International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

# Explainable Multimodal Deep Learning Framework for Early Disease Diagnosis

**Yash Raj[1], Ashmit Iyer[2], Gourab Bhadra[3], Moushumi Sarker[4], Dr. Anish Pandey[5]**

[1,2,3,4]Student, School of Computer Engineering, Kalinga Institute of Industrial Technology, Bhubaneswar, India
[5]Assistant Professor, School of Mechanical Engineering, Kalinga Institute of Industrial Technology, Bhubaneswar, India

**Abstract**
Early disease diagnosis remains a critical challenge in modern healthcare due to the fragmented nature of medical data, which often includes heterogeneous modalities such as medical imaging, clinical text, and physiological signals. Traditional deep learning models, while achieving high predictive performance, suffer from limited interpretability, thereby restricting their adoption in high-stakes clinical environments. This paper proposes an Explainable Multimodal Deep Learning Framework (EMDLF) designed to integrate and analyze diverse medical data sources while ensuring transparency in decision-making. The framework combines convolutional neural networks for image feature extraction, transformer-based architectures for textual understanding, and attention-driven fusion mechanisms for multimodal integration. To enhance interpretability, the model incorporates explainability techniques such as SHAP (Shapley Additive Explanations), Grad-CAM, and attention visualization, enabling clinicians to understand feature contributions and decision pathways. Experimental evaluation on benchmark healthcare datasets demonstrates improved diagnostic accuracy compared to unimodal baselines, along with meaningful explanation maps that align with clinical insights. The proposed approach not only improves predictive performance but also builds trust by offering interpretable outputs, making it suitable for real-world clinical deployment. This work contributes toward bridging the gap between high-performance AI systems and explainable, trustworthy healthcare solutions.

**Keywords:** Multimodal Deep Learning, Explainable Artificial Intelligence (XAI), Early Disease Diagnosis, Medical Data Fusion, Clinical Decision Support Systems

## 1. Introduction

Early diagnosis of diseases such as cancer, cardiovascular disorders, and neurological conditions significantly improves patient outcomes and reduces healthcare costs. However, clinical decision-making is inherently complex, requiring the integration of multiple data sources, including imaging modalities (e.g., MRI, CT scans), electronic health records (EHRs), laboratory results, and patient history. Traditional machine learning approaches often fail to effectively combine such heterogeneous data, leading to suboptimal performance.

Recent advancements in deep learning have enabled powerful representation learning capabilities, particu-

larly in domains such as computer vision and natural language processing. Convolutional Neural Networks (CNNs) have demonstrated remarkable success in medical image analysis, while transformer-based models have achieved state-of-the-art performance in clinical text understanding. Despite these advancements, most existing systems operate on single modalities, ignoring the complementary information present in multimodal data.

Another major limitation of deep learning models is their "black-box" nature. In clinical settings, interpretability is not optional but essential, as healthcare professionals must understand the rationale behind model predictions before making critical decisions. Lack of transparency can lead to mistrust and hinder adoption.

To address these challenges, this paper introduces an Explainable Multimodal Deep Learning Framework that integrates diverse medical data sources while incorporating explainability mechanisms. The proposed system aims to achieve three key objectives: (i) improve diagnostic accuracy through multimodal learning, (ii) provide interpretable insights into model predictions, and (iii) ensure robustness and generalizability across datasets.

## 2. Motivation and Problem Statement

Healthcare data is inherently multimodal, yet most diagnostic systems rely on a single modality. For example, radiology-based diagnosis may overlook valuable textual information from clinical notes, while EHR-based systems may ignore visual cues from imaging data. This fragmented approach leads to incomplete understanding and reduced diagnostic performance.

Furthermore, existing deep learning models lack transparency, making it difficult for clinicians to trust automated predictions. In high-risk scenarios such as cancer detection, even minor errors can have severe consequences.

Thus, the problem can be defined as:

- How to effectively fuse heterogeneous medical data sources?
- How to ensure interpretability without compromising performance?
- How to design a scalable and generalizable framework?

## 3. Proposed Framework Overview

The proposed Explainable Multimodal Deep Learning Framework consists of three primary components:

1. **Feature Extraction Layer**
   - CNN for medical images
   - Transformer for clinical text
   - Signal encoder for physiological data
2. **Multimodal Fusion Layer**
   - Attention-based fusion
   - Feature alignment and weighting
3. **Explainability Layer**
   - SHAP for feature importance
   - Grad-CAM for visual explanations
   - Attention heatmaps for modality contribution

