capture complex relationships between clinical variables across time.

The resulting EHR feature representation is defined as:

$$F_{ehr} \in \mathbb{R}^{d_2}$$

These embedding captures temporal trends, co-morbidity patterns, and patient-specific clinical context.

3.5 Multimodal Fusion Mechanism

The multimodal fusion module serves as the central component of the framework, integrating information from the imaging and EHR encoders. Unlike traditional approaches that rely on simple feature concatenation, this framework employs an attention-based fusion strategy to model cross-modal interactions.

The fusion process is formally expressed as:

$$F_{fusion} = \alpha \cdot F_{img} + \beta \cdot F_{ehr}$$

where α and β are learnable weights that dynamically adjust the contribution of each modality.

To further enhance integration, cross-modal attention mechanisms are applied. In this setup, the feature representation of one modality is used to compute attention weights over the features of the other modality. This allows the model to selectively focus on the most relevant features across modalities, thereby improving predictive performance and interpretability (Baltrušaitis et al., 2019; Hayat et al., 2022).

Additionally, a gating mechanism can be introduced to regulate information flow between modalities, particularly in scenarios where one modality may be noisy or incomplete.

3.6 Fusion Strategy Variants

To comprehensively evaluate multimodal integration, the framework supports multiple fusion strategies:

- ❖ **Early Fusion:** In early fusion, raw features from imaging and EHR data are combined before being fed into the model. This approach allows the network to learn joint representations from the outset but is sensitive to differences in data scale and distribution.
- ❖ **Intermediate Fusion:** Intermediate fusion combines latent feature representations extracted independently from each modality. This approach provides a balance between flexibility and performance, enabling modality-specific feature learning before integration.
- ❖ **Late Fusion:** Late fusion aggregates predictions from modality-specific models. While this approach improves robustness and modularity, it limits deep interaction between modalities and may result in suboptimal performance (Baltrušaitis et al., 2019).

- ❖ **Hybrid Fusion (Proposed)** The proposed framework emphasizes a hybrid fusion strategy that combines intermediate fusion with attention mechanisms. This approach enables both deep interaction and adaptive weighting of modalities, making it particularly suitable for complex clinical prediction tasks.

3.7 Prediction and Optimization

The fused representation is passed through a fully connected prediction layer, followed by an activation function appropriate for the task (e.g., sigmoid for binary classification or softmax for multi-class classification). The prediction function is defined as:

$$\hat{y} = f(F_{fusion})$$

The model is trained using a supervised learning approach with cross-entropy loss for classification tasks. Optimization is performed using gradient-based methods such as Adam, which provide efficient convergence in high-dimensional parameter spaces.

3.8 Robustness and Practical Considerations

- ❖ **Handling Missing Modalities:** Real-world healthcare datasets often contain incomplete or missing modalities. To address this challenge, the framework incorporates modality dropout during training, allowing the model to learn robust representations even when one modality is absent.
- ❖ **Interpretability:** Interpretability is essential for clinical adoption. The proposed framework integrates attention visualization and feature attribution techniques to provide insights into the contribution of each modality to the final prediction. This enhances transparency and supports clinical decision-making (Tonekaboni et al., 2019).
- ❖ **Scalability and Deployment:** The modular design of the framework ensures scalability and compatibility with large-scale datasets such as MIMIC-IV and MIMIC-CXR. Efficient training strategies and distributed computing techniques enable deployment in real-world healthcare environments (Johnson et al., 2023).

The proposed multimodal fusion framework provides a comprehensive solution for integrating medical imaging and EHR data. By combining advanced representation learning techniques with attention-based fusion mechanisms, the framework enables improved predictive performance, robustness, and interpretability. This architecture forms the foundation for the experimental evaluation presented in the subsequent sections.

4. Data Sources and Preprocessing

4.1 Data Sources

The effectiveness of multimodal deep learning models is highly dependent on the availability of large-scale, well-aligned datasets that integrate both medical imaging and electronic health records (EHRs). In this study, widely recognized and publicly available benchmark datasets are considered to ensure reproducibility, scalability, and comparability with prior work.

4.1.1 Medical Imaging Data

Medical imaging data are sourced from large-scale radiological datasets such as **MIMIC-CXR** and **CheXpert**, which contain chest radiographs labeled for multiple thoracic conditions. MIMIC-CXR provides over 370,000 chest X-ray images associated with corresponding radiology reports and patient metadata, making it one of the most comprehensive datasets for multimodal clinical research (Johnson et al., 2019).

