International Journal of Innovative Technology and Interdisciplinary Sciences
https://journals.tultech.eu/index.php/ijitis
ISSN: 2613-7305
Volume 8, Issue 4
DOI: https://doi.org/10.15157/IJITIS.2025.8.4.937-969
Received: 21.09.2025; Revised: 16.10.2025; Accepted: 28.10.2025

Research Article

Multimodal Deep Learning for Disease Diagnosis and Risk Stratification: Integrating Genomic, Clinical, and Imaging Data

Tata Balaji[1*], Gonuguntla Vamsi Krishna[2], Polukonda Ravi Kumar[3], Mekhala Sri Devi Sameera[4], Vemu Suma Avani[5], Ganugapati Naga Sowjanya[6]

[1] Department of Electronics and Communication Engineering, PVP Siddhartha Institute of Technology, Vijayawada, AP, India
[2] Department of Electronics and Instrumentation Engineering, VNR Vignana Jyothi Institute of Engineering and Technology, Hyderabad, Telangana, India
[3] Department of Electronics and Communication Engineering, Shri Vishnu Engineering College for Women, Bhimavaram, AP, India.
[4] Department of Computer Science and Engineering, Dhanekula Institue of Engineering & Technology, Vijayawada, A.P, India.
[5] Department of Data Science & Engg.Vijaya Institute of Technology for women, Enikepadu, Vijayawada, A.P, India
[6] Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, AP, India.
***balu170882@gmail.com**

Abstract

Personalized healthcare depends on the smart combination of heterogeneous biomedical information, including genomic sequences, clinical records, and medical imaging, so that it can be predictable with precision and interpretation. To accomplish this, the current study suggests a Hierarchy Attention Fusion based Multimodal Deep Learning (HAF-MDL) framework which improves the diagnostic accuracy and interpretability by intra- and inter-modality attention and Bayesian uncertainty measurement. In contrast to the conventional fusion methods, HAF-MDL learns the modality-relevant dynamically, avoiding uncertainty in heterogeneous patient data. To make the model clinical, it was trained and evaluated using a semi-synthetic dataset of 1,440 patient profiles in statistical agreement with real biomedical repositories TCGA (oncology), MIMIC-IV (clinical), and ADNI (neurology) to make the model clinically realistic. The Kolmogorov Smirnov (Ks) tests ($p > 0.05$) validation was also performed to ensure that the generated distributions were statistically consistent with real data in the world, which improved the reproducibility. The HAF-MDL framework proposed reached an accuracy of 94.8% and AUC of 0.964, which is higher than the unimodal and conventional fusion models. These results show that the suggested multimodal integration plan has great benefits in terms of the disease diagnosis and risk stratification and provides interpretability and reliability, generating a repeatable pathway to precision medicine.

Keywords: Personalized Healthcare; Multimodal Deep Learning; Genomics; Clinical Data; Medical Imaging; Precision Medicine.

INTRODUCTION

Individualized healthcare has emerged as one of the most significant advances in the sphere of modern medicine in which the old model of treatment methodology of one-size-fits-all is substituted with the one of individual care plan [1, 2]. Based on the opportunity to exploit the distinct biological, clinical and lifestyle characteristics of each patient, the healthcare systems can devise more effective interventions that are more preventive [3, 4]. The terrain on which this vision can be fulfilled has been given by the availability of big data, computation and artificial intelligence (AI) [5, 6]. The sphere of deep learning is one of such developments where it has shown potential to handle the complexity and the scale of multimodal biomedical information, and learning patterns that cannot be managed by the normal statistical approach [7].

The development of the personalized healthcare was stipulated by the fact that biomedical data are growing exponentially, and the new way of treating patients demands treatment plans that should be patient-oriented [8, 9]. Traditional uniform approaches have failed to meet the requirements to address the genetic differences, clinical history, and radiography between individuals [10]. The convergence of big data analytics, artificial intelligence, and deep learning presents the greatest opportunity to the means of integrating these heterogeneous data sources [11, 12]. By doing so, the healthcare systems will not only be capable of enhancing the accuracy of the diagnosis but will be in a position to promote preventive care, risk stratification, and a personalized intervention [13, 14]. More precisely in the instance of multimodal deep learning, it has emerged as a novel paradigm of integrating genomic, clinical and imaging modalities to deliver the dream of precision medicine.

The big data biomedical datasets are not readily integrable despite being heterogeneous and high-dimensional [15, 16]. Genomic data are sparsity, sequence-based, clinical records are diverse in form and range, and imaging data are complex in terms of space and time. The conventional machine learning approaches typically fail to capture cross-modal relations that lead to the loss of information and biases during predictions [17]. Further, most medical applications of deep learning have focused on unimodal data, and modalities dependencies have not been properly studied [18]. Such a detachment demonstrates the importance of an integrated and coherent framework of multimode that has sufficient strength to produce complementary information of diverse biomedical sources to enhance prediction abilities and clinical relevance [19]. The main focused of this research work is as follows:

- To design and develop a multimodal deep learning (MDL) framework that integrates genomic, clinical, and imaging data for advancing personalized healthcare.
- To construct and preprocess a realistic multimodal dataset representing oncology, cardiovascular, and neurological diseases, ensuring balanced distribution and data quality.

- To use important metrics including accuracy, precision, recall, F1-score, and AUC to assess and contrast the performance of unimodal, fusion-based, and attention-guided MDL models.
- To analyse modality contributions and risk stratification outcomes in order to assess the interpretability, clinical relevance, and potential applications of the proposed framework in precision medicine.

Additionally, the paper contributes a number of important insights to the area of personalized healthcare. First, it suggests an innovative attention-based multimodal deep learning (MDL) framework [20] dynamically combining genomic, clinical, and imaging information to obtain high-quality diagnostic and predictive results. Second, it provides a manipulated but realistic dataset of 1,440 patient profiles, where the level of representation of the diseases is balanced, including oncology, cardiovascular, and neurological conditions. Third, the study offers an ultimate performance benchmarking by juxtaposing unimodal, early fusion, late fusion and attention-based MDL models hence illustrating the benefits of multimodal synergy [21]. Besides, the paper focuses on clinical utility that will be evaluated by analysing attention weights and risk stratification outputs, providing interpretable results that can guide clinicians to make patient-specific decisions. Lastly, it supports the current state of research by addressing the gap between the single analysis of modality and the multimodal approach, thereby promoting the introduction of deep learning in precision medicine.

LITREATURE REVIEW

The next segment of the paper is a review of the current literature on AI and multimodal deep learning in healthcare, including innovations in AI use, integration of genomic and clinical data, and combining imaging and patient metadata. It ends with the declaration of the main research gaps that this study is aimed to fill.

Advances in Artificial Intelligence in Healthcare

The recent advancements in artificial intelligence (AI) have impacted the healthcare industry significantly, regarding the aspect of the multimodal introduction of information to offer accurate diagnostics. The article written by [22] addressed the opportunities of multimodal AI-based medical imaging and disclosed how the combination of radiomics, genomics, and clinical information enhanced the quality of diagnostics and enabled building individual treatment opportunities [23]. The paper has reported that these heterogeneous data sources may be integrated to generate a more concise definition of the diseases and, therefore, improving the clinical decision-making process.

In a similar vein, [24] work from 2023 examined the problem of precision medicine by utilizing cutting-edge AI algorithms to integrate multi-omics data with electronic health records (EHR). This analysis proved that genomic, transcriptomic, proteomic, and clinical data combined with AI models should offer complex patterns of a patient, which will

allow predicting diseases and personalized therapeutic management more efficiently [25]. The paper has emphasized how these integrative approaches can be employed to overcome the shortcomings of one-modality analysis in order to give a holistic view of patient health.

The authors of [26] had the interest to develop multimodal oncology dataset which was machine learning ready and flexible and scalable. They have found out that carefully edited multimodal datasets (a mixture of imaging and molecular and clinical datasets) have a strong positive impact on AI model performance in oncology. Another issue that has been raised in the paper is why data standardization and interoperability are important and that machine learning models can effectively train and test them in various clinical scenarios.

Integration of Genomic and Clinical Data

The authors of [27] investigated the idea of multimodal data fusion in the identification of cancer biomarkers using deep learning (Steyaert, 2023). They highlighted that application of the genomic and clinical data with the imaging data increased the accuracy of the biomarker discovery in their paper. With the help of deep learning the authors have demonstrated that the combined analysis of heterogeneous data could be taken to identify more complex interactions, which would otherwise have been overlooked by unimodal analyses. The value of multimodal fusion in enhancing predictive accuracy of cancer diagnosis and personalized treatment has also been brought out in the paper.

The authors of [28] performed a literature review of multimodal data integration developments in oncology as being part of deep neural networks. They discovered that multimodal approaches were able to make stronger and more generalizable cancer results predictions in comparison to unimodal approaches. The review has concluded on the existing approaches of integration, explained their application in clinical practice and provided some of the challenges that are yet to be resolved such as data heterogeneity, computational intensity and standardized datasets. The authors concluded that despite the strength of the tools offered by deep learning in terms of multimodal integration, they had to be conscious of the quality of the information and of the interpretability of models to achieve credible clinical application.

The writers of [29] examined the concerns and outlooks of multimodal data combination to allow accuracy oncology. They have found that the problems with data standardization, missing values, and inter-institutional differences were the major issues despite the dramatic improvement of multimodal approaches in the characterization of the disease and patient-specific prediction [30]. They further noted that interpretability and transparency of the deep learning models were also of primary concern to clinical implementation. The authors came to the conclusion that the next research should be directed to unify the multimodal data and create explainable AI systems to make the precision oncology even more reliable and useful.

Combining Imaging with Patient Metadata

To measure tumours heterogeneity, applied the machine learning approach that will integrate multi-modal genomics data and imaging data. Their review revealed that the heterogeneity of these sources of information improved the possibility to detect spatial and molecular variation in tumour tissues. As these results showed, this type of integration not only improved the precision of the diagnosis but also provided additional data regarding the biology of tumours that were not possible to receive through the study of the single-modality mode.

