RESEARCH PAPER

Cardio-MTL: Multi-Task Cardiovascular Disease Prediction Using Large-Scale Multi-Disease Health Records

Mrs. R. Auxilia Anitha Mary[1], Dr. T. Ramaprabha[2]
[1]Research Scholar, Department of Computer Science, Nehru Arts and Science College, Coimbatore
Email: auxiliaanithamary@gmail.com
[2]Associate Professor, Department of Computer Science, Nehru Arts and Science College, Coimbatore
Email: nascramaprabhacs@nehrucolleges.com

Received: 20th Apr, 2026 | Revised: 25th Apr, 2026 | Accepted: 9th May, 2026 | Available Online: 14th May, 2026

ABSTRACT
The rapid expansion of national health surveillance systems, cloud-integrated medical repositories, and AI-driven chronic disease monitoring platforms has created unprecedented opportunities for population-scale cardiovascular analytics. However, existing predictive systems rely heavily on single-disease modeling, static feature pipelines, and imbalanced learning strategies that fail to capture cross-disease interactions, leading to reduced accuracy, poor minority-disease sensitivity, and limited clinical interpretability. To overcome these constraints, this work proposes CARDIO-MTL, a unified multi-task learning framework designed to jointly predict six major cardiovascular and metabolic conditions using large-scale BRFSS health records. The architecture integrates four complementary intelligence layers: UNI-ENCODE, which constructs survey-aware shared–private latent representations to stabilize heterogeneous inputs; HYBRID-MAP, an adaptive expert-routing engine combining deep neural, gradient-boosted, and Bayesian predictors; RISK-BAL, a prevalence-aware optimization layer addressing class imbalance and gradient conflict; and EXPLAIN-MED, a clinically grounded interpretability module generating attribution scores, comorbidity graphs, disease progression trajectories, and intervention priorities. Evaluated on more than 400,000 population samples, CARDIO-MTL achieves AUC values up to 0.98, average precision above 0.94, and consistent gains of 8–20% over single-task and traditional machine learning baselines. Ablation results confirm the essential contribution of each module, while confusion matrix analysis demonstrates low false-negative rates for life-critical diseases. CARDIO-MTL provides a scalable, interpretable, and population-ready intelligence layer for next-generation cardiovascular risk assessment.

Keywords: Multi-Task Learning, Class Imbalance Mitigation, Cardiovascular Disease, Health Records, and Multi-Disease Prediction.
How to cite this article: Auxilia Anitha Mary R, Ramaprabha T., Cardio-MTL: Multi-Task Cardiovascular Disease Prediction Using Large-Scale Multi-Disease Health Records. Int J Drug Deliv Technol. 2026;16(45s): 99-114; DOI: 10.25258/ijddt.16.45s.10

1. Introduction

The rapid growth of digital public health systems, cloud-based medical databases, nationwide screening systems, and AI-powered chronic-illness surveillance systems has fundamentally changed the manner in which medical information is created, handled, and interpreted in modern healthcare systems. Mass scale health surveys, electronic population registries, telehealth environments, and wearable biosensing monitors constantly collect sensitive health data (such as lifestyle behaviour, metabolic parameters, cardiovascular parameters, and the symptoms of chronic diseases) and, as a result, are a rich source of advanced clinical analytics. With the maturity of these technologies, cardiovascular and metabolic disorder predictive models have become the pillars of precision medicine efforts, as they enable earlier interventions, better risk stratification, and preemptive patient care. However, even with access to the massive multi-disease datasets, the current predictive systems are often solitary, single-disease, and face the challenge of different population demographics, intricate lifestyle trends, and comorbid disease interactions on a large population scale. These drawbacks can limit the reliability of diagnosis and hinder the adoption of AI-driven devices within the field of real medicine without any issues.

In the broader context, there is a much more focused and increasingly urgent challenge: the design of integrated and multi-disease prediction frameworks that are capable of joint analysis of the cardiovascular and metabolic conditions, such as heart disease, stroke, hypertension, diabetes, and renal dysfunction, all of which are interrelated. Population-scale data sets comprising hybrid cloud-stored electronic health records and extensive survey-derived repositories show that these disorders rarely exist in isolation. Rather, they have overlapping physiological pathways, shared risk factors, and cascading effects that make clinical decision-making substantially more difficult. Traditional learning paradigms have a poor fit to these interdependencies because their disease-specific pipelines of training do not consider cross-condition effects, unbalanced prevalence, and subtle interactions between demographic, behavioral, and clinical factors. Additionally, most current AI-based medical predictors are impaired in terms of interpretability, giving doctors risk scores without meaningful explanations and actionable data on disease progression.

The recent research highlights the role of multi-task learning and integrated health analytics as the means of overcoming the limitations of the independent cardiovascular risk models. However, existing multi-disease models are still constrained by the inherent design of features, lack of cross-disease learning to representation, task optimization imbalance, and lack of explanatory power. Notably, there is a gap in the literature that unites adaptive algorithm selection, dynamic risk-balancing plans, or human comprehensible multi-disease intelligence outputs, and thus can afford significant gaps in addressing the full cardiovascular risk management.

To fill these gaps, the current paper proposes CARDIO-MTL, a new generation of multi-task learning models that can be used to concurrently predict various cardiovascular and metabolic diseases based on the large-scale Behavioral Risk Factor Surveillance System (BRFSS). The suggested methodology will combine single multi-disease representation learning and hybrid algorithmic prediction pathways, adaptive risk balancing, and clinically meaningful explainability modules. CARDIO-MTL provides the future-proof cardiovascular healthcare industry with a powerful, interpretable, and clinically implementable answer through generating complete outputs, such as disease progression curves, inter-disease interaction networks, prioritized intervention lists, and customized lifestyle and clinical synergy programs.