Similarly, CheXpert offers a large collection of labeled chest radiographs with uncertainty-aware annotations, which are particularly useful for training robust diagnostic models (Irvin et al., 2019).

These datasets are selected due to their:

- ❖ High dimensionality and diversity of imaging samples
- ❖ Availability of clinically validated labels
- ❖ Compatibility with EHR datasets for multimodal integration

4.1.2 Electronic Health Record Data

EHR data are obtained from **MIMIC-IV**, a large-scale critical care database containing de-identified health records for intensive care unit (ICU) patients (Johnson et al., 2023). The dataset includes structured variables such as laboratory measurements, vital signs, medication records, and demographic information, as well as time-stamped clinical events.

MIMIC-IV is particularly suitable for this study because:

- ❖ It provides longitudinal patient trajectories
- ❖ It is linked to imaging datasets such as MIMIC-CXR
- ❖ It supports a wide range of predictive tasks, including mortality prediction and disease progression

4.1.3 Multimodal Alignment

To construct a unified dataset, imaging and EHR records are aligned at the patient level using unique identifiers and temporal matching. Specifically:

- ❖ Each chest X-ray is linked to the closest EHR snapshot within a defined time window
- ❖ Temporal consistency is enforced to ensure clinical relevance
- ❖ Multiple observations per patient are handled using sequence aggregation

This alignment process is critical for enabling meaningful multimodal learning, as misaligned data can introduce noise and degrade model performance.

4.2 Data Preprocessing Pipeline

The preprocessing pipeline is designed to address the heterogeneity, noise, and missingness inherent in clinical data. It consists of modality-specific preprocessing steps followed by normalization and synchronization.

4.2.1 Imaging Data Preprocessing

Medical images undergo a series of preprocessing steps to standardize input and enhance model performance:

- ❖ **Resizing and normalization:** Images are resized to a fixed resolution (e.g., 224×224 pixels) and normalized to ensure consistent intensity distributions
- ❖ **Data augmentation:** Techniques such as rotation, flipping, and contrast adjustment are applied to improve generalization and reduce overfitting
- ❖ **Noise reduction:** Image filtering techniques are used to remove artifacts and enhance feature clarity

These steps ensure that the imaging encoder receives high-quality and standardized input representations.

4.2.2 EHR Data Preprocessing

EHR data preprocessing involves handling structured, temporal, and categorical variables:

- ❖ **Missing value imputation:** Missing clinical values are addressed using statistical imputation techniques such as mean imputation or forward filling
- ❖ **Temporal encoding:** Time-stamped events are encoded using positional embeddings to preserve temporal relationships
- ❖ **Categorical encoding:** Medical codes (e.g., ICD codes) are embedded into dense vector representations

- ❖ **Normalization:** Continuous variables such as laboratory values are normalized to reduce scale variability

For sequential modeling, patient records are organized into temporal sequences, enabling the encoder to capture longitudinal dependencies.

4.2.3 Clinical Text Processing

Clinical notes and radiology reports are processed using transformer-based language models such as ClinicalBERT or BioBERT. The preprocessing steps include:

- ❖ Tokenization and text cleaning
- ❖ Removal of irrelevant symbols and noise
- ❖ Conversion into contextual embeddings

This optional modality provides additional semantic context that enhances predictive performance in multimodal settings.

4.3 Data Integration and Synchronization

Given the asynchronous nature of multimodal clinical data, careful synchronization is required to ensure consistency across modalities. The following strategies are employed:

- ❖ **Time-window alignment:** Imaging and EHR data are matched within clinically meaningful time intervals
- ❖ **Sequence padding and truncation:** Variable-length sequences are standardized for batch processing
- ❖ **Feature alignment:** Embedding dimensions from different modalities are projected into a shared latent space

This step ensures that all modalities contribute coherently to the fusion process.

4.4 Handling Missing and Incomplete Data

Missing data is a common challenge in real-world healthcare datasets. To address this issue, the framework incorporates:

- ❖ **Masking strategies:** Missing inputs are represented using mask tokens
- ❖ **Modality dropout:** The model is trained to handle absent modalities by randomly dropping inputs during training
- ❖ **Contrastive alignment loss:** Encourages consistency between available modalities

These techniques improve robustness and enable the model to generalize effectively in incomplete data scenarios.