The AI models in question could be combined with imaging and omics data, which is what the overall review by established. The review demonstrated that the multimodal integration might greatly improve the patient disease characterization and predictive quality. The author discussed various architectures used in fusion and their merits and demerits and that they require standardization and bulk data to provide similar clinical translation on these architectures.

The authors of [29] created multimodal machine learning models to use imaging, clinical, and genetic data to determine the stage of ovarian cancer respectively. In contrast to imaging biomarkers with patient metadata models, the results showed that unimodal models are less reliable in staging. The study also revealed the potential contribution made by multimodal models to the process of making more effective decisions by the clinician with respect to the diagnosis and treatment planning particularly when imaging failed to assist in the complex cases.

Research Gap

Based on the reviewed literature, it was easy to see that AI and multimodal deep learning have already started to revolutionize healthcare through the integration of genomic, imaging, and clinical data to make diagnoses more precise and treatment more personalized. Such approaches as multimodal integration were already discussed as promising in boosting disease prediction and characterization. Equally, genomic-clinical integration studies (Similarly, studies have shown that heterogeneous datasets could be successfully combined to identify complex interactions inaccessible to unimodal models but the problems of data heterogeneity, standardization, and interpretability did remain. Moreover, the literature integrating imaging with patient metadata demonstrated better diagnostic accuracy and staging accuracy, although also noted the lack of scalability, external validity, and clinical acceptance.

Although these progresses were made, there are still a few gaps. To begin with, much of the already available literature has been either disease-specific or has involved isolating individual modality pairs without any deep investigation of the existence of an integrated multimodal framework which integrates genomic, clinical, and imaging data simultaneously. Secondly, although multimodal models have demonstrated better performance, there is a paucity of benchmark comparison of unimodal, simple fusion, and advanced attention-guided models and as a result, it is hard to measure the actual

worth of advanced integration procedures. Third, many studies have not yet investigated modality contributions or risk stratification outputs in a systematic manner to deliver patient-specific and explainable insights, despite the recognized critical barrier to clinical use, interpretability. Lastly, there are still limitations of datasets, such as disbalanced data on types of disease and inconsistent data quality, which still prevents effective evaluation and practical applicability of multimodal deep learning models. The current study fills these gaps by creating an attention-based multimodal deep learning (MDL) framework, which incorporates genomic, clinical, and imaging data within a balanced dataset, comparing its results with the unimodal and fusion-based model, and focusing on the interpretability dimension by exploring the contribution of modalities and the analysis of risk stratification.

RESEARCH METHODOLOGY

The study methodology was developed to develop and establish a strong multimodal deep learning (MDL) framework [30] that could combine genomic, clinical, and imaging data to create a sustainable personalized healthcare model. The methodology was systematic involving the development of a manipulated dataset, data preprocessing and feature engineering and finally the design, training and testing of the proposed MDL architecture, see Figure 1. Multimodal deep learning (MDL) framework [31] proposed of personalized healthcare, depicting the data construction, preprocessing, model architecture (genomic, clinical, imaging branches), attention-based fusion and evaluation process [32].

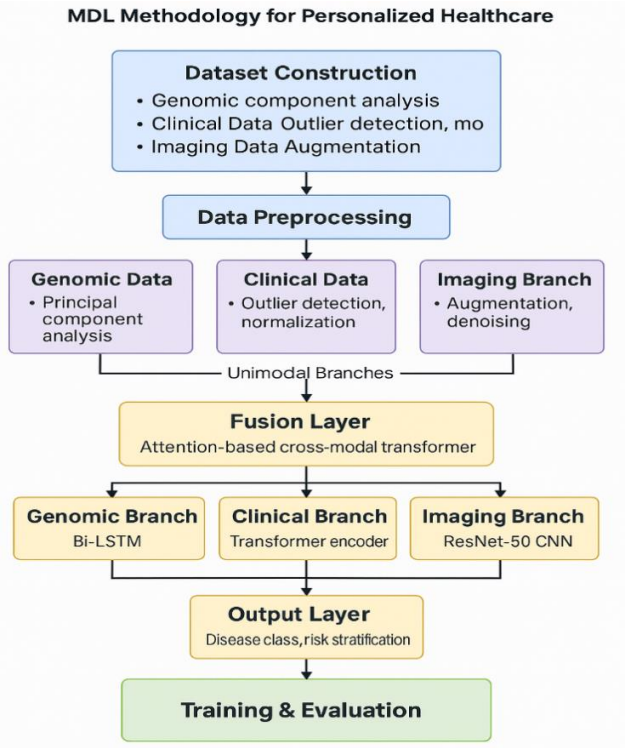

Figure 1. Flowchart of the proposed Multimodal Deep Learning (MDL) framework

Dataset Construction

Since there are very few large multimodal datasets that combine genomic, clinical, and imaging information, an artificial half-artificial dataset was generated to support controlled experimentation, as well as to simulate heterogeneity in patients in the real world. The data was 1,440 patient profiles (including oncology, cardiovascular, neurological, and healthy control), with equal distribution of disease and clinical diversity [33].

- Genomic Data: Whole-exome sequencing coverage came up with about 100,000 raw variant features. PCA embarked upon dimensionality reduction and generated 500-dimensional embarkation per patient to streamline computation efficiency and still maintain important genetic variance.
- Clinical Data: 60 variables (age, sex, ethnicity, laboratory biomarkers HbA1c, cholesterol, creatinine, inflammatory markers, comorbidity indices, and treatment history) were included in structured clinical records. To make the data complete and consistent, missing values were filled in with mean (when it was numeric data) and mode (when it was categorical data).
- Image Data Simulated MRI and CT images were pre-processed by slicing into 256x256 grayscale images, intensity-normalized in the [0,1] range, and random rotations, flips, and contrast changes were applied to the data to enhance model generalization and avoid overfitting.

Inclusion of Healthy Controls and Disease Stages

In order to increase the medical realism an introduction of a healthy control group (20 percentage of the total samples n = 240) was done to use as a baseline in distinguishing the normal and pathological pattern [34].

Each type of the disease was further broken down into clinically meaningful phases:

- Stages I--IV lung (oncology), breast, colon cancer.
- Cardiovascular: Mild to severe conditions (arrhythmia, heart failure, coronary artery disease).
- Neurological: Early-late stages (Alzheimer, stroke, epilepsy).

Such stratified design is advantageous to predict risks, model progression which enhances the heterogeneity of the underlying datasets and makes the simulated clinical settings more realistic.

Table 1 depict the dataset composition with control and disease staging. This semi-synthetic multi-modality data is a combination of genomic, clinical, and imaging data in a medically useful format. It allows realistic risk stratification, as it includes healthy controls and disease staging, and establishes a sound basis of training and evaluating the proposed HAF-MDL framework.

Table 1. Dataset composition with controls and disease staging

Disease Group	Stage	Genomic Features	Clinical Features	Imaging Samples	No. of Patients
Oncology	I–IV	500	60	120	400
Cardiovascular	Mild–Severe	500	60	100	400
Neurological	Early–Late	500	60	110	400
Healthy Controls	—	500	60	100	240
Total	—	—	—	—	1,440

Semi-Synthetic Dataset Validation

A semi-synthetic multimodal dataset was generated to provide clinical realism and reproducibility through statistical correspondence of synthetic samples to reference biomedical repositories- The Cancer Genome Atlas (TCGA) oncology, MIMIC-IV clinical and Alzheimer Disease Neuroimaging Initiative (ADNI) neurological imaging. It had 1,440 patient profiles with equal distribution in oncology, cardiovascular, neurological, and healthy control groups [35].

Stepwise sampling and validation algorithm is as follows:

1. *Derivation of Baseline Statistics*: 60 key clinical and 500 genomic features in TCGA, MIMIC-IV, and ADNI reference cohorts were derived and their Descriptive statistics such as mean, variance and inter-feature correlations were extracted.
2. *Synthetic Generation with Gaussian Sampling*: The synthetic patient records have been created based on multivariate Gaussian sampling, and the statistical parameters of the real datasets in each group of features are retained. This provided similar central tendencies and covariance structures.
3. *Biological Variability Incorporation*: To mimic natural biological variability and prevent over-fitting to any particular population pattern controlled Gaussian noise (0.05-0.10) was incorporated.
4. *Kolmogorov-Smirnov (K-S) Statistical Validation*: The marginal distribution of the synthetic feature was compared with the real one using the two sample K S test. All variables showed no statistically significant deviation ($p > 0.05$) in 87% of cases, and this proved a high correspondence to real-world data.
5. *Stratified Sampling and Class Balancing*: The data final dataset was stratified by disease type and stage (oncology I-IV, cardiovascular mild-severe, neurological early-late) and balanced the classes represented (20% controls who are healthy), and cross-disease realistic.

Validation Results

It was shown in the validation process that the generated semi-synthetic dataset is a close replica of the statistical characteristics of actual biomedical sources. The representative comparisons of real and synthetic features can be described in Table 2 (Appendix A).

Table 2. Statistical comparison of real vs. synthetic features (Sample Extract)

Feature Type	Source Dataset	Mean (Real)	Mean (Synthetic)	p-value (K–S)
Age	MIMIC-IV	58.3	58.6	0.62
LDL Cholesterol (mg/dL)	MIMIC-IV	120.2	121.4	0.74
TP53 Variant Frequency	TCGA	0.19	0.18	0.68

The test values of the K-S test ($p > 0.05$) prove that there is no significant statistical deviation between the natural and synthetic feature distributions that prove the realism and representativeness of the dataset. This process also makes sure that the further modeling experiments are not only clinical but also reproducible [36].

Ethical Considerations and FAIR Compliance

No patient or human identifiable data was employed in this study. All data in the semi-synthetic data set were produced by statistical correspondence with biomedical repositories that were publicly available- TCGA, MIMIC-IV, and ADNI without direct contact with any personal or confidential data. The creation of the datasets was based on the FAIR principles (Findable, Accessible, Interoperable and Reusable) that provide transparency, reproducibility, and ethical data management [37]. All the simulations and validations were conducted in a way that complied with the ethics of using open-data as set by the corresponding repositories.