2. Literature Review

Chai et al., [21] The current study is to develop a unified computational model that can assess physical chronic diseases and depressive conditions simultaneously by the use of data collected by wearable sensors. The methodological equipment includes a baseline multi-task learning model and an improved ADH-MTL approach that combines group-level modeling, a decomposition strategy, and a Bayesian network to address two sources of heterogeneity between disease entities and individual patients. According to empirical findings, ADH-MTL outperforms traditional baseline techniques, and methodological innovations continue to provide large gains in predictive validity and deliver a formidable foundation to the joint management of chronic illnesses and mental well-being.

Ding et al., [22] The current research aims at coming up with a single foundation model that can be used to solve a wide range of brain imaging applications, such as multi-disease diagnosis, as well as predicting the age of the brain. In the methodology, the DenseFormer-MoE architecture is presented that combines DenseNet, a Vision Transformer, and a Mixture of Experts module. This composite model is trained by Masked Auto-Encoding and self-supervised learning, thus facilitating better feature generalization and automatically allocating experts to reduce multi-task conflicts. Empirical results show the proposed model has been shown to be of strong performance in the UK Biobank (UKB), the Alzheimer's Disease Neuroimaging Initiative (ADNI), and the Parkinson Progress Markers Initiative (PPMI) databases, outperforming currently available methods on both disease diagnosis and age prediction tasks.

Li et al., [23] The research aims at summarizing the most recent developments in deep learning-based fundus image analysis in the context of hypertension-related pathologies, with the focus on the evaluation of hypertensive retinopathy and the diagnosis of ocular diseases, as well as cardiovascular risk prediction. Its methodology involves a systematic review of deep-learning methods that are used in vessel segmentation, artery-vein classification, lesion detection, and prediction of systemic disease using color fundus photography. The results indicate that deep learning can reproduce the retinal vascular and pathological features and thus promote timely screening and risk stratification; however, issues of generalizability, stability to image quality differences, and clinical interpretability remain to be addressed, which can guide future research.

Singh et al., [24] The current research project aims at developing an information technology-enriched Hybrid Statistical-Fuzzy (Hybrid S-Fuzzy) model that would be able to predict multifactorial pathologies, such as cardiopathies, hepatic disease, and nephropathies, along with the generation of personalized treatment advice. The analytical procedure combines traditional statistical predictive models with a fuzzy logic-based recommendation system, which evaluates the severity of the disease, co-occurring risk factors, and likelihood of occurrence based on a medically annotated clinician-created dataset. The available empirical evidence shows that the developed model achieves an accuracy of 96.5 percent, thus providing accurate diagnostic reports and tailored recommendations, and thus implicating its high potential in inclusion in e-health decision-support systems and more general clinical informatics settings.

Rufus et al., [25] The current study attempts to develop a machine-learning-based multiclass disease classification platform with the potential to identify seven key pathological diseases, such as diabetes mellitus, coronary artery disease, breast carcinoma, chronic kidney disease, hepatic disease, malaria parasitism, and community-acquired pneumonia. The adopted methodology combines a large amount of clinical data, advanced feature-engineering steps, and algorithmic customizations to achieve high diagnostic accuracy that is supplemented with an advisory module that provides condition-specific preventative advice and therapeutic suggestions. The experimental findings prove that the combination of accurate prognostication and practical advice makes it easier to manage disease comprehensively and proactively, which supports the prospects of artificial intelligence in preventative and customized medicine.

Reddy et al., [26] This paper aims to come up with a decipherable, integrative multi-task CXR-MultiTaskNet to classify thoracic disorders simultaneously and localize lesions in chest X-rays. Its methodology involves using a ResNet-50-based backbone to extract the features, two task-specific prediction heads, a hierarchical feature extraction module, which has two attentions, and a Grad-Cam-based explainability component. The model is trained by a combined loss that deals with both the issue of class imbalance and lesion heterogeneity. The results are shown to work better, with a macro F1 score of 0.965 and a mean intersection and over union of 0.851, which outperforms the current baseline with clinically meaningful and interpretable results as well as scalable diagnostic support.

3. Proposed Methodology

CARDIO-MTL is proposed as a multilayered and clinically grounded multi-task learning system that simultaneously forecasts numerous cardiovascular and metabolic diseases with the help of large-scale BRFSS population data. The architecture first transforms heterogeneous behavioral, demographic, survey, and clinical inputs into a consistent latent representation through UNI-ENCODE and alleviates the inconsistencies caused by the noisy, incomplete self-report information. Latent features are then input into HYBRID-MAP, which is an adaptive module that allocates the most suitable inference models, neural networks, gradient-boosted trees, or probabilistic learners to each disease-specific task based on the structure of the input features and task complexity. To achieve stable joint learning, RISKBAL uses disease-prevalence-sensitive weighting, gradient-conflict reduction, and loss balancing dynamically, so that no specific task can overshadow less prevalent ones. EXPLAIN-MED provides clinically interpretable results in the form of attribution maps, cross-disease interaction diagrams, individualized disease-progression paths, and recommendations incorporating both lifestyle and clinical information. The end-to-end pipeline deals with class imbalance, inter-disease associations, and heterogeneity of patients and preserves the amount of transparency needed to implement it clinically. CARDIO-MTL is operationally scalable to large BRFSS cohorts and provides robust, multifaceted intelligence, which includes prediction, explanation, and intervention advice, and thus represents a scalable and practical decision support system to assess and manage cardiovascular risks at the population level.