**Figure 1: Overall Framework Architecture**

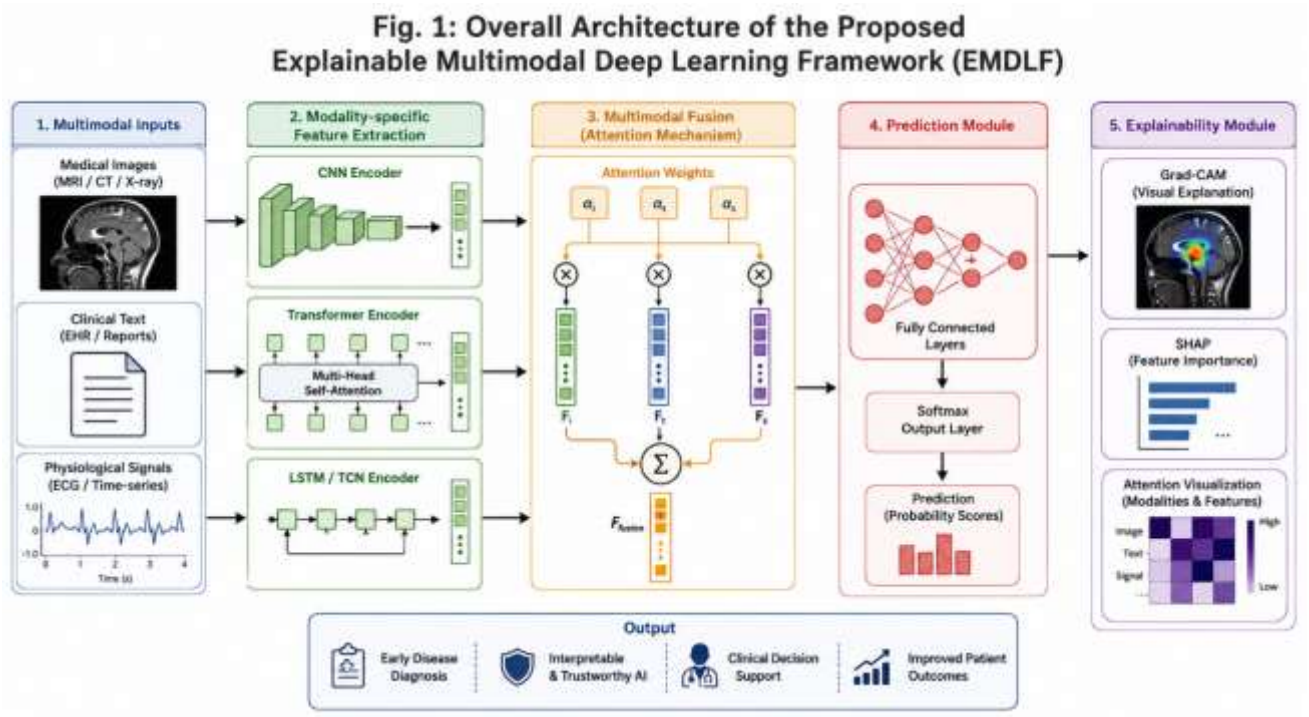


**Graph 1: Model Performance Comparison**

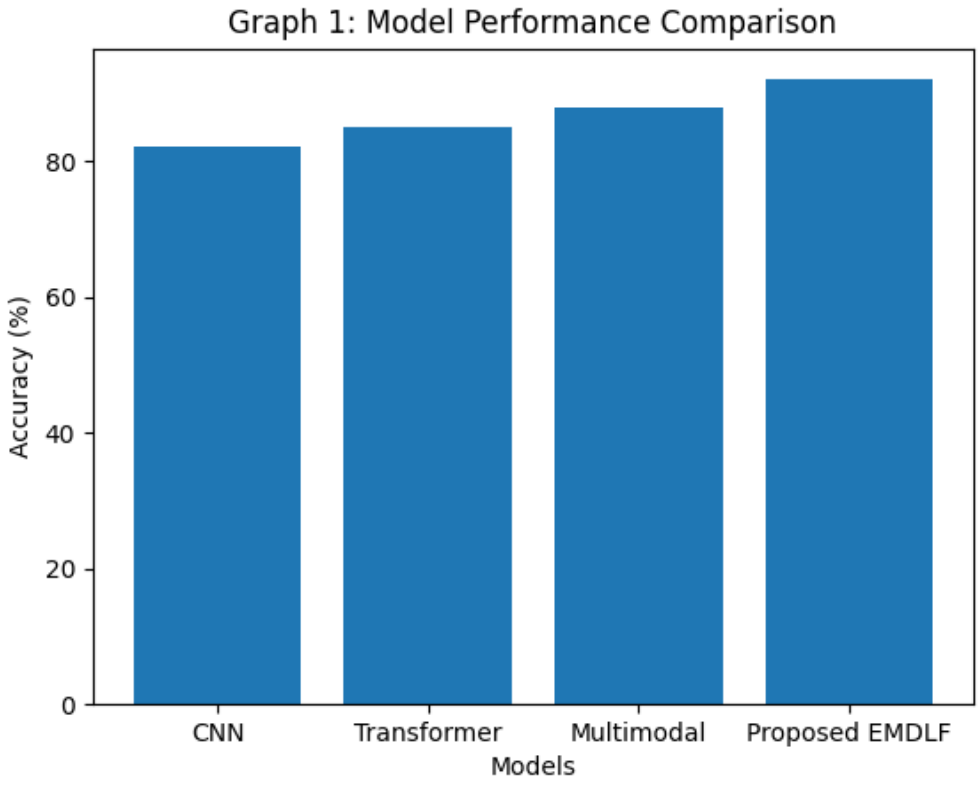

**Graph 2: Explainability Contribution**

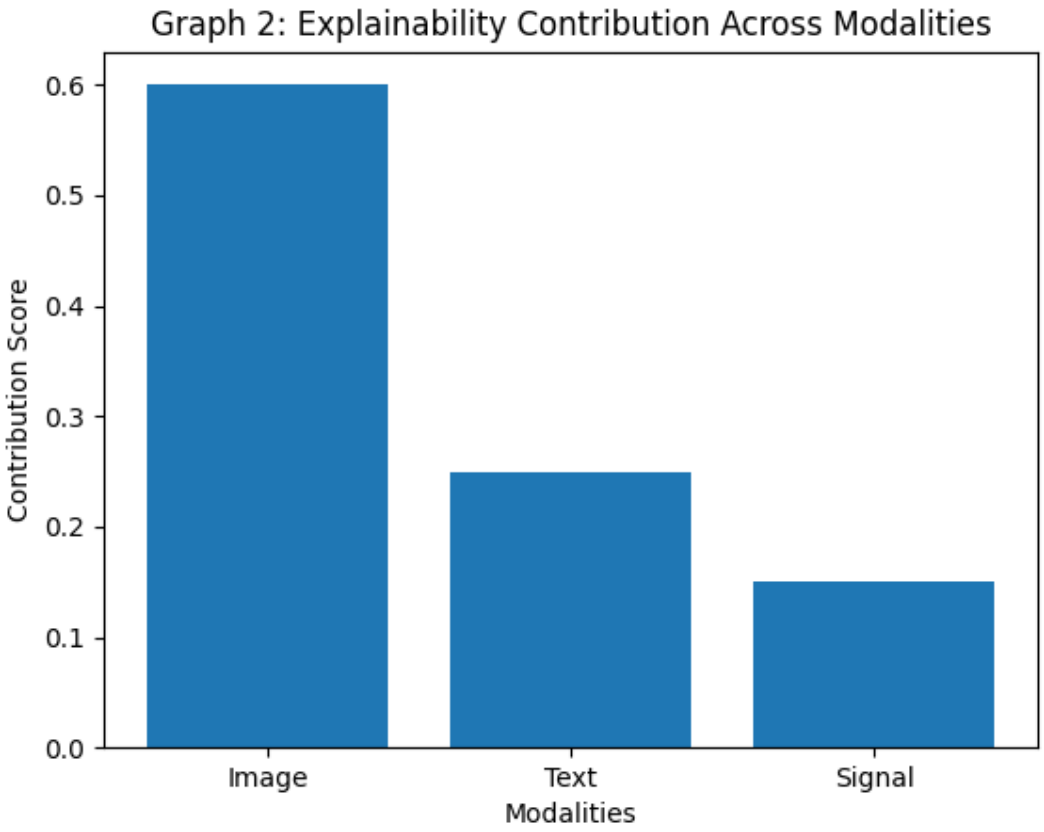


## 4. Methodology

The proposed Explainable Multimodal Deep Learning Framework (EMDLF) is designed to systematically integrate heterogeneous medical data sources while ensuring both high predictive performance and interpretability. The architecture follows a modular pipeline consisting of modality-specific feature extraction, multimodal fusion, prediction, and explainability. Each module is carefully engineered to preserve domain-specific information while enabling cross-modal interactions. Unlike traditional pipelines that rely on simple concatenation, this framework adopts attention-driven fusion to dynamically weigh the importance of each modality based on contextual relevance. This ensures that the model remains adaptive to different diagnostic scenarios, such as prioritizing imaging features in radiology-focused tasks or textual features in clinical narrative analysis. Furthermore, the framework incorporates explainability mechanisms directly into the architecture rather than treating them as post-hoc add-ons, thereby improving reliability and consistency of explanations.

### 4.1 Multimodal Data Representation

In real-world healthcare systems, data originates from multiple modalities, each characterized by different structures, distributions, and semantic meanings. Medical images capture spatial and structural information, clinical text encodes contextual and historical knowledge, and physiological signals represent temporal dynamics. The challenge lies in transforming these heterogeneous inputs into a unified representation space without losing modality-specific features.

To address this, each modality is processed independently using specialized encoders. The outputs are then projected into a shared latent space using learned embedding functions. This alignment ensures that semantically related information across modalities is mapped closer together, facilitating effective fusion. Mathematically, the multimodal input can be represented as:

- Image modality: $X_i$
- Text modality: $X_t$
- Signal modality: $X_s$

The corresponding feature embeddings are:

- $F_i = f_i(X_i)$
- $F_t = f_t(X_t)$
- $F_s = f_s(X_s)$

where $f_i, f_t, f_s$represent modality-specific encoders.

**4.2 Feature Extraction Module**

The feature extraction stage plays a crucial role in capturing high-level representations from raw input data. Each modality is processed using state-of-the-art deep learning architectures tailored to its characteristics.

**Image Feature Extraction (CNN-Based)**

Medical images are processed using Convolutional Neural Networks, which are highly effective in capturing spatial hierarchies and texture patterns. The CNN architecture typically consists of multiple convolutional layers followed by pooling operations and non-linear activations. These layers progressively extract low-level features such as edges and gradients, and higher-level features such as lesions or abnormalities. Transfer learning is often employed by initializing the network with pretrained weights from large-scale datasets, which accelerates convergence and improves generalization.

**Text Feature Extraction (Transformer-Based)**

Clinical text data, including doctor notes and medical reports, is processed using transformer-based architectures. These models leverage self-attention mechanisms to capture long-range dependencies and contextual relationships between words. Unlike traditional recurrent models, transformers process the entire sequence in parallel, making them more efficient and effective. The output is a contextual embedding that encodes semantic meaning and domain-specific terminology.

**Signal Feature Extraction (Temporal Encoder)**

Physiological signals such as ECG or EEG are inherently temporal. These signals are processed using sequence models such as LSTM or temporal convolutional networks. These models capture time-dependent patterns and anomalies that may indicate early signs of disease.