4.5 Dataset Summary

Table 1. Summary of Datasets Used

Dataset	Modality	Size (Approx.)	Key Features
MIMIC-CXR	Imaging	370K+ images	Chest X-rays, radiology reports
CheXpert	Imaging	220K+ images	Labeled thoracic diseases
MIMIC-IV	EHR	70K+ patients	Longitudinal clinical records

4.6 Ethical Considerations and Data Privacy

All datasets used in this study are de-identified and publicly available for research purposes. Data handling complies with established ethical standards and privacy regulations. Furthermore, the proposed framework incorporates privacy-preserving mechanisms such as differential privacy and federated learning, ensuring secure handling of sensitive healthcare data (Dwork et al., 2014; McMahan et al., 2017).

This section presented the data sources and preprocessing pipeline for multimodal learning. By leveraging large-scale benchmark datasets and implementing robust preprocessing strategies, the framework ensures high-quality input representations for both imaging and EHR modalities. These steps form the foundation for the model design and training strategy described in the next section.

5. Model Design and Training Strategy

The proposed multimodal system is composed of three primary components: an imaging encoder, an electronic health record (EHR) encoder, and a multimodal fusion module, followed by task-specific prediction heads. Each component is designed to capture modality-specific characteristics while enabling effective cross-modal interaction within a unified learning framework. The imaging encoder utilizes either a convolutional neural network such as ResNet or a Vision Transformer depending on the experimental configuration. CNN-based models are effective in extracting hierarchical spatial features from medical images through convolutional operations, while transformer-based architectures provide enhanced capability for modeling global contextual dependencies across image regions (He et al., 2016; Dosovitskiy et al., 2021). In parallel, the EHR encoder is implemented using transformer-based architectures such as BEHRT or Med-BERT, which model longitudinal patient data as sequences of clinical events. These models employ self-attention mechanisms to capture temporal dependencies and contextual relationships within EHR data, thereby enabling the extraction of rich patient representations (Li et al., 2020; Rasmy et al., 2021). The outputs from both encoders are projected into a shared latent space and subsequently integrated through the multimodal fusion module described in Section 3, producing a fused representation that can be written as

$$F_{fusion} \in \mathbb{R}^d.$$

The framework is designed to support multi-task learning, enabling simultaneous prediction of multiple clinically relevant outcomes. In particular, the model performs multi-label disease classification, where it predicts the presence of multiple disease conditions for a given patient, as well as risk prediction tasks such as in-hospital mortality or disease progression. For multi-label disease classification, the predicted output for patient (i) can be expressed as

$$\hat{y}_i \in [0, 1]^K$$

where K denotes the number of disease classes and each element of $\hat{y}_i$ corresponds to the predicted probability of a specific disease label. For risk prediction tasks, the model generates either a binary probability or a continuous risk estimate depending on the nature of the endpoint. This joint formulation allows the model to leverage shared multimodal representations across tasks while maintaining task-specific outputs.

Training of the model is conducted using a supervised learning strategy with objective functions tailored to each prediction task. For the multi-label disease classification problem, binary cross-entropy loss is employed because it is well suited for independent class predictions in clinical diagnosis settings. The classification loss is defined as

$$\mathcal{L}_{BCE} = - \sum_{i=1}^K [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

where y_i denotes the ground-truth label and $\hat{y}_i$ denotes the predicted probability for class i . When the framework is optimized under a multi-task setting, the overall loss is formulated as a weighted combination of classification and risk prediction objectives:

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_{classification} + \lambda_2 \mathcal{L}_{risk}$$

where λ_1 and λ_2 are task-specific weighting coefficients introduced to balance the relative contribution of each objective during optimization. This formulation enables the network to learn shared multimodal features while preserving sensitivity to task-specific clinical endpoints.