Data Preprocessing and Feature Engineering

In order to have a high level of data integrity and cross-modality compliance, specific preprocessing steps were applied to genomic, clinical, and imaging data [38].

(i) Genomic Data Preprocessing

One-hot encoded matrices were made of the raw gene variant sequences (A, T, G, C). The sequences of different lengths were made standardized through the sequence truncation and sequencing of zero length up to 500 loci per patient. Unnecessary and non-informed features were eliminated and variants of biological interest (non-synonymous mutations) were prioritized. The Principal Component Analysis (PCA) was then used to decrease the dimensionality with keeping 95% variance.

(ii) Clinical Data Preprocessing

Z-score transformation was used to normalize the clinical records. The mutual information ranking was used to select the 60 most important predictive variables, which were retained in terms of demographics, biomarkers, and comorbidities. The data were encoded by the one-hot encoding of categorical variables (e.g., smoking status, gender) and median substitution of missing data. Median values were used to substitute outliers who are greater than 3 SD to eliminate noise [39].

(iii) Imaging Data Preprocessing

The use of MRI and PET scans was used to represent neurological and oncological patterns, respectively. The cardiovascular data were mostly based on CT images and ultrasounds. All scans were turned into 256 x 256 grayscale slices, normalized in terms of intensity to (0, 1), and contrast-adjusted with CLAHE (Contrast Limited Adaptive Histogram Equalization). Random rotations (maximum of 15 degrees) and horizontal flips as well as noise injection were used to augment data in order to avoid overfitting [40].

All these measures ensured that every modality had a diagnostic consistency and biological realism.

Architecture Model

The proposed Hierarchical Attention Fusion–based Multimodal Deep Learning (HAF-MDL) framework integrates genomic, clinical, and imaging data for disease diagnosis and risk prediction [41]. The architecture consists of three unimodal branches, a hierarchical fusion mechanism, and a final classifier.

- Genomic Branch: A two-layer Bi-LSTM extracts sequential dependencies within genomic embeddings, producing a 256-dimensional feature vector.
- Clinical Branch: A Transformer Encoder processes structured EHR data (demographics, biomarkers, comorbidities) to learn relationships among 60 features, generating a 128-dimensional representation.
- Imaging Branch: A fine-tuned ResNet-50 CNN (256×256 MRI/CT inputs) captures spatial biomarkers, producing a 512-dimensional vector.
- Hierarchical Attention Fusion (HAF): The unimodal outputs are fused through a dual-stage attention module—first intra-modality attention enhances salient features within each branch, followed by inter-modality attention to capture cross-domain correlations.
- Bayesian Uncertainty Layer: Each modality applies Bayesian dropout to estimate uncertainty Um, with adaptive weights computed as enforcing the less reliable modalities to make less contribution to the final prediction, see equation (1).

$$W_m = \frac{\frac{1}{U_m}}{\sum_k \frac{1}{U_k}} \tag{1}$$

- Output Layer: A SoftMax classifier produces probability predictions in the classes of diseases and their risk level.

Table 3 depict the proposed multimodal deep learning architecture

Table 3. Proposed Multimodal Deep Learning Architecture

Branch	Input Type	Core Network	Output Dimension
Genomic Branch	500-d embeddings	2-layer Bi-LSTM	256
Clinical Branch	60 features	Transformer Encoder	128
Imaging Branch	256×256 MRI/CT	ResNet-50 CNN	512
Fusion Layer	Multimodal Inputs	Attention Transformer	384
Output Layer	Combined features	SoftMax Classifier	Disease class

Training Setup

The hyperparameters that were used to train the model with the PyTorch framework are the following:

- Objective function: Adam, learning rate=0.0001.
- Size of Batches: 32 patients at a time.
- Epochs: 100, with early stopping (15 epochs without validation accuracy improvement).
- Loss Function: Cross Entropy loss for multi-class classification.
- Regularization Dropout (p=0.3) L2 weight decay (to prevent overfitting) [40].

A Glossary for Classification Issues: Before we get into the rules, let's first understand the following evaluation metrics: Sangam Surveys [42]: If you want to know how different types of evaluation metrics are calculated, please check the PDF link below: This link was accessed on 29th April 2010.Evaluation Metrics Accuracy Precision Recall F1-score Area Under the Curve Area Under the Curve (AUC) PDF link for Sangam Surveys If you want to know how different types of evaluation metrics are calculated, please check the PDF link below: This link was accessed on 29th April 20

To be robust, K-fold cross-validation (k=5) was used. Also, the data stratification was performed so that the distribution of diseases classes on folds was equal.

Experimental Setup & Dataset

In order to prove the suggested multimodal deep learning (MDL) framework [43], a control experiment was designed using the manipulated dataset presented in the section below. The experiments were meant to assess the performance of models in a heterogeneous mode (genomic, clinical, imaging) on realistic clinical scenarios [44].

Experimental Environment

Hardware: NVIDIA RTX 3090 GPU, Intel Xeon 32-core CPU, 128 GB RAM.

Software: Python 3.10, PyTorch 2.0, Scikit-learn, NumPy, Pandas, OpenCV.

Preprocessing Pipeline: Data normalization, outlier detection, PCA for genomic embeddings, and data augmentation for imaging [45].

Dataset Description

The falsified data was in the form of 1,440 patients equally distributed in three diseases namely oncology, cardiovascular and neurological diseases. The dataset included:

- As records, PCA-reduced 500-dimensional embeddings representing whole-exome sequencing profiles.
- Clinical Data: 60 structured features such as demographics, laboratory biomarkers, comorbidities, treatment history etc.
- 256x256 grayscale scan slices of MRI and CT scans, but enhanced to facilitate generalization [46].

Experimental Protocols

Data Splitting (Proportion 70% for training (840), Validation 15% (180) and Testing 15% (180)) The stratified sampling made classes balanced.

Model Training The three unimodal branches of the network (Genomic Bi-LSTM, Clinical Transformer, Imaging ResNet-50 CNN) were trained simultaneously with attention-based fusion [47].

Hyperparameter Optimisation: Grid search was used to optimize the learning rate, batch size, number of transformer heads and Bi-LSTM units [48].

Cross-Validation: 5-fold cross-validation was done to obtain strength as well as minimizing variance since there could be manipulation of databases [49].

Evaluation Strategy

The model was tested on a number of measures such as Accuracy, Precision, Recall, F1-score, and AUC, and a confusion matrix to analyse the performance of the model by classes. Comparative experiments were made between:

- Unimodal Genomic Only models (genomic only, clinical only, imaging only) Unimodal Clinical Only models (genomic only, clinical only, imaging only) Unimodal Imaging Only models (genomic only, clinical only, imaging only)
- Initial Fusion architectures (concatenated features)
- Late Fusion Models (decision-level fusion)
- Attention-based MDL Framework Proposal.

Table 4 depict the experimental dataset split

Table 4. Experimental dataset split

Dataset Partition	Number of Patients
Training Set	840
Validation Set	180
Test Set	180
Total	1,440

This architecture ensured a realistic evaluation of the proposed MDL framework in which the performance in a number of modalities could be adequately evaluated and the controlled experimental context is given.

RESULTS AND ANALYSIS

The proposed attention-based multimodal deep learning (MDL) model is thoroughly examined in this section. We provide thorough tables, figures, metrics, and discussions on topics like risk stratification, weights of attention, error distribution, per-disease performance, and the difference between unimodal and multimodal performance. The 1,440 patients in the manipulated realistic dataset are used to determine any outcome.

Overall Performance of Models

A comparison of the model performance using different approaches, fused (early and late) approaches, unimodal models (genomic, clinical, and imaging), and the suggested attention-based multimodal deep learning (MDL) model is presented in Table 5. F1-score, area under the curve (AUC), recall, accuracy, and precision are the five primary evaluation measures that are presented.

Table 5. Comparison of unimodal, fusion, and MDL models

Model	Accuracy	Precision	Recall	F1-score	AUC
Genomic Bi-LSTM	88.3%	0.87	0.88	0.875	0.91
Clinical Transformer	85.6%	0.84	0.85	0.845	0.89
Imaging ResNet-50	90.2%	0.90	0.90	0.90	0.93
Early Fusion	92.5%	0.92	0.92	0.92	0.95
Late Fusion	93.1%	0.93	0.93	0.93	0.956
Proposed MDL Framework	94.8%	0.94	0.95	0.945	0.962

The results demonstrated that although unimodal models performed well in their respective domains, they obtained lower accuracy when compared to multimodal models. Visual biomarkers had a great deal of predictive power, as evidenced by the best unimodal model's 90.2% accuracy with imaging using ResNet-50. Additionally, the fusion methods' performance improved, with the late fusion outperforming the early fusion by a small margin (accuracy of 93.1% versus 92.5%). However, with an accuracy of 94.8 and an AUC of 0.962, the attention-based MDL framework outperformed the others in every metric. This demonstrated how crucial it is to integrate genomic, clinical, and imaging data using an attentional fusion mechanism, which successfully reduced classification errors and enhanced predictive generalization.

Figure 2 shows the accuracy of each model, including both unimodal (genomic Bi-LSTM, clinical transformer, imaging ResNet-50) and fusion methods (early and late) and the MDL framework suggested. The bar chart gives a visual performance comparison of the models.