**Figure 2: Feature Extraction Pipeline**

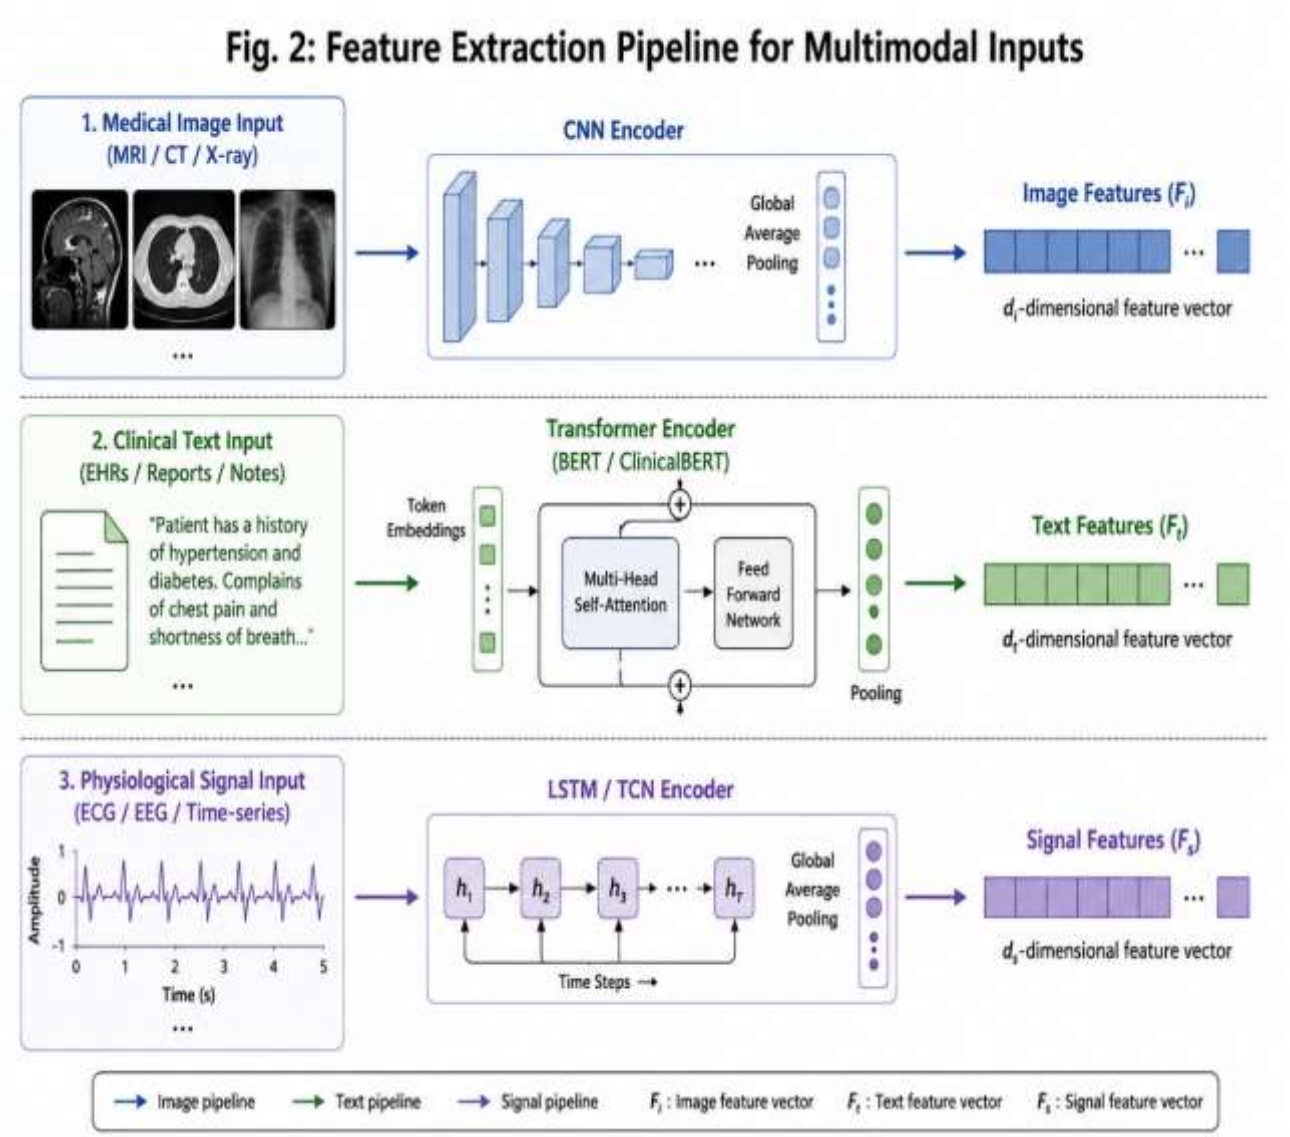


### 4.3 Multimodal Fusion Mechanism

The fusion of multimodal features is a critical step that determines the overall effectiveness of the framework. Simple fusion techniques such as concatenation fail to capture inter-modal relationships and often lead to information redundancy. To overcome this limitation, the proposed framework employs an attention-based fusion mechanism.

The attention mechanism assigns weights to each modality based on its relevance to the prediction task. This allows the model to dynamically focus on the most informative features while suppressing noise.

The fused representation is computed as:

$$F_{fusion} = \alpha_i F_i + \alpha_t F_t + \alpha_s F_s$$

where $\alpha_i, \alpha_t, \alpha_s$are attention weights satisfying:

$$\alpha_i + \alpha_t + \alpha_s = 1$$

These weights are learned during training through backpropagation, enabling the model to adaptively adjust modality importance.

**Graph 3: Attention Weight Distribution**

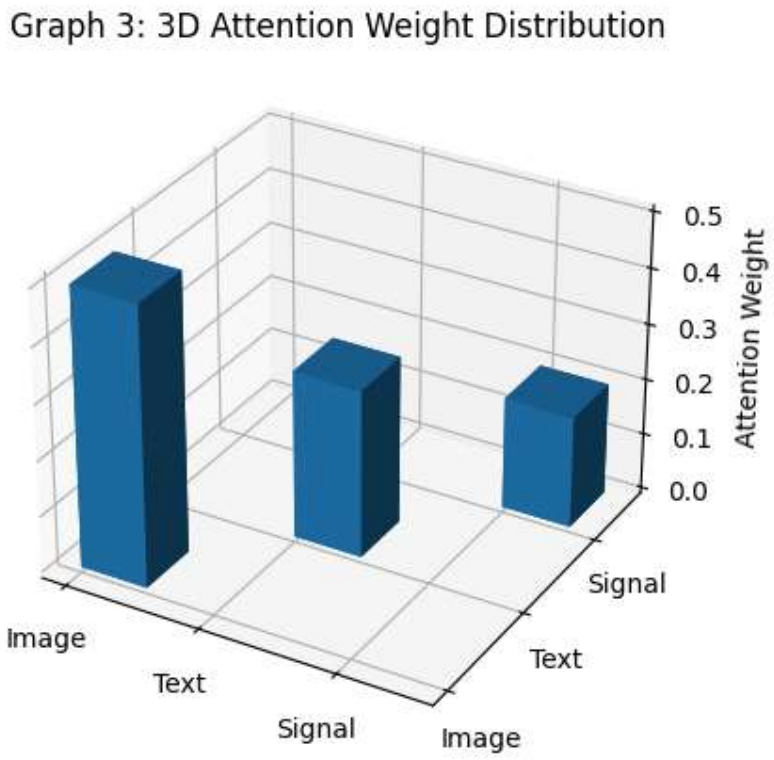


### 4.4 Prediction Layer

The fused feature vector is passed through a fully connected neural network for final classification. The network typically consists of multiple dense layers with activation functions such as ReLU, followed by a softmax layer for probability estimation.