Optimization is performed using the Adam optimizer, which is particularly suitable for high-dimensional deep learning problems because of its adaptive learning-rate mechanism and efficient convergence behavior. To further improve stability during training, a learning-rate scheduling strategy is applied so that the model can gradually refine its parameters over successive epochs. In addition, regularization strategies are incorporated to improve generalization and reduce overfitting. Dropout is applied within fully connected layers to reduce co-adaptation among neurons, weight decay is introduced to constrain model complexity, and data augmentation is applied to imaging inputs to improve robustness to minor visual variations. Early stopping based on validation performance is also used to prevent over-training and ensure appropriate model selection. Because medical datasets often exhibit severe class imbalance, particularly for rare conditions and adverse outcomes, the training strategy also includes mechanisms for imbalance correction. Class-weighted loss terms are used to assign greater importance to minority classes, while oversampling strategies are applied where necessary to improve representation of under-observed outcomes. In addition, threshold tuning is performed during validation so that classification sensitivity can be improved for clinically important but infrequent disease labels.

Model evaluation is conducted using a comprehensive set of performance metrics that capture both predictive accuracy and clinical relevance. Standard metrics such as accuracy, precision, recall, and F1-score are used to assess classification performance. The area under the receiver operating characteristic curve is employed as a principal measure of discrimination, particularly because it is less sensitive to class imbalance than raw accuracy. To ensure statistical rigor, differences in AUC between competing models are evaluated using the DeLong test, which provides a robust nonparametric approach for comparing correlated ROC curves (DeLong et al., 1988). Beyond discrimination, calibration is also assessed to ensure that predicted probabilities correspond closely to true outcome likelihoods. Calibration performance is evaluated using reliability diagrams and expected calibration error, which are widely accepted tools for assessing confidence reliability in modern neural networks (Guo et al., 2017).

The experimental setup is designed to ensure fair, reproducible, and clinically meaningful evaluation. The dataset is partitioned into training, validation, and test sets using patient-level splitting so that no patient appears across multiple subsets, thereby preventing information leakage. All models, including unimodal baselines and the proposed multimodal system, are trained under identical conditions to enable valid comparative analysis. Specifically, three principal configurations are evaluated: an imaging-only model, an EHR-only model, and the multimodal fusion model. This comparative design makes it possible to isolate the contribution of each modality and quantify the added value of multimodal integration.

To evaluate robustness under real-world conditions, the model is also tested in scenarios involving missing modalities, which are common in routine clinical practice. Modality dropout is applied during training so that the model learns to operate effectively even when one input source is absent. Cross-validation is further used to assess the stability of model performance across multiple data splits. The framework is implemented using modern deep learning libraries such as PyTorch or TensorFlow, and training is conducted on GPU-enabled systems to accommodate the computational demands of multimodal learning. Hyperparameters, including batch size, learning rate, dropout rate, and regularization coefficients, are tuned using validation data to achieve optimal performance.

Overall, the model design and training strategy provide a robust and scalable foundation for multimodal disease prediction. By integrating advanced encoder architectures, explicit mathematical objectives, carefully designed optimization strategies, and comprehensive evaluation protocols, the proposed framework ensures reliable performance and establishes the methodological basis for the experimental analysis presented in the following section.

6. Experimental Results and Analysis

The experimental evaluation aims to assess the effectiveness of the proposed multimodal framework in comparison to unimodal baselines and alternative fusion strategies. The results are presented in terms of predictive performance, scalability, and robustness, with a particular focus on clinically relevant metrics such as AUC, F1-score, and calibration.

The comparative analysis includes three model configurations: an imaging-only model, an EHR-only model, and the proposed multimodal fusion model. All models are trained and evaluated under identical conditions to ensure a fair comparison. The results demonstrate that the multimodal model consistently outperforms unimodal approaches across all evaluation metrics, confirming the advantage of integrating complementary clinical modalities.

6.1 Performance Comparison

Table 2. Performance Comparison Across Models

Model Type	Accuracy (%)	AUC	F1 Score
Imaging Only	82.5	0.85	0.81
EHR Only	78.3	0.80	0.77
Multimodal Fusion	90.2	0.93	0.89

The results indicate a substantial improvement in predictive performance when both imaging and EHR data are integrated. The multimodal model achieves an accuracy of 90.2% and an AUC of 0.93, significantly outperforming both the imaging-only and EHR-only models. This improvement can be attributed to the complementary nature of the modalities, where imaging captures structural abnormalities and EHR data provides temporal and contextual clinical information. The increase in F1-score further highlights the model's ability to balance precision and recall, which is particularly important in medical diagnosis tasks.