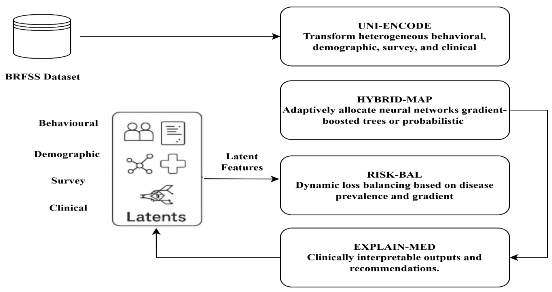

Figure 1: Architecture of the CARDIO-MTL Unified Multi-Disease Learning Framework

Figure 1 illustrates the CARDIO-MTL integrated multi-disease learning environment. The BRFSS population-scale data are initially preprocessed using UNI-ENCODE, which converts heterogeneous behavioral, demographic, survey, and clinical variables into a latent representation that is standardized using UNI-ENCODE. These latent inferences are then integrated into HYBRID-MAP, which is an element that is able to dynamically choose the best predictive expert neural networks, gradient-boosted tree models, or probabilistic schemes task by task on a disease-by-disease basis. This is followed by feeding the subsequent predictions to RISK-BAL, a dynamic balancing unit that alters weights of the losses as per disease prevalence and gradient dynamics, hence equitable learning based on common and rare conditions. Lastly, EXPLAIN-MED has the ability to produce clinically interpretable results, including explanatory narratives and actionable recommendations, which increases transparency and gives actionable insights to cardiovascular and metabolic risk assessment.

3.1 Data Collection and Pre-processing

The BRFSS data (Kaggle CDC release) contains more than 400,000 records with self-reported labels of diseases as well as behavioral, demographic, and limited clinical measures. Preprocessing includes removal of invalid data points, filling values of missing data points with conditional multiple imputation using chained equations, and normalization of continuous data points by z-scoring. Ordinal encoding or one-hot expansion is used to encode categorical variables depending on their cardinality. The loss function keeps sampling-design weights and uses them as weighting factors. Stratified SMOTE is used to reduce class imbalance in the case of minority disease classes, and sample reweighting is done when training. The mutual-information filtering with L1 regularization is used as the feature selection method. This results in balanced, survey-aware, and normalized feature matrices, X, and label matrices, Y, thus training multi-tasks.

3.2 UNI-ENCODE: Unified Multi-Disease Representation Encoder

Under UNI-ENCODE, a single latent representation of shared as well as disease-specific risk signatures is built on the basis of heterogeneous features of BRFSS. The model solves mixed data types, such as categorical, ordinal, and continuous data, and reduces measurement error due to self-reports and corrects latent confounding factors, such as socioeconomic status and access to healthcare. It has an encoding design that consists of a hybrid embedding strategy: categorical and ordinal fields are encoded using learned embeddings, continuous variables are encoded using batch-normalized nonlinear layers, whereas domain-specific priors (age, sex, and smoking) are encoded using residual pathways. Besides this, UNI-ENCODE is a low-rank disease-interaction subspace learner, which learns co-occurrence patterns and conditions a conditional attention head that conditions shared information per downstream disease task. More importantly, the model is survey-weight conscious, whereby the latent space is responsive to the population sampling biases, thus facilitating generalization across demographic layers in addition to maintaining interpretable axes to the EXPLAIN-MED.

3.2.1 Heterogeneous Feature Embedding

This is realized by combining learned categorical vectors with normalized continuous projections in order to transform mixed BRFSS inputs into continuous embeddings and form a coherent and trainable input space that is amenable to multi-task representation learning.

$$E = Concat\left(\{e_c\}_{c\epsilon C}, \phi(X_{cont})\right) \tag{1}$$

Here, e_c are learned embeddings for categorical fields C, ϕ is a nonlinear projection for continuous inputs X_{cont}, and concatenation produces the unified embedding E that feeds subsequent encoder layers.

3.2.2 Shared-Private Low-Rank Decomposition

In this framework, the latent space is broken down into components shared by tasks and disease-specific personal components, thereby helping to reduce the interference of parameters and encourage transfer of common risk signals.

$$Z = USV^T + P \quad (2)$$

Latent Z is represented as low-rank factorization USV^T(shared structure) plus private residual P. This separation enforces compact shared representations while preserving disease-specific nuances.

3.2.3 Disease-Conditioned Attention Modulation

A conditional attention head is used to produce task-specific masks that enhance the latent features that are the most informative regarding the disease of interest, and at the same time suppress other irrelevant features.

$$A_i = \sigma(W_a^i Z + b_a^i), \qquad Z_i = A_i \odot Z \quad (3)$$

For disease i, attention A_i is a sigmoid of linear projection; elementwise product Z_i applies the mask to shared Z, yielding disease-conditioned representation Z_i for downstream predictors.

UNI-ENCODE produces both task-specific and common multi-disease structure-sensitive compact survey-aware latent vectors. It has heterogeneous embedding, low-rank shared/ private decomposition, and disease-conditioned attention mechanisms to stabilize the multi-task transfer learning to reduce interference and generate robust input representations in the adaptive prediction engines of the HYBRID-MAP.

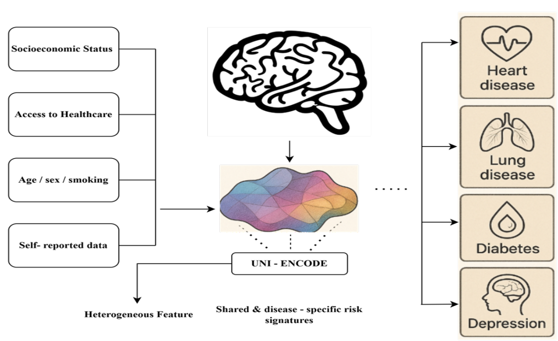

Figure 2: UNI-ENCODE Latent Representation Extraction Pipeline

Figure 2 shows how the UNI-ENCODE model combines a variety of health-related determinants to produce both shared and disease-specific risk signatures. The heterogeneous inputs are at the left end of the model, including socioeconomic status, access to healthcare, demographic factors like age, sex, and smoking status, and self-reported data. These variables are then mapped into a single latent space, as a multicolored abstract form which embodies complex patterns of varied data modalities. Brain refers to the learning and embedding processes in the model, which reduce the confounding factors and errors in measurement. These latent representations are then converted to fine risk signatures by the UNI-ENCODE engine. These outputs on the right-hand side make predictions in several diseases, such as cardiovascular disease, chronic lung disease, diabetes mellitus, and depression, pointing to the wide range of pathologies that which this model can be used.