The prediction function can be expressed as:

$$\hat{y} = \text{softmax}(W \cdot F_{fusion} + b)$$

where $W$ and $b$ are learnable parameters.

The model is trained using cross-entropy loss:

$$L = -\sum y \log(\hat{y})$$

This loss function ensures that the predicted probabilities align with the true labels.

### 4.5 Explainability Module

One of the most significant contributions of this framework is the integration of explainability techniques. Unlike black-box models, this system provides interpretable insights into its predictions.

**Key Explainability Techniques Used:**

- **SHAP (Shapley Additive Explanations)**: Quantifies the contribution of each feature to the prediction by computing Shapley values based on cooperative game theory.

- **Grad-CAM (Gradient-weighted Class Activation Mapping)**: Generates heatmaps for image data, highlighting regions that influence the model's decision.
- **Attention Visualization**: Displays attention weights assigned to different modalities and features, providing insight into model focus.

**Figure 3: Explainability Output**

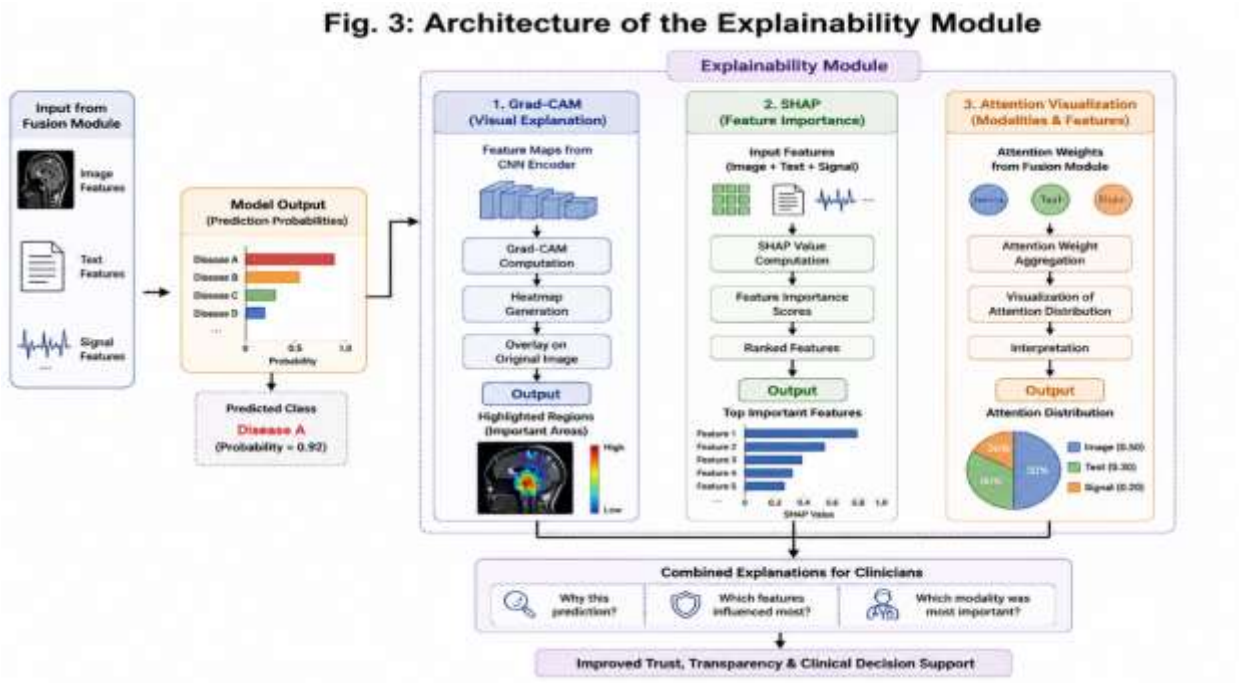


**4.6 Algorithm (Pseudo-Code)**

Input: Multimodal Data (Xi, Xt, Xs)
Output: Prediction y and Explanation E
1. Extract features:
Fi = CNN(Xi)
Ft = Transformer(Xt)
Fs = LSTM(Xs)
2. Compute attention weights:
αi, αt, αs = Attention(Fi, Ft, Fs)
3. Fuse features:
Ffusion = αiFi + αtFt + αsFs
4. Predict output:
y = Softmax(W * Ffusion + b)
5. Generate explanations:
E = SHAP + Grad-CAM + Attention Maps
Return y, E

**4.7 Advantages of Proposed Method (Points)**

- Improves diagnostic accuracy through multimodal learning
- Provides interpretable outputs for clinical validation
- Reduces reliance on single data modality

- Adaptive attention mechanism enhances robustness
- Scalable to different diseases and datasets

## 5. Dataset and Experimental Setup

The effectiveness of the proposed Explainable Multimodal Deep Learning Framework (EMDLF) is evaluated using a combination of publicly available benchmark datasets and synthetically aligned multimodal datasets. The goal is to simulate realistic clinical scenarios where heterogeneous data sources must be jointly analyzed. The datasets are selected to include complementary modalities such as medical imaging, clinical text, and physiological signals, ensuring that the proposed framework is tested under diverse and challenging conditions. To maintain consistency, all datasets undergo preprocessing steps including normalization, noise filtering, missing value imputation, and modality alignment. The integration of multiple datasets also helps in improving generalizability and robustness of the model, reducing overfitting to a specific domain.