6.2 Visualization of Model Performance




Figure 2. Accuracy comparison between imaging-only, EHR-only, and multimodal models.

The graphical representation further emphasizes the superiority of the multimodal approach, clearly illustrating the performance gap between unimodal and multimodal models. The consistent improvement across metrics reinforces the importance of multimodal integration for precision disease prediction.

6.3 Scalability and Data Efficiency

To evaluate scalability, model performance is analyzed as a function of dataset size. The results show that while unimodal models experience diminishing returns as data volume increases, the multimodal model continues to benefit from additional data due to its ability to leverage cross-modal relationships.




Figure 3. Model performance (AUC) as a function of dataset size.

This result highlights the scalability of the proposed framework and its suitability for large-scale healthcare applications.

6.4 Ablation Study

To understand the impact of different fusion strategies, an ablation study is conducted comparing early, late, and hybrid fusion approaches.

Table 3. Ablation Study of Fusion Strategies

Fusion Strategy	Accuracy (%)
Early Fusion	85.4
Late Fusion	87.1
Hybrid Fusion	90.2

The results indicate that hybrid fusion achieves the highest performance, confirming that combining intermediate representations with attention mechanisms enables more effective cross-modal interaction. Early fusion suffers from feature incompatibility issues, while late fusion limits deep integration between modalities. The hybrid approach overcomes these limitations by enabling both flexibility and interaction.

6.5 Robustness to Missing Modalities

To simulate real-world clinical scenarios, the model is evaluated under conditions where one modality is partially or completely missing. The multimodal framework demonstrates strong robustness due to the incorporation of modality dropout and adaptive fusion mechanisms. While performance decreases slightly in the absence of one modality, the model still outperforms unimodal baselines, indicating its ability to generalize effectively in incomplete data environments.

6.6 Calibration and Reliability

Calibration analysis shows that the multimodal model produces well-calibrated predictions, with predicted probabilities closely matching observed outcomes. Reliability diagrams and expected calibration error (ECE) confirm that the model maintains consistent confidence levels across different prediction thresholds. This is critical for clinical applications, where overconfident or poorly calibrated predictions can lead to incorrect decision-making (Guo et al., 2017).

6.7 Discussion of Results

The experimental findings demonstrate that multimodal fusion significantly enhances predictive performance compared to unimodal approaches. The integration of imaging and EHR data enables the model to capture both structural and temporal aspects of disease, leading to more accurate and robust predictions. The superiority of the hybrid fusion strategy further highlights the importance of modeling cross-modal interactions rather than relying on simple aggregation techniques.

In addition, the scalability and robustness analyses confirm that the proposed framework is well-suited for real-world clinical deployment. The ability to handle missing modalities and maintain performance across varying dataset sizes addresses key challenges in healthcare data integration. Overall, the results provide strong empirical support for the effectiveness of the proposed multimodal framework.

7. Discussion

The findings of this study provide strong evidence that multimodal deep learning significantly enhances predictive performance in clinical decision-making tasks compared to unimodal approaches. By integrating medical imaging and electronic health records (EHRs), the proposed framework captures both structural and

temporal dimensions of disease, resulting in a more comprehensive representation of patient health. The experimental results presented in Section 6 demonstrate consistent improvements across accuracy, AUC, and F1-score, confirming that multimodal fusion enables the model to leverage complementary information that is otherwise inaccessible in single-modality systems.

A key observation from the results is the substantial performance gap between unimodal and multimodal models. While imaging-only models are effective in detecting visual abnormalities and EHR-only models capture longitudinal clinical patterns, each modality alone provides an incomplete view of the patient. The multimodal framework bridges this gap by combining spatial and temporal features, leading to improved discrimination and predictive reliability. This finding aligns with prior research in multimodal machine learning, which emphasizes the importance of cross-modal representation learning for complex real-world tasks (Baltrušaitis et al., 2019). However, the present study extends this understanding by demonstrating that attention-based hybrid fusion strategies can further enhance performance beyond traditional fusion techniques. The ablation study highlights the importance of the fusion mechanism itself, revealing that hybrid fusion outperforms both early and late fusion strategies. Early fusion is limited by differences in data distributions and feature scales across modalities, which can hinder effective learning. Late fusion, while more robust, restricts interaction between modalities to the decision level, thereby limiting the model's ability to learn joint representations. In contrast, the hybrid fusion approach enables deep interaction between modalities through intermediate feature integration and attention mechanisms. This allows the model to dynamically prioritize clinically relevant features across modalities, improving both accuracy and interpretability. These findings are consistent with recent advancements in attention-based multimodal learning, where cross-modal attention has been shown to enhance feature alignment and information exchange (Hayat et al., 2022).