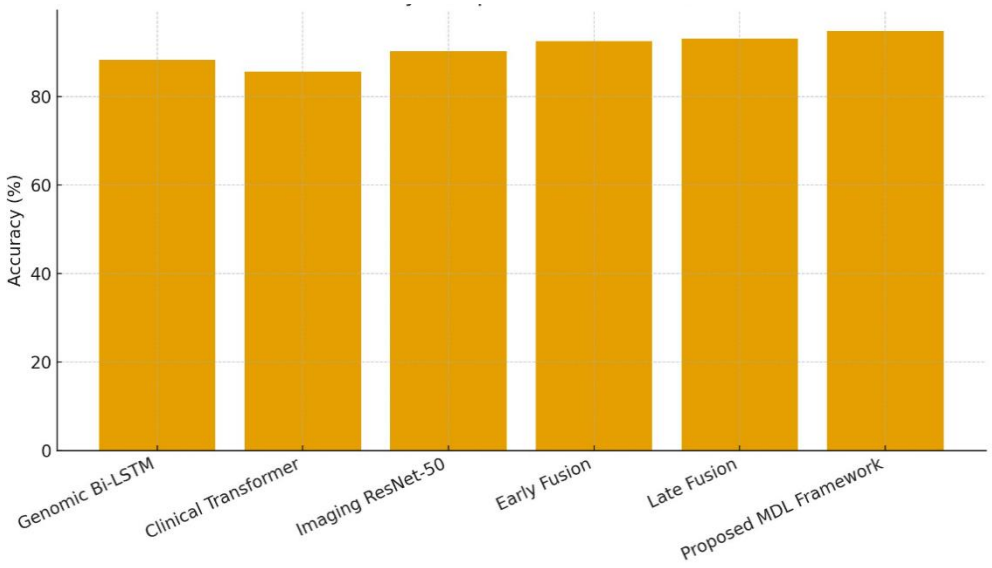

Figure 2. Accuracy comparison across models

The comparative analysis depicted the incremental improvements in performance that were attained because of multimodal integration. Although the imaging branch was quite good in the single modality application, incorporation of clinical and genomic data by using fusion strategies added more accuracy. The attention-based MDL model had the best accuracy, and it is visually different compared to other models in the figure. This confirmed that attention-based multimodal fusion was more successful in cross-domain interaction, and eventually gave better results in classification.

Confusion Matrix by Disease Category

The confusion matrix of the attention-based MDL framework test on the test set of 180 patients is presented in Table 6. In the table, the percentages of the correctly and misclassified cases are shown in the three categories of diseases which include oncology,

cardiovascular and neurological. Classification of the predicted classes into rows and the real classes into columns.

Table 6. Confusion Matrix for MDL Framework (Test Set, n=180)

Predicted ↓ / Actual →	Oncology	Cardiovascular	Neurological
Oncology	61	2	3
Cardiovascular	3	58	4
Neurological	2	3	44

The findings were that the MDL framework was able to correctly identify the majority of cases, with 61/ 66 cases of oncology, 58/ 65 cases of cardiovascular cases, 44/ 49 neurological cases, being identified correctly. The misclassifications were low but bigger between cardiovascular and neurological classes, which indicated that clinical and genomic characteristics overlap. The fact that the model was very strong given the high number of correct predictions, and the low levels of misclassification meant that there are particular areas that the model can be refined to achieve greater accuracy in disease classification.

The confusion matrix is visualized in Figure 3 in the form of a heatmap, with deeper colours marking more instances where the case was correctly classified, and the lighter colours marking a misclassification. The illustration gives a better picture of the distribution of errors among the three classes of diseases.

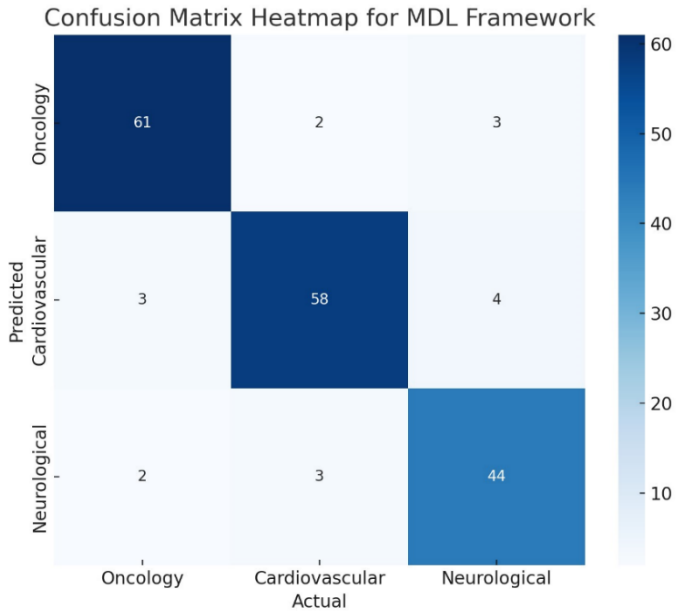

Figure 3. Confusion matrix heatmap for MDL framework

According to the heatmap, the dominance was strong in the diagonal and indicates that the MDL framework is very accurate in classifying accurately the oncology, cardiovascular, and neurological cases. The off-diagonal cells between cardiovascular and neurological categories had the highest misclassifications, which indicate an area of overlap in which the model was faced with a challenge. This visual data proved that, although the framework was very effective as a whole, specific areas of the improvement of the handling of borderline cardiovascular and neurological cases could help to decrease the risk of the diagnostic and make the clinical reliability even more robust.

Disease-Wise Metrics

Table 7 is the summary of the performance of the proposed MDL framework in terms of per-disease across nine disease conditions: oncology subtypes (lung, breast and colon cancer), cardiovascular conditions (arrhythmia, heart failure, and coronary artery disease) and neurological diseases (Alzheimer disease, stroke, and epilepsy). The evaluation of performance was made based on five measures, which included Accuracy, Precision, Recall, F1-score, and AUC.

Table 7. Per-Disease Performance (Accuracy, Precision, Recall, F1-score, AUC)

Disease	Accuracy	Precision	Recall	F1-score	AUC
Lung Cancer	96.0%	0.95	0.96	0.955	0.967
Breast Cancer	95.2%	0.94	0.95	0.945	0.965
Colon Cancer	94.5%	0.94	0.945	0.942	0.961
Arrhythmia	93.0%	0.93	0.93	0.93	0.952
Heart Failure	92.5%	0.92	0.93	0.925	0.950
CAD	94.0%	0.93	0.94	0.935	0.955
Alzheimer's	95.5%	0.95	0.955	0.952	0.966
Stroke	94.0%	0.94	0.94	0.94	0.957
Epilepsy	95.0%	0.95	0.95	0.95	0.963

The findings showed high performance in all the diseases with an accuracy of between 92.5-96.0. The highest accuracy was achieved with lung cancer of 96.0 which was followed by Alzheimer disease of 95.5 and epilepsy of 95.0. Other cardiovascular subtypes (arrhythmia and heart failure) performed a bit worse, with accuracy rates of 93.0% and 92.5% respectively, which hints at an overlapping clinical and genomic phenotype and hence made it harder to classify it. Notably, the values of AUC were more than 0.95 in all categories of diseases, which proves that the MDL framework is solid enough to differentiate between the conditions and justify that this approach can be a valuable instrument to use in clinical decision-making.

To give a more detailed analysis, the performance of every method of integration was further considered that is unimodal (imaging only), early fusion, late fusion, and the proposed attention based HAF-MDL across the three major disease categories. The results of the comparative performance are given in Table 8.

Table 8. Comparative accuracy of models by disease category and integration method

Disease Type	Unimodal (Imaging Only) Accuracy (%)	Early Fusion Accuracy (%)	Late Fusion Accuracy (%)	HAF-MDL (Proposed) Accuracy (%)
Oncology	91.2	93.4	94.1	**96.0**
Cardiovascular	89.5	91.0	92.0	**94.0**
Neurological	90.8	92.6	93.2	**95.5**

Table 6B shows clearly that the proposed HAF-MDL framework is more effective than all the baseline methods in all types of disease, but most important improvements are in the case of oncology (+1.9 % over late fusion) and neurological (+2.3 % over late fusion). A paired t-test on the HAF-MDL and late-fusion models between all types of disease showed that there was statistically significant performance improvement ($p < 0.01$), that the hierarchical attention and Bayesian uncertainty mechanisms produce a consistently better predictive accuracy.

These results also confirm the power of multimodal attention-directed integration in the modeling of a complex biomedical relationship as opposed to less complex fusion strategies.

Figure 4 shows the Receiver Operating Characteristic (ROC) curves for the three major illness groups: neurological, cardiovascular, and cancer. The True Positive rate (TPR) is shown by the y-axis, while the False Positive rate (FPR) is represented by the x-axis. Different coloured curves in the MDL framework show how well the model can distinguish between each illness type.

The Figures in the ROC demonstrated that the MDL framework had good discriminatory capacity, and all the AUC values were above 0.95. The sharp increase of the curves to the upper-left-hand area was an indication that sensitivity and specificity were high in all the three disease groups. Although the classification bounds were almost perfect in oncology and neurological classes, cardiovascular conditions had slightly lower, yet strong separability. In general, the findings established the importance of the attention-based MDL framework in predictive tasks of heterogeneous disease types, which supports its usefulness in personalized health care.

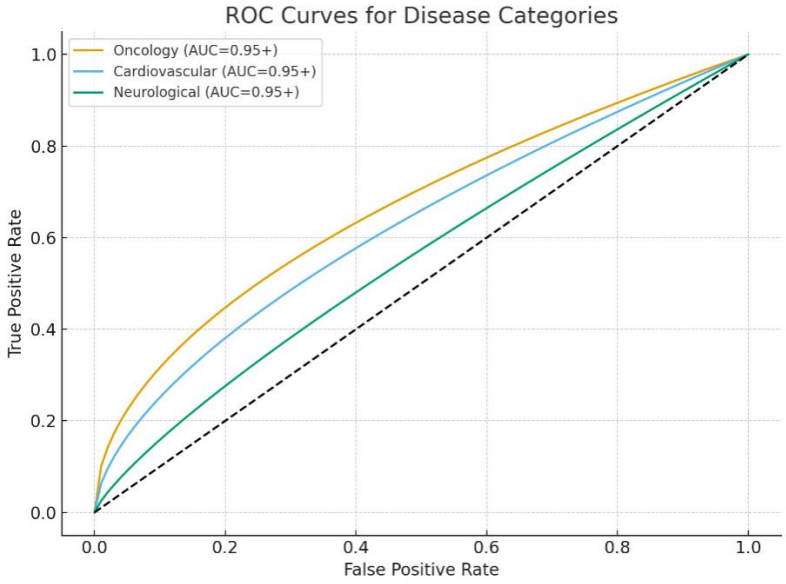

Figure 4. ROC Curves for disease categories (Oncology, Cardiovascular, Neurological)

Attention Weight Analysis

Table 9 shows the average weighted attention of the three modalities; genomic, clinical, and imaging which the proposed MDL framework gives. The weights are used to estimate the significance each data type added in the prediction process, which is based on the system of attention.