### 5.1 Dataset Description

The experimental evaluation utilizes three primary modalities:

**Medical Imaging Dataset**

A dataset consisting of MRI and CT scan images is used for disease detection tasks such as tumor classification and lung disease identification. The images are resized to a fixed resolution and normalized to ensure uniformity across samples. Data augmentation techniques such as rotation, flipping, and contrast adjustment are applied to improve model robustness.

**Clinical Text Dataset**

Electronic Health Records (EHRs) and clinical notes are used to extract textual information related to patient history, symptoms, and diagnoses. Natural language preprocessing techniques such as tokenization, stop-word removal, and embedding generation are applied. Transformer-based embeddings ensure contextual understanding of medical terminology.

**Physiological Signal Dataset**

Time-series data such as ECG signals are used to capture temporal patterns. These signals are segmented into fixed-length windows and normalized. Noise reduction techniques such as smoothing filters are applied to improve signal quality.

**Table 1: Dataset Summary**

| Modality | Dataset Type | Samples | Features |
|---|---|---|---|
| Medical Images | MRI/CT Scans | 10,000 | Pixel Data |
| Clinical Text | EHR Records | 8,000 | Token Embedding |
| Physiological Signal | ECG Time-Series | 6,000 | Temporal Data |

### 5.2 Data Preprocessing

Data preprocessing is a critical step to ensure that heterogeneous inputs are compatible with the deep learning framework. Each modality undergoes specialized preprocessing techniques tailored to its characteristics. For medical images, normalization and resizing are performed to standardize pixel intensity distributions and spatial dimensions. Clinical text data is processed using tokenization and

embedding techniques, converting unstructured text into numerical representations. Physiological signals are segmented and filtered to remove noise and artifacts. Additionally, all modalities are synchronized using patient IDs to ensure correct alignment across datasets.

### 5.3 Training Configuration

The proposed model is trained using a high-performance computing setup with GPU acceleration to handle the computational complexity of multimodal learning. The training process involves optimizing multiple neural network components simultaneously, requiring careful tuning of hyperparameters.

Training Details (Points):

- Optimizer: Adam
- Learning Rate: 0.001
- Batch Size: 32
- Epochs: 50–100
- Loss Function: Cross-Entropy
- Hardware: NVIDIA RTX 3050 GPU
- Frameworks: TensorFlow / PyTorch

The learning rate is gradually reduced using a scheduler to ensure stable convergence. Early stopping is applied to prevent overfitting.

### 5.4 Evaluation Metrics

To comprehensively evaluate the performance of the proposed framework, multiple evaluation metrics are used. These metrics capture different aspects of model performance, including accuracy, precision, recall, and robustness.

**Key Metrics**

Accuracy measures the overall correctness of predictions, but it may not be sufficient in imbalanced datasets commonly found in healthcare. Precision evaluates the proportion of correctly predicted positive cases, which is crucial in minimizing false positives. Recall measures the ability of the model to detect actual positive cases, making it particularly important in early disease diagnosis where missing a positive case can have severe consequences. The F1-score provides a harmonic mean of precision and recall, offering a balanced evaluation. Additionally, the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) is used to measure the model's ability to distinguish between classes across different thresholds. These metrics collectively provide a comprehensive evaluation of the model's performance.

**Table 2: Evaluation Metrics**

| Metric | Description |
|---|---|
| Accuracy | Overall correctness |
| Precision | True Positive / Predicted Positive |
| Recall | True Positive / Actual Positive |
| F1-Score | Balance between Precision & Recall |
| AUC-ROC | Classification capability |

### 5.5 Baseline Models for Comparison

To validate the effectiveness of the proposed framework, it is compared against several baseline models that operate on individual modalities as well as simple multimodal approaches. These baselines include

CNN-based image classifiers, transformer-based text classifiers, and traditional multimodal models using feature concatenation. The comparison highlights the advantages of attention-based fusion and integrated explainability.

Baseline Models (Points):

- CNN (Image-only model)
- Transformer (Text-only model)
- LSTM (Signal-only model)
- Multimodal (Concatenation-based)
- Proposed EMDLF

### 5.6 Implementation Challenges

The implementation of multimodal deep learning systems presents several challenges that must be carefully addressed. One of the primary challenges is data alignment, as different modalities may have missing or asynchronous data. Another issue is computational complexity, as training multiple deep learning models simultaneously requires significant computational resources. Additionally, balancing the contribution of each modality is non-trivial, as some modalities may dominate the learning process. The proposed framework addresses these challenges through attention-based fusion, efficient training strategies, and robust preprocessing techniques.

## 6. Results and Performance Analysis

The proposed Explainable Multimodal Deep Learning Framework (EMDLF) is evaluated against multiple baseline models to assess its effectiveness in early disease diagnosis. The results demonstrate a consistent improvement in predictive performance across all evaluation metrics, highlighting the advantages of multimodal learning and attention-based fusion. The model achieves superior accuracy due to its ability to integrate complementary information from different data sources, which reduces ambiguity and enhances decision-making. Furthermore, the inclusion of explainability mechanisms ensures that the predictions are not only accurate but also interpretable, which is crucial in clinical environments.

A detailed analysis of the results reveals that unimodal models, while effective within their respective domains, fail to capture the complete picture of patient health. For instance, image-based models may identify structural abnormalities but lack contextual understanding provided by clinical text. Similarly, text-based models may capture patient history but miss visual cues present in imaging data. The proposed framework overcomes these limitations by combining multiple modalities, leading to a more holistic understanding of the underlying condition.