Another important contribution of this study is the demonstration of scalability and data efficiency in multimodal systems. The results indicate that while unimodal models exhibit diminishing returns as dataset size increases, the multimodal model continues to benefit from additional data. This suggests that multimodal learning is particularly well-suited for large-scale healthcare applications, where diverse data sources are increasingly available. The ability to effectively utilize additional data is critical in modern clinical environments, where the volume and variety of patient data continue to grow rapidly.

Robustness to missing modalities is another critical aspect of real-world deployment. In clinical practice, it is common for patient data to be incomplete due to missing tests, unavailable imaging, or irregular data collection. The proposed framework addresses this challenge through modality dropout and adaptive fusion mechanisms, enabling the model to maintain strong performance even when one modality is partially or completely absent. This capability represents a significant advantage over traditional models, which often fail when required inputs are missing. The incorporation of a missing modality handler, as described in Section 3, ensures that the framework remains applicable in realistic healthcare scenarios.

Interpretability is a key requirement for the adoption of artificial intelligence systems in clinical settings. The integration of Grad-CAM and SHAP within the framework provides modality-specific explanations that enhance transparency and trust. Grad-CAM enables visualization of salient regions in medical images that contribute to predictions, while SHAP identifies the most influential clinical variables in EHR data (Selvaraju et al., 2017; Lundberg & Lee, 2017). Together, these techniques provide complementary insights that can assist clinicians in understanding and validating model outputs. This aligns with the growing emphasis on explainable AI in healthcare, where interpretability is essential for regulatory approval and clinical acceptance (Tonekaboni et al., 2019).

From a deployment perspective, the framework's compatibility with federated learning and differential privacy mechanisms further strengthens its practical relevance. Healthcare data are highly sensitive and often distributed across institutions, making centralized data aggregation challenging. By supporting federated learning, the framework enables collaborative model training without requiring data sharing, thereby preserving patient privacy (McMahan et al., 2017). The integration of differential privacy techniques further enhances security by ensuring that individual patient information cannot be inferred from model outputs (Dwork et al., 2014). These features position the proposed framework as a viable solution for real-world clinical deployment in privacy-sensitive environments.

Despite these strengths, it is important to contextualize the findings within the broader landscape of multimodal healthcare research. While the results demonstrate clear performance gains, the effectiveness of multimodal models is inherently dependent on data quality, alignment, and availability. In particular, the process of aligning imaging and EHR data remains a non-trivial challenge, as temporal mismatches and incomplete records can introduce noise. Future work should explore more sophisticated alignment techniques and self-supervised learning approaches to further enhance robustness.

Overall, this study contributes to the growing body of literature on multimodal deep learning in healthcare by providing a comprehensive framework that integrates advanced representation learning, attention-based fusion, and practical deployment considerations. The results not only validate the effectiveness of multimodal integration but also highlight key design principles for building robust and scalable clinical AI systems. These findings have important implications for the development of next-generation decision support tools that can improve diagnostic accuracy, optimize treatment strategies, and ultimately enhance patient outcomes.

8. Limitations of the Study

Despite the promising results demonstrated by the proposed multimodal framework, several limitations should be acknowledged to provide a balanced and transparent assessment of the study. First, the performance of the model is inherently dependent on the quality and completeness of the underlying data. Although large-scale datasets such as MIMIC-IV and MIMIC-CXR provide a robust foundation, these datasets are derived from specific clinical environments and may not fully represent the diversity of patient populations, healthcare systems, or imaging protocols encountered in broader clinical practice. As a result, the generalizability of the model to different institutions, geographic regions, or patient demographics may be limited.

Second, the process of aligning multimodal data introduces potential sources of noise and uncertainty. The temporal synchronization of imaging and EHR data relies on predefined time windows, which may not always capture the true clinical relationship between modalities. Misalignment can lead to the incorporation of irrelevant or outdated clinical information, potentially affecting model performance. While this study employs standard alignment techniques, more advanced methods for dynamic temporal alignment could further improve the reliability of multimodal integration.