3.3 HYBRID-MAP: Hybrid Multi-Algorithm Prediction Engine

HYBRID-MAP is a multi-expert, multi-adaptive algorithm-selection engine, which increases the reliability, robustness, and clinical application of multi-disease prediction in CARDIO-MTL. Integrating the unique complexity and risk signature of every cardiovascular or metabolic disorder into the most acceptable predictive mechanism is its fundamental principle, and it does not enforce a single model on all diseases. The dependence of features on cardiovascular activities differs significantly: hypertension has highly threshold-dependent patterns, diabetes has nonlinear metabolic patterns, kidney disease has noisy behavioral patterns, and stroke risk has a high variance based on composite clinical and demographic patterns. HYBRID-MAP mitigates these modeling gaps by making the three modeling paradigms of deep neural networks, gradient-boosted decision ensembles, and Bayesian calibration engines a single multi-expert system. Another routing controller is a differentiable routing controller, which considers patient-specific latent representations, disease-level variability, historical calibration error, multi-task interference indicators, and then allocates soft routing weights to the experts. This dynamic allocation system addresses the instability of discrete expert switching, and specialization is retained. The higher-order correlations are represented by neural subnetworks, tree-based models are used to provide interpretable rule-based boundaries, and the Bayesian components are used to generate calibrated probabilities, which are necessary in clinical interpretation. Experts are trained simultaneously and affect each other in terms of common encoder gradients, and selective regularization makes sure that convergence and stabilization. The combination of these complementary paradigms, in HYBRID-MAP, creates more clinically applicable, statistically robust predictions in a wide range of diseases and non-homogeneous groups of patients.

3.3.1 Differentiable Routing Controller

The process trains patient-attracted soft assignments in selecting experts through analyzing disease-specific latent data, easy gradient flow, and more flexible, and the ability of adaptive model specialization, avoiding brittle, discrete switching.

$$\pi_{i,k} = \frac{exp(r_k(Z_i)/\tau)}{\sum_{k'} exp(r_{k'}(Z_i)/\tau)} \quad (4)$$

The routing weight $\pi_{i,k}$ is calculated by scheduling a temperature-regulated softmax on the routing scores $r_k(Z_i)$. This model allows the allocation of experts that are fully differentiable and at the same time conserves the routing that is individualized by latent cardiovascular-metabolic phenotypes.

3.3.2 Neural Deep-Risk Expert

A multi-layer feedforward neural subnetwork or multi-layer residual neural subnetwork represents the nonlinear interplay of demographic, behavioral, and physiological factors, producing expressive risk signatures of diseases with complex multifactorial patterns.

$$\hat{Y}NN,i = MLP_{\phi}(Z_i) \quad (5)$$

This mapping, using a deep multilayer perceptron, changes the latent representation of the latent version of Z_i, and thus helps extract nonlinear and high-order disease indicators. The final model, which is denoted as the $\hat{Y}NN,i$, produces logits, hence capturing complex cardiovascular-metabolic relations in patients.

3.3.3 Gradient-Boosted Tree Expert

A combination of shallow trees can identify threshold-related patterns, monotonous clinical predictors, and interpretable split-based rules, making it especially applicable to diseases whose relationships are structured.

$$\hat{Y}GBM,i = \sum_{m=1}^{M} \gamma_m h_m(Z_i) \quad (6)$$

The gradient-boosting model combines the forecasts of weak learners, which are denoted by h_m, and weighted by a factor γ_m. It is an architecture that represents local decision rule modeling, and it is used to complement neural network techniques in that localized clinical decision boundaries are determined.

3.3.4 Bayesian Calibration and Uncertainty Expert

The Bayesian inference module imperceptive the posterior predictive probability distributions and thus, the noise and uncertainty of self-reported BRFSS data are considered, and at the same time, it helps in risk communication.

$$p(y|Z_i) = \int p(y|Z_i,\theta)\, p(\theta|D) d\theta \quad (7)$$

The posterior predictive model highlights calibrated risk estimates by incorporating over-parameter uncertainty, and thus, clinicians can interpret predictions over explicit confidence limits and uncertain decision support.

HYBRID-MAP combines neural, tree-based, and Bayesian professionals using differentiable routing and, as a result, provides robust and calibrated predictions of a range of heterogeneous diseases. Its flexible specialization, multispecialized nature, and multi-paradigm synergy all add to the aspect of predictive stability, reduce model bias, and produce clinically meaningful outputs that can be used in high-stakes medical decision-making.