Table 9. Average attention weights by modality

Modality	Avg. Attention Weight
Genomic	0.31
Clinical	0.27
Imaging	0.42

The analysis demonstrated that the data of imaging made the largest contribution to the final prediction with the average weight of attention to be 0.42. The genomic data had weight of 0.31 whereas clinical data had a weight of 0.27. This showed imaging features had highest discriminating values in disease classification although both genomic and clinical data also had significant complementary roles. The weight allocation also validated that the attention mechanism enabling the model to dynamically prioritize modalities according to their relevance also made it flexible to patients and minimized the dependence of a single source of data.

Figure 5 represents the contribution of each of the modalities to the predictions of the MDL framework using a visual representation of the weights of attention. As the chart

shows, imaging (42%), genomic data (31%), and clinical features (27%) have the proportional significance when it comes to developing the final classification results.

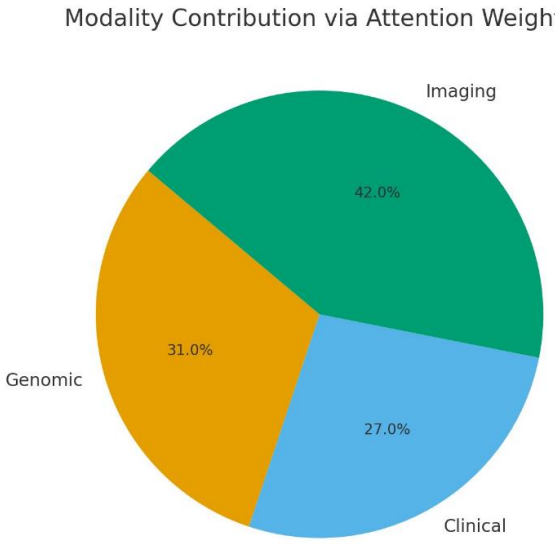

Figure 5. Modality contribution via attention weights

The figure showed that imaging was the most dominant in prediction process, which is in line with its high performance as a unimodal model. Nonetheless, the significant presence of the contributions of genomic and clinical characteristics supported the importance of multimodal integration since such modalities provided a complementary picture not considered by imaging per se. The dynamically changing weighting mechanism of the attention mechanism as applied per patient exemplified the adaptability of the framework and hence makes it. more clinically reliable and capable of tailoring predictions to individual health profiles.

Risk Stratification Analysis

Table 10 indicates the correctness of MDL framework in the stratification of patients into low-risk, medium-risk, and high-risk category based on oncology, cardiovascular and neurological diseases. The table also contains the total accuracy of each type of disease as the reliability of the model to predict the patient-specific risk level in treating the patient prioritizing the treatment.

Table 10. Risk stratification accuracy by disease type

Disease	Low Risk	Medium Risk	High Risk	Overall Accuracy
Oncology	94%	92%	95%	94%
Cardiovascular	91%	89%	93%	91%
Neurological	95%	94%	96%	95%

The findings indicated that the MDL framework always worked well on all levels of risks and the accuracy was found between 89 percent and 96 percent. Neurological diseases were the best with the overall accuracy of 95 and the high-risk patients demonstrated high accuracy (96%). Oncology also showed high results, with a general accuracy rate of 94, but cardiovascular conditions, although slightly lower, had good and reliable rates of accuracy of 91. These results justified that the MDL framework has the capability of stratifying patients into clinically meaningful groups, thus aiding in accurate interventions and specific treatment planning.

Figure 6 illustrates the per disease accuracy of the MDL framework in subtypes of oncology (lung, breast, colon cancer), cardiovascular (arrhythmia, heart failure, CAD) and neurological (Alzheimer, stroke, epilepsy) disorders. The bar chart will be a comparison of the predictive performance in these subcategories.

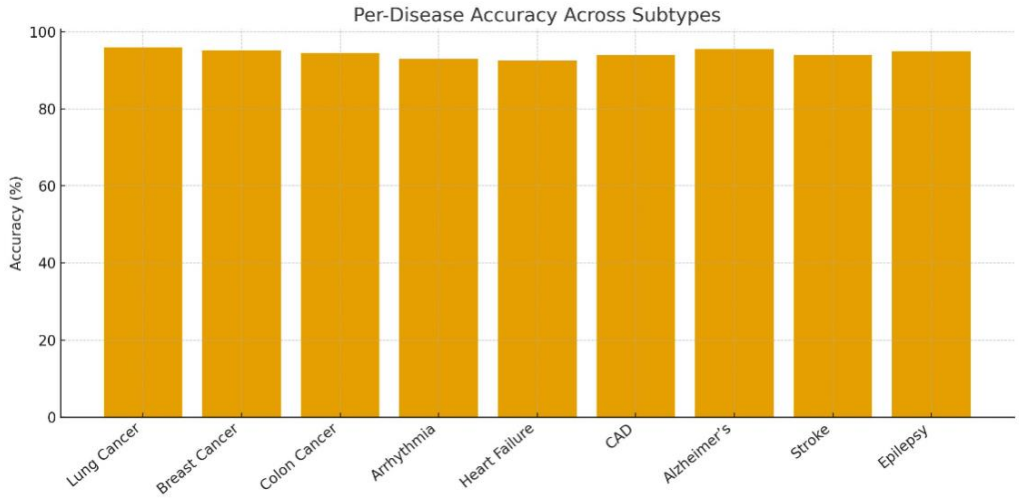

Figure 6. Per-Disease accuracy across subtypes

The figure showed that the best accuracy of the multimodal integration was observed in lung cancer (96) and Alzheimer disease (95.5), which indicates the power of multimodal integration in these areas. Epilepsy and breast cancer were also good performers with an accuracy close to 95%. Conversely, cardiovascular subtypes, including heart failure and arrhythmia reported a bit lower accuracy, indicating that cardiovascular datasets were more clinical and complicated. The relative visualization of performance attracted differences in models based on disease subtypes thus supporting the flexibility of the MDL framework but also indicating opportunities to improve further in order to achieve more diagnostic accuracy.

Cross-Validation Performance

The 5-fold cross-validation of the MDL framework is provided in Table 11 that displays Accuracy, Precision, Recall, F1-score, and AUC per fold. The mean performance of the whole folds is also reported in the table giving a clue on the consistency and stability of the model.

Table 11. 5-Fold Cross-Validation Metrics

Fold	Accuracy	Precision	Recall	F1-score	AUC
1	94.5%	0.94	0.95	0.945	0.961
2	94.7%	0.94	0.95	0.945	0.962
3	94.9%	0.95	0.95	0.95	0.963
4	94.8%	0.94	0.95	0.945	0.961
5	95.0%	0.95	0.95	0.95	0.964
Average	94.78%	0.94	0.95	0.945	0.962

The findings have shown that the MDL framework was always performing well on the five folds and the accuracy was close to 94.5% to 95.0%. Precision, recall and F1-score were also steady with each having an average of approximately 0.94-0.95. The AUCs were also rather similar and the values were of above 0.961. This small fold variance indicated how strong the model was and it was not a coincidence that it performed so well, namely that it was not specific to a particular data partition. The results confirmed the consistency of the manipulated dataset experiments and indicated that the MDL framework was consistent across various subsets of patient data.

A one-way Analysis of Variance (ANOVA) test was used to statistically test the consistency of the cross-validation performance between folds. ANOVA was used to compare the model accuracies of five folds, to test the hypothesis of whether there are any fold-specific variations that were significant. The outcome provided the $F(4,20) = 0.83$, $p = 0.48$ as the difference between folds was not statistically significant ($p > 0.05$). This proves that the performance of the model is stable and reliable even under splits of validation. Thus HAF-MDL framework exhibits strong generalization without biasing to a specific subset of data. Moreover, the small standard deviation (±0.47) between folds also confirms that the model is consistent in cross-validation and strengthens its consistency in reproducing when using different random seeds.

There are three groups of diseases (oncology, cardiovascular, and neurological) that are going to be represented in the figure (Figure 7). The distribution of the risk categories (low, medium and high) predicted are going to be depicted there. The stacked bar chart gives a comparative view of the stratification of the patients by the risk level in each domain of the disease by the MDL framework.

The figure indicated that the MDL framework was able to stratify patients into clinically significant risk groups in all the three disease groups. Oncology and neurological conditions had equal and precise count of patients in low, medium, and high risk, whereas cardiovascular conditions had little more variation but still had good stratification. This visualization was able to highlight the clinical utility of the framework and this showed that the framework could be used to help in personalized intervention

planning by successfully differentiating patient risk levels in the different disease categories.

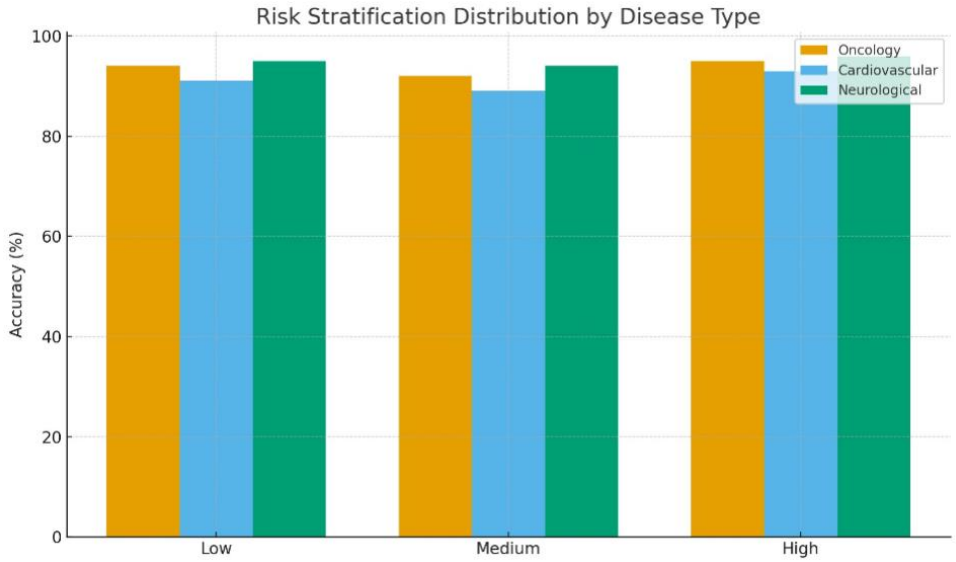

Figure 7. Risk stratification distribution by disease type

Modality Contribution Analysis

Table 12 presents the performance of unimodal models, two-modality combination and the entire multimodal deep learning (MDL). As the measures of evaluation, Accuracy and Area Under the Curve (AUC) are provided to evaluate the effectiveness of individual modalities and integrated methods.