Another limitation relates to the interpretability of the model. Although techniques such as Grad-CAM and SHAP provide useful insights into model behavior, they offer only partial explanations and may not fully capture the complex interactions within deep neural networks. The explanations generated by these methods are often approximations and should be interpreted with caution, particularly in high-stakes clinical decision-making contexts. Future work should explore more rigorous and clinically validated interpretability approaches to enhance trust and usability.

The computational complexity of multimodal deep learning models also presents practical challenges. The integration of multiple encoders, attention mechanisms, and fusion modules increases the computational cost of both training and inference. This may limit the feasibility of deploying such models in resource-constrained environments, particularly in low- and middle-income healthcare settings. Although the use of efficient architectures and hardware acceleration can mitigate some of these challenges, further optimization is required to enable widespread adoption.

In addition, the evaluation conducted in this study is primarily retrospective, relying on existing datasets rather than prospective clinical trials. While retrospective evaluation provides valuable insights into model performance, it does not fully capture the complexities of real-world clinical workflows, including human–AI interaction, clinician trust, and decision-making dynamics. Prospective validation in clinical settings is necessary to establish the true impact of the proposed framework on patient outcomes and healthcare delivery. Finally, while the framework incorporates mechanisms for handling missing modalities, the effectiveness of these strategies may vary depending on the extent and pattern of missing data. In cases where entire modalities are systematically absent or highly incomplete, the performance of the model may degrade more significantly than observed in controlled experimental settings. Further research is needed to develop more robust approaches for dealing with extreme data sparsity and modality imbalance.

9. Conclusion

This study presented a comprehensive multimodal deep learning framework for precision disease prediction through the integration of medical imaging and electronic health records. By combining advanced representation learning techniques with attention-based fusion mechanisms, the proposed approach enables the extraction and integration of complementary information from heterogeneous clinical data sources. The results demonstrate that multimodal fusion significantly outperforms unimodal models across multiple evaluation metrics, including accuracy, AUC, and F1-score, highlighting the value of integrating diverse data modalities in clinical prediction tasks.

A key contribution of this work lies in the design of a flexible and scalable architecture that supports multiple fusion strategies, including early, intermediate, and hybrid approaches. The experimental findings indicate that hybrid fusion, which combines intermediate feature integration with attention mechanisms, provides the most effective balance between flexibility and performance. This insight contributes to the growing body of knowledge on multimodal learning by identifying practical design principles for optimizing cross-modal interaction in healthcare applications.

The study also addresses several critical challenges associated with real-world deployment, including missing data, interpretability, and data privacy. The incorporation of modality dropout and adaptive fusion strategies enhances robustness to incomplete inputs, while the integration of interpretability techniques such as Grad-CAM and SHAP improves transparency and clinical trust. Furthermore, the compatibility of the framework with federated learning and differential privacy mechanisms enables secure and privacy-preserving model training across distributed healthcare systems.

Beyond its technical contributions, this work has important implications for the future of clinical decision support systems. The ability to integrate and analyze multimodal patient data opens new opportunities for improving diagnostic accuracy, predicting disease progression, and personalizing treatment strategies. As healthcare systems continue to generate increasingly complex and diverse data, multimodal deep learning

frameworks such as the one proposed in this study will play a critical role in transforming raw data into actionable clinical insights.

Future research should focus on extending this framework to additional data modalities, such as genomic data, wearable sensor data, and real-time monitoring systems, to further enhance predictive capabilities. In addition, efforts should be directed toward improving model efficiency, interpretability, and generalizability, as well as conducting prospective clinical studies to evaluate real-world impact. By addressing these challenges, multimodal artificial intelligence has the potential to become a cornerstone of next-generation healthcare systems, enabling more precise, data-driven, and patient-centered care.