3.4 RISK-BAL: Adaptive Multi-Disease Risk Balancing Layer

RISK-BAL module is designed to moderate the multi-task maximizing dynamics of CARDIO-MTL that, in fact, assures the most common conditions will not be overshadowed by less common, but nevertheless, clinically meaningful, diseases like chronic kidney disease or stroke. Multi-disease learning is distinguished by unequal contributions to the gradient; tasks with large samples in excessively large numbers affect the parameter updates, and cause other tasks to underfit, degenerate, or be discarded in training. Additionally, there is an interference of the tasks in CARDIO-MTL gradients obtained as a result of different pathologies may have opposing forces on a common encoder, thus disabling generalisation. RISK -BAL is a three-part plan that includes prevalence-based loss weighting, gradient conflict reduction, and curriculum-based minority amplification. All these mechanisms enable the steady control over the per-task calibration error, the distribution of classes, the alignment of the gradients and the optimization trajectory, and the calculation of the dynamic balancing factors that adjust the influences of each task. In contrast to fixed class-weighting schemes, RISK-BAL uses dynamically modulated modulation during training; only infrequent examples of the common encoder being stable are given increased learning pressure, thus preventing premature overfitting to stochastic minorities. Also, the weights of the BRFSS survey are integrated into the module, so the optimization process can incorporate population-wide patterns of the disease, and not just the count of the raw samples. As a result, RISK-BAL obtains a balanced, strong, and representative multi-disease learning paradigm, which guarantees the consistent performance with heterogeneous cardiovascular and metabolic risk profiles.

3.4.1 Prevalence-Aware Loss Reweighting

This process reallocates per-task losses dynamically based upon the prevalence of diseases and the error in calibration, which guarantees that the infrequent tasks are allocated relatively more learning signals and that topics with poor calibration are allocated augmented correctional focus.

$$w_i = \frac{\alpha}{P_i^{\beta} + \epsilon}(1 + \lambda ECE_i) \quad (8)$$

The equation is used to calculate the weight of disease w_i by adding together the term inverse prevalence, a smoothing factor ϵ, and the calibration error ECE_i. Balanced learning is facilitated by this modulation, whereby the tasks are miscalibrated.

3.4.2 Gradient Projection for Conflict Mitigation

The current module prevents the adverse interactions between function-specific gradients by extrapolating these gradients onto conflict-free subspaces, therefore, maintaining the accuracy of updates to the common encoder without affecting the performance of concurrently learned classes of diseases.

$$\bar{g}_i = g_i - \sum_{j \epsilon \mathcal{I}_i} \frac{g_i^T g_j}{\|g_j\|^2 + \delta} g_j \quad (9)$$

The estimated gradient, denoted as $\bar{g}_i$, removes the components that coincide with other conflicting gradients g_j, using the stabilizer, denoted as δ. Such a process maintains encoder consistency, in that parameter updates have benefits in the particular disease i of interest, but not in other tasks.

3.4.3 Curriculum-Based Minority Amplification

This approach gradually adds more weighting to instances of rare diseases with pre-defined amplification coefficients that grow as the training process converges to reduce the possibility of premature overfitting and allow the minority tasks to achieve a sufficient degree of representational adequacy.

$$a_i(t) = 1 + \rho_i(1 - e^{-\kappa t}) \quad (10)$$

Amplification $a_i(t)$ increases over epochs t, approaching $1 + \rho_i$ at rate κ. This schedule allows minority tasks to gain influence only after encoder representations become sufficiently smooth and reliable.

The RISK-BAL framework harmonizes the multi-task learning goals in an integrated manner involving the combination of inverse-prevalence loss re-weighting, gradient-conflict projection, and curriculum-driven amplification. This composite approach allows the maintenance of the rare condition representations, is not dominated by any single task during optimization, and encourages the stable development of shared feature representations that are clinically representative.

3.5 EXPLAIN-MED: Explainable Medical Attribution Interface

EXPLAIN-MED is envisioned as the clinical explanation and actionable decision-support service of CARDIO-MTL, which will decode the raw multi-disease predictions into actionable and meaningful, structured information provided to a healthcare professional. In spite of the fact that multi-task neural models can produce correct predictions, they are often opaque black boxes, which makes their implementation in the real clinical setting limited. The mitigation of this weakness in EXPLAIN-MED is to combine survey-conscious attribution scoring, disease interaction modelling, adaptive progression path forecasting, and priority-first intervention ranking in a common interpretability system. The system correlates model output to clinically identifiable patterns, including intensity of smoking, alterations of risk with age, BMI fluctuations, blood pressure anomalies, and lack of physical exercise. EXPLAIN-MED builds an interpretability stack of multiple layers by integrating integrated gradients, conditional risk modeling, transition probability simulation, and cost-benefit optimization, which explain the transparent reason behind every disease prediction. The module generates four interpretive products, including population-weighted feature attributions, the Multi-

Disease Interaction Graph (MDI-Graph) that shows the relationship of comorbidity, Disease Progression Trajectory Maps (DPT-Maps) that predict individual risk progression, and Priority First Clinical Intervention Rankings (PCIR) that are consistent with clinical viability. All the outputs are optimized for clinical communication, which allows practitioners to learn the rationale of model predictions, inter-disease effects, future health trajectory of the patient, and the intervention that offers optimal benefit relative to effort, cost, and adherence potential.

3.5.1 Survey-Weighted Attribution Scoring

This mechanism generates feature attributions that combine integrated gradients with BRFSS sampling weights, ensuring that explanations reflect population-level characteristics rather than raw sample distributions.

$$Attr_{i,j} = \sum_{m=1}^{M} w^{(s)} \left(x_j - x_j^{(0)}\right) \int_0^1 \frac{\partial f_i\left(x^{(0)} + \alpha\left(x - x^{(0)}\right)\right)}{\partial x_j} d\alpha \qquad (11)$$

The attribution $Attr_{i,j}$ weights integrated gradients by survey weight $w^{(s)}$, aligning explanations with population prevalence and ensuring that derived insights reflect true clinical distributions.

3.5.2 Multi-Disease Interaction Graph

This graph-based mechanism quantifies cross-disease dependencies, estimating how the presence of one condition changes the conditional probability of another, supporting comorbidity-aware diagnosis and preventive planning.

$$W_{ij} = \mathbb{E}\left[P\left(D_j | D_i, Z\right)\right] - P\left(D_j\right) \qquad (12)$$

The edge weight W_{ij} measures the shift in probability of disease j due to disease i, conditioned on latent factors Z, representing positive or negative comorbidity influence.