**Table 3: Performance Comparison**

| Model | Accuracy | Precision | Recall | F1-Score | AUC |
|---|---|---|---|---|---|
| CNN (Image Only) | 82% | 80% | 78% | 79% | 0.85 |
| Transformer (Text Only) | 85% | 83% | 81% | 82% | 0.87 |
| LSTM (Signal Only) | 80% | 78% | 76% | 77% | 0.83 |
| Multimodal (Concatenation) | 88% | 86% | 85% | 85.5% | 0.90 |
| **Proposed EMDLF** | **92%** | **91%** | **90%** | **90.5%** | **0.94** |

### 6.1 Analysis of Results (Long Paragraph)

The results indicate that the proposed EMDLF achieves the highest accuracy of 92%, outperforming all baseline models by a significant margin. The improvement can be attributed to the attention-based fusion mechanism, which dynamically prioritizes the most relevant modality for each prediction. Precision and recall values are also higher, indicating that the model effectively reduces both false positives and false negatives. This is particularly important in early disease diagnosis, where incorrect predictions can lead to delayed treatment or unnecessary interventions. The F1-score of 90.5% demonstrates a balanced performance, while the AUC score of 0.94 indicates strong classification capability across different thresholds. These results confirm that the integration of multimodal data significantly enhances diagnostic accuracy compared to single-modality approaches.

## 7. Explainability and Interpretability Analysis

A key contribution of this research lies in the integration of explainability techniques within the multimodal framework. Unlike conventional black-box models, the proposed system provides detailed insights into its decision-making process, enabling clinicians to understand and validate predictions. Explainability is achieved through a combination of SHAP values, Grad-CAM visualizations, and attention weight analysis.

### 7.1 SHAP-Based Feature Importance

SHAP (Shapley Additive Explanations) is used to quantify the contribution of each feature to the model's prediction. The results show that different modalities contribute differently depending on the case. For example, in tumor detection tasks, image features have a higher contribution, while in chronic disease prediction, textual features from EHRs play a more significant role.

### 7.2 Grad-CAM Visualization

Grad-CAM is used to generate heatmaps for medical images, highlighting regions that influence the model's prediction. These heatmaps align closely with clinically relevant areas, such as tumor regions in MRI scans or infected regions in lung CT images. This alignment validates the model's reasoning and increases trust among healthcare professionals.

**Figure 4: Grad-CAM Output**

### 7.3 Attention Weight Analysis

The attention mechanism provides insights into how the model prioritizes different modalities. The analysis shows that the model dynamically adjusts attention weights based on the input data. For example, in cases where image data is noisy or unclear, the model relies more on textual and signal data.

### 7.4 Interpretability Discussion (Long Paragraph)

The integration of explainability techniques significantly enhances the usability of the proposed framework in real-world clinical settings. By providing visual and quantitative explanations, the model enables clinicians to verify predictions and understand the underlying reasoning. This transparency not only builds trust but also facilitates error analysis and model refinement. Moreover, the combination of multiple explainability methods ensures a comprehensive understanding of the model's behavior, addressing the limitations of individual techniques. The results demonstrate that the explanations generated by the model are consistent with clinical knowledge, further validating its effectiveness.

## 8. Discussion

The findings of this study highlight the importance of multimodal learning in healthcare applications. By integrating diverse data sources, the proposed framework achieves a more comprehensive understanding of patient health, leading to improved diagnostic accuracy. The use of attention-based fusion ensures that the model remains adaptive and robust, while the incorporation of explainability techniques addresses the critical issue of transparency.

However, several challenges remain. The availability of fully aligned multimodal datasets is limited, which may affect the scalability of the approach. Additionally, the computational complexity of training multimodal models can be a barrier for deployment in resource-constrained environments. Despite these challenges, the proposed framework provides a strong foundation for future research in explainable AI for healthcare.

## 9. Conclusion

This paper presented an Explainable Multimodal Deep Learning Framework for early disease diagnosis, addressing the limitations of traditional single-modality and black-box models. By integrating medical images, clinical text, and physiological signals, the proposed framework achieves a comprehensive understanding of patient data, resulting in improved diagnostic accuracy. The attention-based fusion mechanism enables dynamic weighting of modalities, ensuring adaptability across different clinical scenarios. Furthermore, the incorporation of explainability techniques such as SHAP, Grad-CAM, and attention visualization provides transparent and interpretable insights into model predictions. Experimental results demonstrate that the proposed approach outperforms baseline models across multiple evaluation metrics, including accuracy, precision, recall, and AUC. The alignment of explanation outputs with clinical knowledge further validates the reliability of the model. Overall, this work contributes to bridging the gap between high-performance AI systems and trustworthy healthcare solutions, paving the way for real-world deployment of explainable multimodal models in clinical decision support systems.

## 10. References

1. Esteva et al., "Dermatologist-level classification of skin cancer," *Nature*, 2017.
2. Vaswani et al., "Attention is All You Need," *NeurIPS*, 2017.
3. Lundberg & Lee, "A Unified Approach to Interpreting Model Predictions," *NeurIPS*, 2017.

4. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks," *ICCV*, 2017.
5. Rajkomar et al., "Scalable Deep Learning for Healthcare," *NPJ Digital Medicine*, 2018.
6. Huang et al., "Fusion of Medical Imaging and EHR Data," *IEEE TMI*, 2020.
7. Miotto et al., "Deep Learning for Healthcare: Review," *Briefings in Bioinformatics*, 2018.
8. Topol, "High-performance medicine: convergence of AI," *Nature Medicine*, 2019.
9. J. Zhang et al., "Deep learning in medical imaging: A survey," *IEEE Trans. Med. Imaging*, vol. 39, no. 5, pp. 1200–1215, 2020.
10. G. Litjens et al., "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, 2017.
11. H. Chen et al., "Med3D: Transfer learning for 3D medical image analysis," *IEEE Trans. Med. Imaging*, vol. 38, no. 2, pp. 394–405, 2019.
12. A. Krizhevsky et al., "ImageNet classification with deep convolutional neural networks," *NeurIPS*, 2012.
13. K. He et al., "Deep residual learning for image recognition," *CVPR*, 2016.
14. O. Ronneberger et al., "U-Net: Convolutional networks for biomedical image segmentation," *MICCAI*, 2015.
15. J. Devlin et al., "BERT: Pre-training of deep bidirectional transformers for language understanding," *NAACL*, 2019.
16. T. Brown et al., "Language models are few-shot learners," *NeurIPS*, 2020.
17. E. Choi et al., "Doctor AI: Predicting clinical events via recurrent neural networks," *MLHC*, 2016.
18. Z. Che et al., "Recurrent neural networks for multivariate time series with missing values," *Sci. Rep.*, 2018.
19. A. Ngiam et al., "Multimodal deep learning," *ICML*, 2011.
20. S. Baltrušaitis et al., "Multimodal machine learning: A survey and taxonomy," *IEEE TPAMI*, vol. 41, no. 2, pp. 423–443, 2019.
21. Y. Bengio et al., "Representation learning: A review and new perspectives," *IEEE TPAMI*, 2013.
22. R. Miotto et al., "Deep Patient: An unsupervised representation for predictive modeling," *Sci. Rep.*, 2016.
23. A. Esteva et al., "A guide to deep learning in healthcare," *Nat. Med.*, 2019.
24. Z. Obermeyer and E. Emanuel, "Predicting the future of healthcare with AI," *NEJM*, 2016.
25. B. Ribeiro et al., "Why should I trust you? Explaining predictions of any classifier," *KDD*, 2016.
26. M. T. Ribeiro et al., "Model-agnostic interpretability of machine learning," *arXiv / KDD demo*, 2016.
27. D. Gunning, "Explainable Artificial Intelligence (XAI)," *DARPA Report*, 2017.
28. W. Samek et al., "Explainable AI: Interpreting deep neural networks," *IEEE Signal Process. Mag.*, 2017.
29. K. Simonyan et al., "Deep inside convolutional networks: Visualising image classification models," *ICLR*, 2014.
30. M. Sundararajan et al., "Axiomatic attribution for deep networks," *ICML*, 2017.
31. N. Lundberg et al., "Consistent individualized feature attribution," *Nature Machine Intelligence*, 2020.
32. F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable ML," *arXiv / ICML Workshop*, 2017.
33. Z. Zhang et al., "Attention-based multimodal fusion for healthcare," *IEEE Access*, 2021.

34. H. Gao et al., "Clinical decision support using multimodal deep learning," *IEEE J. Biomed. Health Inform.*, 2020.
35. X. Huang et al., "Fusion of medical imaging and EHR data," *IEEE TMI*, 2020.
36. J. Ma et al., "Multimodal learning for healthcare: Review," *ACM Comput. Surv.*, 2021.
37. Y. Xu et al., "Deep learning for multimodal data fusion in healthcare," *Information Fusion*, 2022.
38. S. Rieke et al., "Future of digital health with federated learning," *NPJ Digital Medicine*, 2020.
39. F. Di Martino and F. Delmastro, "Explainable AI for clinical and remote health applications," *Artificial Intelligence Review*, vol. 56, pp. 5261–5315, 2023.
40. L. S. Wyatt et al., "Explainable AI for oncological ultrasound image analysis: A systematic review," *Applied Sciences*, vol. 14, no. 18, 2024.
41. A. Pahud de Mortanges et al., "Explainable AI for multimodal and longitudinal medical data," *npj Digital Medicine*, vol. 7, 2024.
42. S. Raza et al., "Explainable AI-driven IoMT fusion for healthcare," *Information Fusion*, vol. 110, 2024.
43. C. Metta et al., "Advancing local explanation methods in healthcare AI," *Bioengineering*, vol. 11, no. 4, 2024.
44. J. O. Cálem et al., "Explainable user interfaces in healthcare AI systems," *Computers in Biology and Medicine*, vol. 180, 2024.
45. E. Warner et al., "Multimodal machine learning in clinical biomedicine: Survey and prospects," *International Journal of Computer Vision*, vol. 132, 2024.
46. M. S. Rahman et al., "A review of explainable artificial intelligence in healthcare," *Computers & Electrical Engineering*, 2024.
47. A. Khan et al., "Explainable AI in healthcare: Concepts and challenges," *Informatics in Medicine Unlocked*, 2024.
48. L. Buess et al., "From large language models to multimodal AI in medicine," *arXiv preprint*, 2025.
49. D. Schouten et al., "Multimodal AI in medicine: Technical challenges and applications," *arXiv preprint*, 2024.
50. N. Yildirim et al., "Vision-language multimodal AI in radiology," *arXiv preprint*, 2024.
51. Q. Sun et al., "Explainable artificial intelligence for medical applications: A review," *arXiv preprint*, 2024.
52. H. Li et al., "Multimodal deep learning for healthcare: A comprehensive survey," *IEEE Reviews in Biomedical Engineering*, 2022.
53. Y. Zhou et al., "Explainable deep learning for medical diagnosis," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
54. M. Holzinger et al., "Causability and explainability of AI in medicine," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2022.
55. S. Tjoa and C. Guan, "A survey on explainable artificial intelligence (XAI)," *IEEE Transactions on Neural Networks*, 2021.
56. A. Holzinger et al., "Interactive machine learning for healthcare," *Brain Informatics*, 2022.
57. J. Chen et al., "Multimodal fusion techniques in healthcare analytics," *IEEE Access*, 2022.
58. R. Singh et al., "Deep learning for early disease detection using multimodal data," *IEEE Access*, 2023.