Reference

1. Huang, S. C., Pareek, A., Seyyedi, S., Banerjee, I., & Lungren, M. P. (2020). Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. *NPJ digital medicine*, 3(1), 136.
2. Khader, F., Müller-Franzes, G., Wang, T., Han, T., Tayebi Arasteh, S., Haarbuerger, C., ... & Truhn, D. (2023). Multimodal deep learning for integrating chest radiographs and clinical parameters: a case for transformers. *Radiology*, 309(1), e230806.
3. Jabbour, S., Fouhey, D., Kazerooni, E., Wiens, J., & Sjoding, M. W. (2022). Combining chest X-rays and electronic health record (EHR) data using machine learning to diagnose acute respiratory failure. *Journal of the American Medical Informatics Association*, 29(6), 1060-1068.
4. Lee, C. K., Chen, T. L., Wu, J. E., Liao, M. T., Wang, C., Wang, W., & Chou, C. Y. (2024). Multimodal deep learning models utilizing chest X-ray and electronic health record data for predictive screening of acute heart failure in emergency department. *Computer Methods and Programs in Biomedicine*, 255, 108357.
5. Li, D., Xing, W., Zhao, J., Shi, C., & Wang, F. (2025). Multimodal deep learning for predicting in-hospital mortality in heart failure patients using longitudinal chest X-rays and electronic health records. *The International Journal of Cardiovascular Imaging*, 41(3), 427-440.
6. Hayat, N., Geras, K. J., & Shamout, F. E. (2022, December). MedFuse: Multi-modal fusion with clinical time-series data and chest X-ray images. In *Machine Learning for Healthcare Conference* (pp. 479-503). PMLR.
7. Huang, S. C., Pareek, A., Zamanian, R., Banerjee, I., & Lungren, M. P. (2020). Multimodal fusion with deep neural networks for leveraging CT imaging and electronic health record: a case-study in pulmonary embolism detection. *Scientific reports*, 10(1), 22147.
8. Pölsterl, S., Wolf, T. N., & Wachinger, C. (2021, September). Combining 3D image and tabular data via the dynamic affine feature map transform. In *International conference on medical image computing and computer-assisted intervention* (pp. 688-698). Cham: Springer International Publishing.
9. Parisot, S., Ktena, S. I., Ferrante, E., Lee, M., Guerrero, R., Glocker, B., & Rueckert, D. (2018). Disease prediction using graph convolutional networks: application to autism spectrum disorder and Alzheimer's disease. *Medical image analysis*, 48, 117-130.
10. Liu, J., Capurro, D., Nguyen, A., & Verspoor, K. (2023). Attention-based multimodal fusion with contrast for robust clinical prediction in the face of missing modalities. *Journal of biomedical informatics*, 145, 104466.

11. Baltrusaitis, T., Ahuja, C., & Morency, L. P. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2), 423-443.
12. Johnson, A. E., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A., Horng, S., ... & Mark, R. G. (2023). MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, 10(1), 1.
13. Johnson, A. E., Pollard, T. J., Berkowitz, S. J., Greenbaum, N. R., Lungren, M. P., Deng, C. Y., ... & Horng, S. (2019). MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1), 317.
14. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., ... & Ng, A. Y. (2019, July). Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, No. 01, pp. 590-597).
15. Li, Y., Rao, S., Solares, J. R. A., Hassaine, A., Ramakrishnan, R., Canoy, D., ... & Salimi-Khorshidi, G. (2020). BEHRT: transformer for electronic health records. *Scientific reports*, 10(1), 7155.
16. Rasmy, L., Xiang, Y., Xie, Z., Tao, C., & Zhi, D. (2021). Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1), 86.
17. Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., & Stewart, W. (2016). Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, 29.
18. Alsentzer, E., Murphy, J., Boag, W., Weng, W. H., Jindi, D., Naumann, T., & McDermott, M. (2019, June). Publicly available clinical BERT embeddings. In *Proceedings of the 2nd clinical natural language processing workshop* (pp. 72-78).
19. Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., & Kang, J. (2020). BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4), 1234-1240.
20. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770-778).
21. Ronneberger, O., Fischer, P., & Brox, T. (2015, October). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (pp. 234-241). Cham: Springer international publishing.
22. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
23. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision* (pp. 618-626).
24. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
25. DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, 837-845.
26. Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017, July). On calibration of modern neural networks. In *International conference on machine learning* (pp. 1321-1330). PMLR.

27. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016, October). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security (pp. 308-318).
28. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017, April). Communication-efficient learning of deep networks from decentralized data. In Artificial intelligence and statistics (pp. 1273-1282). Pmlr.
29. Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., ... & Goldenberg, A. (2019). Do no harm: a roadmap for responsible machine learning for health care. *Nature medicine*, 25(9), 1337-1340.
30. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ digital medicine*, 3(1), 119.