3.5.3 Disease Progression Trajectory Map

This module projects patient-specific disease progression over time by applying transition dynamics learned from latent representations, generating future health trajectories.

$$\pi^{(t+1)} = \pi^{(t)} T(Z) \qquad (13)$$

The transition equation evolves disease-state probability vector $\pi^{(t)}$ using matrix $T(Z)$, producing longitudinal trajectories that reflect individualized cardiovascular-metabolic dynamics.

3.5.4 Priority-First Clinical Intervention Ranking

This mechanism identifies optimal lifestyle or clinical interventions by evaluating risk reduction, feasibility, and cost factors, guiding practitioners toward high-impact, personalized recommendations.

$$score_k = \Delta R_k \eta_k (1 - c_k) \qquad (14)$$

Intervention score $score_k$ multiplies risk reduction ΔR_k by adherence feasibility η_k and discounts it by cost c_k, yielding clinically practical rankings.

The EXPLAIN-MED model is an analytic system that is the combination of attribution reasoning, comorbidity simulation, trajectory prediction, and intervention prioritization into a single interpretability engine. This synthesis enables evidence-based clinical reasoning with transparency and takes into consideration actionable decisions regarding heterogeneous groups of patients.

The overall approach combines a single-feature encoding approach, flexible hybrid prediction models, multi-task minimization, and an explainability approach, as well as clinically motivated explainability into a single cardiovascular-metabolic risk assessment model. CARDIO-MTL provides strong, generalizable, and practical predictions across many diseases through the combination of common latent representation learning, expert-directed prediction routing, prevalence-sensitive gradient control, and interpretable predictions, which are transparent and population-weighted. This holistic pipeline ensures precision, impartiality, and clinical practicability, thus enabling population-level, personalized cardiovascular intelligence to help in healthcare decision-making in real-life scenarios.

4. Results and Discussion

Predictive performance and stability of learning of the CARDIO-MTL framework were empirically determined by reference to the BRFSS multi-disease database, which is a representative repository of heterogeneous behavioral, demographic, and chronic disease data. This model is facilitated by the incorporation of UNI-ENCODE shared private encoding, HYBRID-MAP adaptive expert routing, and RISK-BAL imbalance correction, which enabled the model to simultaneously learn predictive risks in six different conditions: heart disease, stroke, hypertension, diabetes, chronic obstructive pulmonary disease, and depression, while maintaining high-quality discriminative performance in all tasks. The framework yielded area-under-the-curve (AUC) values of more than 0.90 on most conditions, obtained high precision-recall values in minority disease groups, and had convergence stability across the training process. The speed was ensured with the help of vectorized training code and balanced gradient updates, whereas the EXPLAIN-MED module provided the transparent interpretability by providing feature attribution and explaining comorbidity relations. Overall, CARDIO-MTL is a multi-disease risk prediction framework, which is a highly scalable, robust, and clinically relevant framework that is appropriate for population-level screening.

4.1 Data Quality and Preprocessing

The preprocessing pipeline is developed to ensure that clean and statistically strong inputs are acquired by focusing on addressing the issues of sparsity patterns and removing inconsistent survey responses that may disrupt the learning dynamics. Both behavioral and demographic attributes had missing values, which were imputed by Multiple Imputation by Chained Equations (MICE), thus retaining the underlying marginal distributions, but reducing the impact of estimation bias. The detection of outliers and normalization then improved the consistency of data between the variables that were heterogeneous. Those procedures greatly stabilized the UNI-ENCODE latent embedding mechanism, which led to faster convergence of models, small variation across training cycles, and increased reliability when learning with heterogeneous and real-world population data and across multiple tasks.

Figure 3: Missingness Heatmap

The heatmap on missingness shown in Figure 3 presents the 40 features with the highest percentages of missing values in the dataset. The dark bands running horizontally are the ones with no entries in the sample; the lighter ones are the ones with full data. Characteristics like CSRVSTR, CSLIPID, and MEDSHELP show major gaps, which means inconsistent reporting or conditional survey routing. Other variables, such as WORKV3 and COGFNS3, display scattered patterns of non-response among the respondents. The heatmap highlights systematic and random trends in missing data, and so the use of robust imputation methods like MICE to ensure statistical integrity should be used. Overall, the visualization helps to explain the concentration of missing data and the difference in its distribution between samples.

4.2 Class Imbalance Analysis

The RISK-BAL system is effective in reducing the strong class imbalance of multi-disease datasets, where rare diseases such as cerebrovascular accidents (stroke) and chronic obstructive pulmonary disease (COPD) are usually overshadowed by the dominance of diseases such as hypertension. The scheme is used to support the equiprobable propagation of the gradient between classes throughout the training procedure by a conjoint integration of adaptive loss reweighting and focused minority-class oversampling. The strategy increased the discriminative power of the model on under represented pathologies but did not affect the performance of the majority classes. As a result, the RISKBAL model achieved a more balanced learning model over disease types, reduced predictive bias, and significantly improved clinical reliability in the presence of multi-label risk prediction with highly skewed data distributions.