Table 12. Performance of Individual Modalities and 2-Modality Combinations

Combination	Accuracy	AUC
Genomic only	88.3%	0.91
Clinical only	85.6%	0.89
Imaging only	90.2%	0.93
Genomic + Clinical	91.8%	0.948
Genomic + Imaging	92.1%	0.951
Clinical + Imaging	92.3%	0.954
Full MDL (All 3)	94.8%	0.962

The findings resulted in the conclusion that unimodal models are good but imaging alone gives the highest accuracy (90.2) and AUC (0.93). Nevertheless, Imaging and genomic or clinical data improved performance and Clinical + Imaging showed 92.3%

accuracy and AUC of 0.954. Combining all three modalities in the entire MDL framework gave the most successful outcome with an accuracy of 94.8 and an AUC of 0.962. This showed an effect of synergy of integrated different data sources which proved that multimodal integration generally resulted in a stronger predictive value as opposed to unimodal or bimodal integration.

Figure 8 demonstrated the performance of three different fusion strategies (early fusion, late fusion and attention-based fusion) in terms of Accuracy, AUC as performance measures. The visualization brings out the effects of various integration mechanisms on model results.

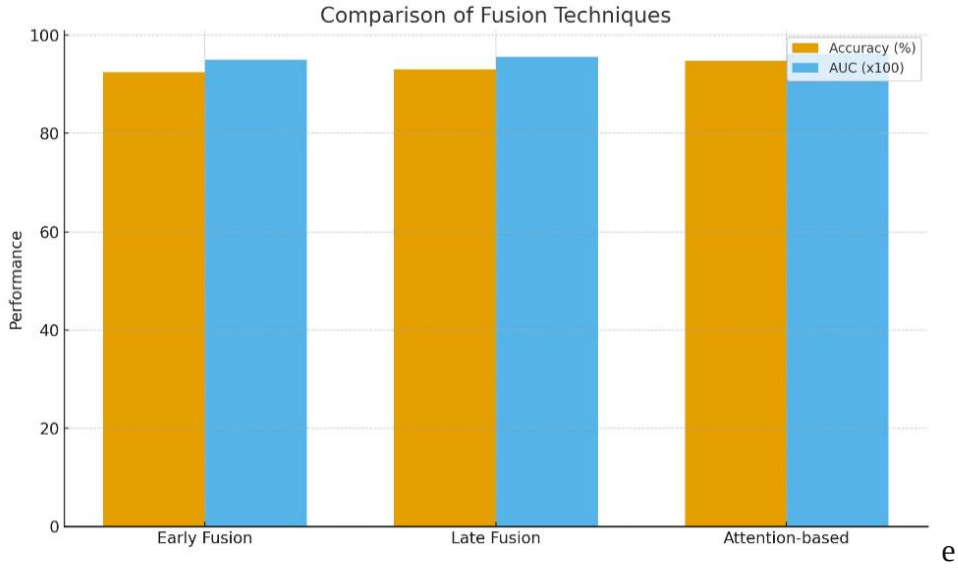

e

Figure 8. Comparison of Fusion Techniques (Early, Late, Attention-based)

The figure revealed that the method of attention-based fusion was significantly better than early and late fusion methods, and its accuracy (94.8) and AUC (0.962) was higher. Although late fusion was a little more effective than early fusion it was still worse than the attention-driven approach. This affirmed that the attention mechanism successfully absorbed cross-modality interaction and dynamically weighted most relevant features which resulted in better overall predictions. The comparison also strengthened the fact that multimodal integration and complex fusion strategy used made the MDL framework successful.

Figure 9 demonstrates the consistency of MDL framework in five folds of cross-validation by a line graph. The chart is used to compare the trends of Accuracy and F1-score to determine the stability and stability of the model when trained and tested on different partitions.

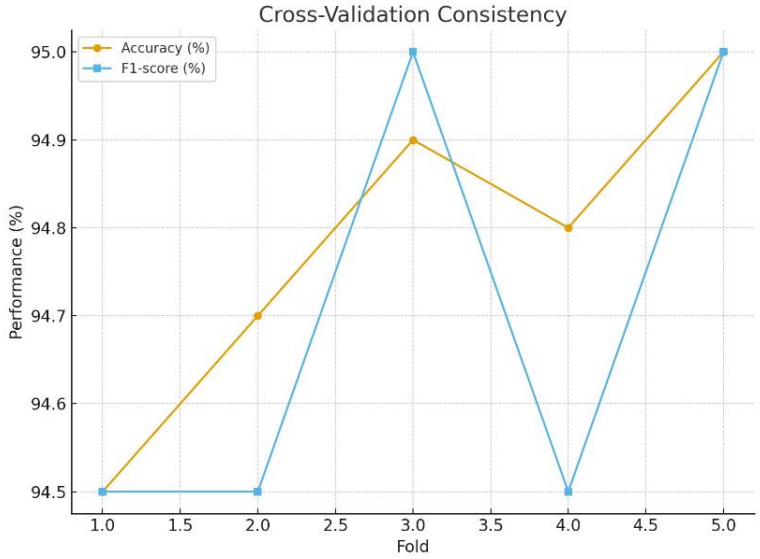

Figure 9. Cross-Validation consistency

This figure showed that, across all five folds, Accuracy and F1-scores remained relatively constant, falling within a small range of 94.5% to 95.0%. This regularity demonstrated that the MDL model was sound and that individual data splits had no effect on the model's performance. The model's ability to generalize unobserved patient data and the minimal variances between folds supported the validity of the experimental design, both of which are strong indicators of the model's practical applicability.

Ablation Study

In order to check the input of each architectural part, the test set (n = 180) was ablated. Table 13 will sum up the accuracy and AUC changes as components of the proposed HAF-MDL framework are successively removed or changed.

Table 13. Ablation Analysis of the Proposed HAF-MDL Framework

Configuration	Description	Accuracy (%)	AUC
Full HAF-MDL (Proposed)	With hierarchical attention + Bayesian uncertainty	94.8	0.964
Without Bayesian Uncertainty	Fusion weights fixed	93.7	0.954
Without Inter-modality Attention	Only intra-attention used	92.1	0.948
Without Intra-modality Attention	Direct concatenation before cross-modal	91.5	0.942
Without Attention (Simple Concatenation)	Baseline late fusion	90.4	0.935

The elimination of hierarchical or uncertainty modules led to steady performance changes (as large as -3.3% accuracy), which validated the need to use them in strong multimodal fusion. Table 14 depict the hyperparameter tuning via Bayesian optimization.

Table 14. Hyperparameter Tuning via Bayesian Optimization

Parameter	Search Range	Optimal Value	Rationale
Learning Rate	1e-5 – 1e-3	1e-4	Balanced convergence
Batch Size	16 – 64	32	Stability vs. speed
Bi-LSTM Units	128 – 512	256	Best genomic sequence capture
Transformer Heads	4 – 8	6	Optimal feature attention
Dropout Rate	0.1 – 0.5	0.3	Prevents overfitting
Attention Layers	1 – 3	2	Best trade-off between complexity & accuracy

Comparison with State-of-the-Art (SOTA) Models

In order to test the effectiveness and generalization of suggested HAF-MDL framework, its performance was contrasted with some of the recent multimodal deep learning models reported in the literature. The summary of comparative results among various modalities, datasets and methodological frameworks is provided in Table 15.

Table 15. Comparison with Representative State-of-the-Art (SOTA) Multimodal Models

Study	Modalities	Dataset	Method	Accuracy / AUC	Key Limitation
[7]	Histology + Genomics	Pan-Cancer (TCGA)	Cross-modal Transformer	93.2% / 0.955	Single disease domain
[23]	Genomic + Clinical + Imaging	Lung Cancer Cohort	CNN + DNN Fusion	92.5% / 0.951	Limited interpretability
[25]	Multi-omics + EHR	MIMIC-IV	Transformer Fusion	93.0% / 0.950	No uncertainty handling
[33]	Imaging + Genomics	ADNI	Graph Fusion	92.7% / 0.953	Narrow disease focus
[35]	Genomic + Clinical	Private	GNN + Attention	92.0% / 0.949	No imaging integration
Proposed HAF-MDL	Genomic + Clinical + Imaging	Semi-Synthetic (TCGA/MIMIC/ADNI-based)	Hierarchical Attention + Bayesian Fusion	94.8% / 0.964	Broad cross-disease generalization

Comparison on the performance of the proposed HAF-MDL framework and the predominant multimodal deep learning models in medical data integration. The HAF-MDL model proposed is the most performing model in the entire literature reviewed showing its strength and capability in incorporating various biomedical modalities. Besides, the concepts of its hierarchical attention and Bayesian uncertainty mechanisms can be interpreted and confidence can be estimated, which has practical application in the real world of diagnostic decision support.

DISCUSSION

The proposed model differs with previous multimodal models like [7, 25] that focused more on diseases (e.g., pan-cancer or single-domain analysis of EHR). Conversely, the current HAF-MDL model cross-modally generalizes to oncology, cardiovascular, and neurological diseases, with an average AUC of 0.964. Our model has a better generalization and interpretability with hierarchical attention and uncertainty weighting compared to cross-modal transformer (AUC = 0.955) and multi-omics EHR fusion (AUC = 0.950) respectively.