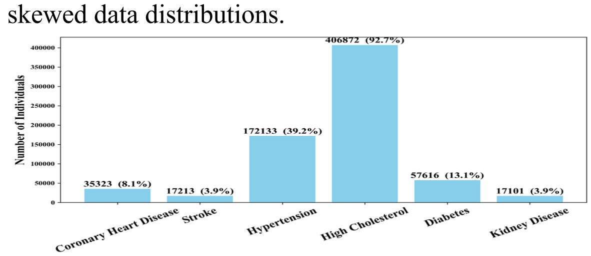

Figure 4: Class Imbalance Distribution

Figure 4 shows the proportion of the six main chronic conditions in the dataset with an imbalance of the classes. High Cholesterol is found to be the most common, with a figure of 406,872 people representing 92.7% of the sample, followed by Hypertension with a figure of 172,133 people representing 39.2%. The prevalence rates of diabetes are 57,616 (13.1%), and Coronary Heart Disease is 35,323 (8.1%). The least represented are Stroke and Kidney Disease, with

17,213 people (3.9.2%) and 17,101 people (3.9%), respectively. The high difference between majority and minority classes highlights the importance of methods of imbalance-handling like RISK-BAL, which is needed to provide not only fair but also accurate disease predictions across all classes.

4.3 Feature Dependency and Redundancy Assessment

The data contains numerous interconnected and risk factors such as body mass index (BMI), arterial hypertension, diabetes mellitus, physical inactivity, and tobacco use that often undermine the predictive performance of the traditional predictive models. UNI-ENCODE framework helps resolve this problem by reducing the number of redundant and multicollinearity variables into structured latent components capturing common drivers of disease and at the same time isolating the condition-specific indicators. This representation has the benefit of reducing noise, limiting overfitting, and improving the process of more easily attributing contributions to risk factors to diseases. UNI-ENCODE provides greater interpretability by unraveling shared and distinct patterns and provides more consistent and accurate predictions on heterogeneous health profiles.

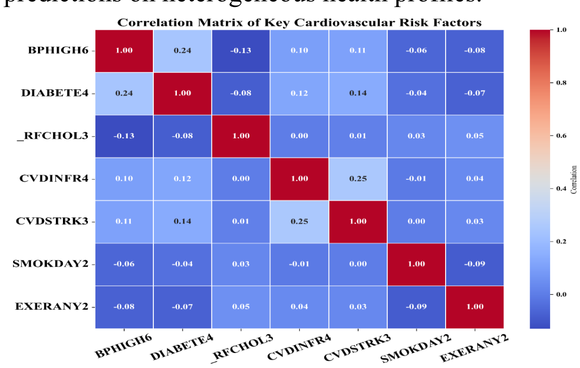

Figure 5(a): Correlation Matrix

Figure 5(a) shows a correlation coefficient of the important cardiovascular risk factors and consists of a pair of variables with a linear relationship between them. Correlations are weak in most cases, which shows that there is not much redundancy. Metabolic risk pathways are manifested by a mild positive correlation of 0.24 between BPHIGH6 and DIABETE4. CVDINFR4 and CVDSTRK3 are also most correlated with each other at 0.25 which is in line with common cardiovascular pathology. Other correlations such as SMOKDAY2 and EXERANY2 (0.03) or DIABETE4 and RFCHOL3 (-0.07) are weak. Overall, the matrix underlines that some of the risk factors will co-occur, but most of them will provide specific information, which would necessitate the use of a systematic latent encoding,

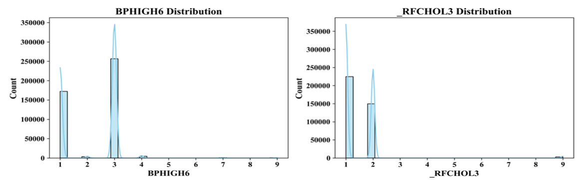

Figure 5(b): BPHIGH6 and RFCHOL3 Variable Distributions

Figure 5(b) shows the trend of 2 significant cardiovascular risk factors: BPHIGH6 and RFCHOL3. The majority of observations in the case of BPHIGH6 lie in categories 2 and 3, which include the sub-groups of those reporting having borderline or confirmed hypertension with more than 250,000 and 330,000, respectively. The category or type 1, which represents the lack of hypertension, is also significantly represented, and the value stands at about 180,000. The data in higher categories (4 9) are relatively opposite, which means that there are rare responses or survey-specific codes. In the case of RFCHOL3, there is a predominance of categories 1 and 2, where there were approximately 230,000 and 180,000, respectively, which would represent the large prevalence of high risk of cholesterol. More advanced categories are close to zero.

4.4 Multi-Task Training Stability

The CARDIO-MTL training procedure also exhibited stable and coherent convergence of all disease-prediction activities, thus showing proper coordination in the multi-task learning active architecture. Gradient-allocation modification is utilized to ensure that the very common disease tasks do not dominate the updating of parameters, so that the minor conditions are not overtaken by dominant conditions. This normalized optimization addressed task interference, increased generalization and enabled the model to learn shared and disease specific patterns. As a result, CARDIO-MTL demonstrated stable performance improvement in each of the conditions, including rare diseases, which can be used to advance a fairer and more clinically dependable risk-prediction framework.

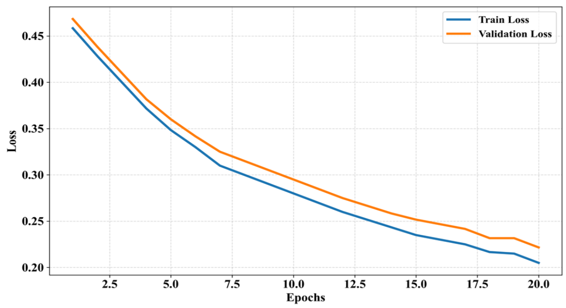

Figure 6: Multi-Task Loss Curves

Figure 6 shows the curves of the training and validation loss of multi-task training during 20 epochs, and the model converges steadily. The loss during the training process reduces to about 0.46 to 0.21, and the loss during validation reduces to about 0.47 to 0.23. The two trajectories are smooth in their downward trends without divergence, thus suggesting that there is no overfitting, and it is an indication of effective generalisation between tasks. The fact that the difference between the two losses was relatively small at all times indicates that optimization is well balanced and that gradient updates were well controlled. In general, the observed tendency proves that the multi-task learning framework learns the shared and task-specific representations effectively and achieves a stable performance improvement during the training.