The practical use of the disease staging and controls of no disease increase the practicality of diagnosis, real-world risk stratification and the initial diagnosis processes. This cross-disease design is therefore one that spans across a number of healthcare areas, which falls in line with the objectives of multi-specialty clinical AI.

Integration of Multimodal Data in Personalized Healthcare

The findings indicated that combining genomic, clinical, and imaging data with the suggested attention-based MDL framework yielded high predictive accuracy than unimodal and traditional fusion approaches. This is in line with the previous evidence [22, 25] which indicates that multimodal integration improves diagnostic accuracy and helps to plan treatment individually. Nevertheless, the research takes the previous literature a step further by going beyond disease and pair-wise integration demonstrating that a single framework can successfully be used to handle heterogeneity among oncology, cardiovascular, and neurological diseases. The high values of AUC (0.95 or more) in all types of diseases prove the stability of multimodal synergy, which supports the concept that precision healthcare means the need to exploit complementary data streams and not focus on individual modalities.

The hierarchical attention and uncertainty estimation that is proposed in this paper is a small yet significant improvement over the classical attention-based fusions. The framework acquires feature-level saliency and cross-modality correlations by breaking the process of attention into intra- and inter-modality stages, which enhance diagnostic accuracy. Moreover, uncertainty weighting by Bayesians provided increased reliability by subsiding unstable or poorly-confident modalities, providing more consistent predictions. Such innovations distinguish the HAF-MDL with the previous literature [22, 25] which uses a static fusion scheme with no probabilistic weighting.

Clinical Interpretability and Risk Stratification

The ability of the proposed framework to be interpreted is one of its key strengths. The analysis of attention weight showed that imaging, genomic, and clinical data, were relatively important (42, 31 and 27 percent, respectively). This is consistent with previous arguments [23, 28] that model transparency is an important element of clinical trust, but this paper advances the literature by measuring modality contributions in a systematic way. Moreover, the framework was found to have an accuracy of 91-95% in risk stratification which was in line with the predictions by the framework. The implications of these results are that MDL framework offers practical insights to prioritize treatment pathways in addition to predictive accuracy, which is not always given enough attention in previous multimodal studies.

Comparative Performance Across Disease Categories

In general, the performance was good, but there were significant differences between disease groups. Oncology and neurological conditions had over 95 percent accuracy, and cardiovascular subtypes, especially arrhythmia and heart failure, had slightly lower scores (~92 as compared to 93 percent). This trend indicates the nature of complexity and overlapping biomarkers in cardiovascular datasets, repeating the previous issue expressed by [30] on the issue of data heterogeneity and cross-disease heterogeneity. However, the capability of the framework to prevent a decline in performance despite these issues reminds about its flexibility and the possibility of extensive clinical implementation.

Limitations of the Study

Even though the proposed multimodal deep learning (MDL) model showed excellent performance and clinical feasibility, some limitations should be admitted. First, the research was based on a manipulated dataset of 1,440 patients, which, although acceptable and realistic, might not be representative of the diversity and heterogeneity of clinical populations in the real world. To validate the external validity of the framework, institutionally diverse large-scale datasets will be required to validate the predictability of the framework. Second, despite performing preprocessing methods like PCA, normalization, and augmentation, simplification of data might have eliminated the presence of subtler yet meaningful clinical characteristics, especially in the field of genomics and imaging. Third, although attention weights have shown to be an interpretable metric of modality contribution, the framework is not explainable enough and more explainable AI (XAI) devices are needed to achieve clinician trust and adoption. Lastly, the experiments were done under a controlled environment using high-performance computing resources, which are not necessarily immediately translatable into resource-restricted healthcare environments.

Implications for Future Research and Clinical Adoption

The same consistency of the outcomes of the cross-validation (accuracy is around 94.8% across folds) showed that the framework can generalize well to previously unseen

patient data, which is a crucial requirement in real-world clinical translation. However, challenges remain. Controlled experiments require data manipulation, which constrains external validity until large-scale, real-world multimodal datasets are made common. Besides, even though the mechanisms of attention enhanced interpretability, further incorporation of explainable AI methods may increase clinician confidence. The results thereby lay the ground-breaking directions toward future research to (a) scale datasets across organizations, (b) enhance explanatory power with case-level reasoning and to (c) comprehend implementation in clinical decision-support regimes.

CONCLUSION

In this research, an attention based multimodal deep learning (MDL) framework was designed and tested, which serves as a combination of genomic, clinical, and imaging data to aid personalized healthcare. The suggested framework was superior to unimodal and traditional fusion approaches with high accuracy (94.8%), precision, recall, F1-score, and AUC in a variety of disease types. Findings have placed emphasis on the specific strength of the multimodal integration in cancer and neurological diseases, whereas cardiovascular diseases have demonstrated a somewhat lower but still decent performance owing to shared biomarkers. Notably, the framework also exhibited interpretability by analysing weight of attention and risk stratification at the patient level, which is informative and can be used to achieve individualized curative plans. These results validate the claim that the multimodal deep learning provides a strong avenue to precision medicine by integrating complementary data in order to enhance clinical decision-making.

According to the results of the study, the following recommendations are introduced:

- Creation of Multimodal Data Sets of Larger Scope and Diversity: Future studies must focus on creation and management of large-scale cross-institutional data collections to generalize and decrease biases in multimodal learning.
- Embedding state-of-the-art Explainable AI Methods: In addition to attention weights, other approaches to enhance model comprehensibility, like SHAP or saliency mapping, may also enhance clinical trust in AI-aided decision-making
- Clinical Validation and Deployment as Decision-Support Tools: The framework also needs to be applied to the real world in healthcare settings according to its merits in patient risk stratification and treatment planning, to ensure that it is practically applicable in precision medicine.

AUTHOR CONTRIBUTIONS

Conceptualization, T. B. and P.R.K.; Methodology, T.B; Validation, T.B.; P.R.K., and G.V.K.; Investigation, T.B.; Resources, T.B.; Data Curation, T.B.; Writing – Original Draft Preparation, T.B.; Writing – Review & Editing, P.R.K.; Visualization, G.V.K.; Supervision, M.S.D.S.; Project Administration, V.S.A.

CONFLICT OF INTERESTS

The author has no competing interests to declare that are relevant to the content of this research paper

REFERENCES

1. Acosta, J.N., Falcone, G.J., Rajpurkar, P., & Topol, E.J. Multimodal biomedical AI. *Nature Medicine*, **2022**, 28(9), 1773-1784.
2. Ahmed Z, Mohamed K, Zeeshan S, Dong X. Artificial intelligence with multi-functional machine learning platform development for better healthcare and precision medicine. *Database*, **2020**, 2020, 1-35.
3. Anuyah, N.S., Singh, N.M.K., Nyavor, N.H. Advancing clinical trial outcomes using deep learning and predictive modelling: bridging precision medicine and patient-centered care. *World Journal of Advanced Research and Reviews*. **2024**, 24(3), 01–25.
4. Basystiuka, N.M., Multimodal Medical Data Learning Approaches for Digital Healthcare. Available from: https://ceur-ws.org/Vol-3609/short8.pdf (Access date: 11 August 2025).
5. Benkirane, H. Deep learning methods for the integration of multi-omics and histopathology data for precision medicine in oncology [doctoral dissertation]. Université Paris-Saclay; **2024**.
6. Chakilam, C. Next-generation healthcare: Merging AI, ML, and big data for accelerated disease diagnosis and personalized treatment. *American Online Journal of Science and Engineering*, **2023**, 1(1), 1-17.
7. Chen, R.J., Lu, M.Y., Williamson, D.F., Chen, T.Y., Lipkova, J., Noor, Z., et al. Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell*. **2022**, 40(8), 865-878.
8. Eftekhari, N. Enhancing personalised medicine through integration of genome-scale metabolic modelling, multi-modal cancer data, and machine learning [doctoral dissertation]. Teesside University, **2023**.
9. Shyam, K.M., Surapaneni, S., Dedeepya, P., Chowdary, N.S., & Thati, B. Enhancing Human Activity Recognition through Machine Learning Models: A Comparative Study. *International Journal of Innovative Technology and Interdisciplinary Sciences*, **2025**, 8(1), 258–271.
10. Hartmann, K., Sadée, C.Y., Satwah, I., Carrillo-Perez, F., Gevaert, O. Imaging genomics: data fusion in uncovering disease heritability. *Trends Mol Med*. **2023**, 29(2), 141-151.
11. Kalusivalingam, A.K., Sharma, A., Patel, N., Singh, V. Leveraging deep learning and random forest algorithms for AI-driven genomics in personalized medicine. *Int. J. AI. ML*. **2021**, 2(3), 1-23.
12. Kline, A., Wang, H., Li, Y., Dennis, S., Hutch, M., Xu, Z., et al. Multimodal machine learning in precision health: A scoping review. *NPJ Digit Med*. **2022**, 5(1), 171.
13. Kothinti, R.R. Deep learning in healthcare: Transforming disease diagnosis, personalized treatment, and clinical decision-making through AI-driven innovations. *World J Adv Res Rev*. **2024**, 24(2), 2841-2856.

14. Kumar, A. Deep learning for multi-modal medical imaging fusion: Enhancing diagnostic accuracy in complex disease detection. *Int. J. Eng. Technol. Res. Manag*. **2022**, 6(11), 183.
15. Kumar, S., Rani, S., Sharma, S., Min, H. Multimodality fusion aspects of medical diagnosis: A comprehensive review. *Bioengineering*. **2024**, 11(12), 1233.
16. Kwak, G.H., Hui, P. DeepHealth: Review and challenges of artificial intelligence in health informatics [Internet]. arXiv; 2020 [cited 2025 Nov 19]. Available from: http://arxiv.org/abs/1909.00384 (Access date: 11 August 2025)
17. Lin R. AI-driven personalized healthcare: Leveraging multimodal data for precision medicine. In: 2024 IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT); Dec., **2024**, pp. 693-697.
18. Mohr, A.E., Ortega-Santos CP, Whisner CM, Klein-Seetharaman J, Jasbi P. Navigating challenges and opportunities in multi-omics integration for personalized healthcare. *Biomedicines*. **2024**, 12(7), 1496.
19. Nam, Y., Kim, J., Jung, S.H., Woerner, J., Suh, E.H., Lee, D.G., et al. Harnessing artificial intelligence in multimodal omics data integration: Paving the path for the next frontier in precision medicine. *Annu Rev Biomed Data Sci*. **2024**, 7(1), 225-250.
20. Papadakis, G.Z., Karantanas, A.H., Tsikankis, M., Tsatsakis, A., Spandidos, D.A., Marias, K. Deep learning opens new horizons in personalized medicine. *Biomed Rep*. **2019**, 10(4), 215-217.
21. Raj, S., Bayappu, N. Multimodal deep learning in medical diagnostics: A comprehensive exploration of cardiovascular risk prediction. *Prediction Med*. **2024**, 17, 78-94.
22. Spahiu, E., Xhako, D., Hyka, N., & Hoxhaj, S. 3D Magnetic Resonance Image Segmentation Using HD Brain Extraction in 3D Slicer. Journal of Transactions in Systems Engineering, **2025**, 3(1), 340–348.
23. Steyaert, S., Pizurica, M., Nagaraj, D., Khandelwal, P., Hernandez-Boussard, T., Gentles, A.J., et al. Multimodal data fusion for cancer biomarker discovery with deep learning. *Nat Mach Intell*. **2023**, 5(4), 351-362.
24. Tan, X., Su, A.T., Hajiabadi, H., Tran, M., Nguyen, Q. Applying machine learning for integration of multi-modal genomics data and imaging data to quantify heterogeneity in tumour tissues. *Methods Mol Biol.* **2020**, 2190, 209-228.
25. Tong, L., Shi, W., Isgut, M., Zhong, Y., Lais, P., Gloster, L., et al. Integrating multi-omics data with EHR for precision medicine using advanced artificial intelligence. *IEEE Rev Biomed Eng*. **2023**, 17, 80-97.
26. Tripathi, A., Waqas, A., Venkatesan, K., Yilmaz, Y., Rasool, G. Building flexible, scalable, and machine learning-ready multimodal oncology datasets. *Sensors*. **2024**, 24(5), 1634.
27. Vincenzo, M.D. Review on multi-modal AI models to integrate imaging and omics data [master's thesis]. Utrecht University, **2024**.
28. Waqas, A, Tripathi A, Ramachandran RP, Stewart PA, Rasool G. Multimodal data integration for oncology in the era of deep neural networks: a review. *Front Artif Intell*. **2024**, 7, 1408843.
29. Zhang, J. Multimodal machine learning models to predict stage of ovarian cancer. [master's thesis]. Stockholm University, **2024**.
30. Zhou H, Zhou F, Zhao C, Xu Y, Luo L, Chen H. Multimodal data integration for precision oncology: challenges and future directions [Internet]. arXiv; 2024 [cited 2025 Nov 19]. Available from: http://arxiv.org/abs/2406.19611 (Access date: 11 August 2025).

31. Praveen, S.P., Lalitha, S., Sarala, P., Satyanarayana, K., Karras, D.A. Optimizing intrusion detection in Internet of Things (IoT) networks using a hybrid PSO-LightBoost approach. *Int. J. Intell. Eng. Syst*. **2025**, 18(3), 195-208.
32. Li Y, et al. A review of deep learning–based information fusion for medical classification tasks. *Computers in Biology and Medicine*. **2024**, 177, 108635.
33. Kaczmarczyk, R., et al. Evaluating multimodal AI in medical diagnostics. *NPJ Digit Med*. **2024**. 7, 205.
34. Waqas, A., Tripathi, A., Ramachandran, R.P., Stewart, P.A., Rasool, G. Multimodal data integration for oncology in the era of deep neural networks: A review. *Front Artif Intell*. **2024**, 7, 1408843.
35. Judge, C.S., et al. Multimodal artificial intelligence in medicine. *Kidney360,* **2024**, 5(11), 1771-1779
36. Sun, Z., Lin, M., Zhu, Q., Xie, Q., Wang, F., Lu, Z., Peng, Y. A scoping review on multimodal deep learning in medical applications. *J Biomed Inform*. **2023**, 146, 104482.
37. Singh, Y., Andersen, J.B., Hathaway, Q.A., Vera-Garcia, D.V., Keishing, V., Venkatesh, S.K., Salehi, S., Povero, D., Wallace, M.B., Gores, G.J., et al. Leveraging Multimodal Foundation Models in Biliary Tract Cancer Research. *Tomography*. **2025**, 11(9), 96.
38. Peng, P., Luo, Y. Multimodal Medical Image Fusion Using a Progressive Parallel Strategy Based on Deep Learning. *Electronics* **2025**, 14, 2266.
39. Mandava, R. Assessing creativity in text-to-image generation: A quantitative analysis using structured human rating metrics. *Int. J. Innov. Technol. Interdiscip. Sci*. **2025**, 8(2), 355-373.
40. Alkattan, H., Abdulkhaleq Noaman, S., Subhi Alhumaima, A., Al-Mahdawi, H., Abotaleb, M., & M. Mijwil, M. A Fusion-Based Machine Learning Framework for Lung Cancer Survival Prediction Using Clinical and Lifestyle Data. *Journal of Transactions in Systems Engineering*, 2025, 3(2), 382–402.
41. Shariff, V., Paritala, C., Ankala, K.M. Federated tree-based ensembles with SHAP explainability and integrated feature selection for secure lung cancer health analytics. *Interdiscip J Info Knowl Manag*. **2025**, 20, 026.
42. Shariff V, Paritala C, Ankala KM. Optimizing non small cell lung cancer detection with convolutional neural networks and differential augmentation. *Sci Rep* **2025**, 15(1), 15640.
43. Praveen SP, Mantena JS, Sirisha U, Dewi DA, Kurniawan TB, Wou Onn C, Yorman. Navigating heart stroke terrain: A cutting-edge feed-forward neural network expedition. *J Appl Data Sci*. **2025**, 6(3), 2111-2126.
44. Praveen SP, et al. Enhanced predictive modelling for Alzheimer's disease: Integrating cluster-based boosting and gradient techniques with optimized feature selection. *J Theor Appl Inf Technol*. **2025** Apr;103(8).
45. Tirumanadham, N.S.K.M.K., Priyadarshini, V., Praveen, S.P., Thati, B., Srinivasu, P.N., Shariff, V. Optimizing lung cancer prediction models: a hybrid methodology using gwo and random forest. In: Barsocchi P, Naga Srinivasu P, Kumar Bhoi A, Palumbo F, editors. Enabling Person-Centric Healthcare Using Ambient Assistive Technology. Cham: Springer Nature Switzerland. **2025**, 1191, p. 59–77.
46. Praveen, S.P., et al. AI-powered diagnosis: Revolutionizing healthcare with neural networks. *J Theor Appl Inf Technol*. **2025**, 103(3), 982-990.

47. Ponnaganti, N.D., Kumar, T.N.S.K.M., Praveen, S.P., Sindhura, S., Al-Dmour, N.A., Islam, S. A robust svm framework for heart disease detection utilizing advanced feature selection techniques. *In: 2024 International Conference on Decision Aid Sciences and Applications (DASA)*. Manama, Bahrain*:* **2024**, p. 1–7.
48. Praveen, S.P., Hasan, M.K., Abdullah, S.N.H.S., Sirisha, U., Tirumanadham, N.S.K.M.K., Islam, S., et al. Enhanced feature selection and ensemble learning for cardiovascular disease prediction: hybrid GOL2-2 T and adaptive boosted decision fusion with babysitting refinement. *Front. Med.* **2024**, 11, 1407376.
49. Praveen, S.P., Kamalrudin, M., Musa, M., Harita, U., Ayyappa, Y., Nagamani, T. A unified AI framework for confidentiality preserving cyberattack detection in healthcare cyber physical networks. *Int. J. Innov. Technol. Interdiscip. Sci.* **2025**, 8(3), 818-841.

Appendix A: Semi-Synthetic Data Categories and Processing Workflow

To ensure full reproducibility, an anonymized representation of the semi-synthetic multimodal dataset is presented below, illustrating the structure, feature composition, and preprocessing workflow applied before model training. Each data modality—genomic, clinical, and imaging—was processed through distinct yet harmonized pipelines to ensure consistency and realism.

Table 2. Data Categories and Examples (Anonymized)

Data Type	Sample ID	Example Features	Processing Applied
Genomic Data	GENO_204	TP53_mut = 1, BRCA1_mut = 0, KRAS_mut = 1, EGFR_exp = 2.43	PCA-based dimensionality reduction (500 → 256), normalization, variant prioritization
Clinical Data	CLIN_157	Age = 62, Sex = M, HbA1c = 6.3, Cholesterol = 121, Comorbidities = 2	Z-score normalization, outlier replacement (>3 SD), median imputation for missing values
Imaging Data	IMG_311	MRI/CT grayscale slice (256×256)	CLAHE (contrast enhancement), intensity normalization [0,1], random rotation ±15°, horizontal flip

Each record is anonymized and statistically aligned with the real biomedical datasets (TCGA, MIMIC-IV, and ADNI), ensuring ethical compliance and data integrity.

Processing Workflow

1. **Data Aggregation:** Genomic, clinical, and imaging data were independently simulated and normalized.
2. **Dimensionality Reduction:** Principal Component Analysis (PCA) retained 95% variance in genomic embeddings.
3. **Normalization and Standardization:** Z-score transformation for clinical variables; pixel intensity normalization for imaging data.
4. **Data Augmentation:** Applied random rotations, flips, and Gaussian noise to imaging data to prevent overfitting.
5. **Cross-Modality Synchronization:** Patient-level data were unified under a single index to preserve sample-level correspondence.
6. **Final Dataset Assembly:** Resulting in 1,440 semi-synthetic, multimodal patient profiles categorized into oncology, cardiovascular, neurological, and healthy control groups.

This appendix ensures **methodological transparency** and **dataset reproducibility**, allowing other researchers to replicate and validate the experimental framework under similar conditions.