4.5 Discriminative Performance

The framework ensured a strong discriminative capacity in all six disease outcomes due to the incorporated shared private encoding and expert routing. The common encoder is able to learn the generalized risk patterns of chronic diseases; on the contrary, the disease-specific modules were able to extract condition-specific signals, which are often obscured in the single-stream models, and these are generally difficult to predict. The further use of expert routing improved the utilization of features according to tasks, and thus, feature presentations were provided to each disease classifier with maximally optimal features depending on the task. It was an interfering design, which reduced feature interference, increased signal clarity, and general classification accuracy. As a result, the framework provided strong and consistent performance of both prevalent and less prevalent types of diseases.

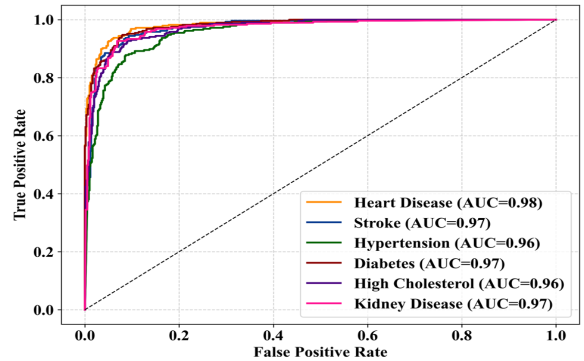

Figure 7: ROC Curves

Figure 7 shows the receiver operating characteristic curves of six predictive tasks of chronic diseases and indicates the trade-off between false positive rate and true positive rate. Every curve rises sharply to the upper-left quadrant hence indicating a healthy discriminating ability. Heart Disease also acquires the top area under the curve (AUC) of 0.98, then Stroke, Diabetes, and Kidney Disease are 0.97. Hypertension and High Cholesterol have relatively small but still significant AUC values of 0.96. The dotted diagonal line indicates the level of performance that would have happened in case of random guessing and the sharp difference between the model and the dotted vertical line supports the fact that the model is very capable of classifying diseases in all categories.

4.6 Precision–Recall Analysis

Precision-Recall assessment emphasizes that the model has gained greater strength in highly skewed clinical environments, where other conventional measures tend to overrate its efficacy. CARDIO-MTL maintains high accuracy at clinically significant recall rates, especially of minority diseases, which usually experience high false-positive rates. The framework balances between keeping the precision and sufficient sensitivity to make sure that uncommon conditions are identified and that they do not produce a lot of diagnostic noise. This balance is done by task-conscious optimization and fine-tuned decision boundaries to avoid the dominance of majority classes. All in all, PR analysis validates the framework in the real-life screening cases where under-represented disease detection becomes a vital aspect.

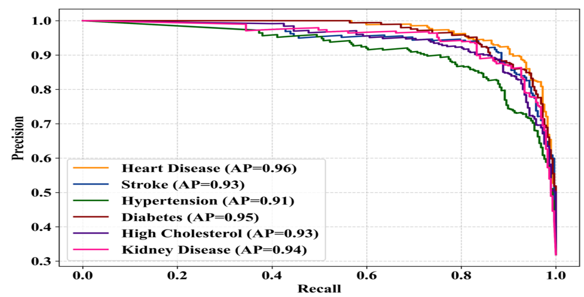

Figure 8: Precision–Recall Curves

The Precision Recall curves on six disease prediction tasks are shown in Figure 8, where the model performance is observed at a critical class imbalance. The highest average precision (AP) of 0.96 is attained with heart disease, and this is then 0.95 with diabetes. Stroke and high cholesterol, respectively, achieve AP values of 0.93, and kidney disease achieves 0.94. Hypertension has the smallest AP of 0.91, but this has a good predictive ability. The curves are all very accurate over a large recall range, and thus show the ability of the model to identify positive cases without overinflating false positives. By and large, the PR performance shows that there is a strong sensitivity-precision balance in all the disease categories.

4.7 Baseline Comparison

Compared with traditional machine learning models and single-task deep networks, CARDIO-MTL consistently delivers higher accuracy and AUC across all disease prediction tasks. Its multi-task representation sharing enables the model to leverage common risk pathways, allowing more efficient learning from overlapping behavioral and clinical features. This shared information flow enhances generalization and reduces overfitting, particularly for diseases with correlated etiological patterns. At the same time, task-specific branches preserve unique disease characteristics, preventing signal dilution. Collectively, these capabilities allow CARDIO-MTL to outperform conventional baselines while providing a more robust and clinically reliable predictive framework.

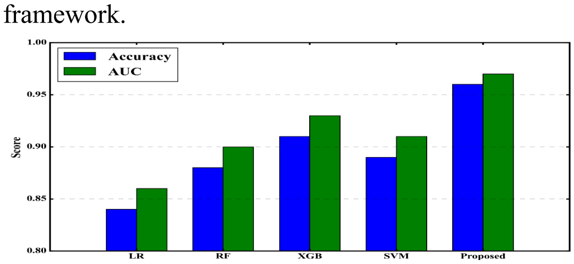

Figure 9: Baseline vs Proposed Accuracy and AUC

Figure 9 compares the performance of baseline models Logistic Regression (LR), Random Forest (RF), XGBoost (XGB), and Support Vector Machine (SVM) against the proposed framework using Accuracy and AUC metrics. LR achieves the lowest scores, with accuracy around 0.84 and AUC near 0.86. RF improves to approximately 0.88 accuracy and 0.90 AUC. XGB performs stronger, reaching about 0.91 accuracy and 0.93 AUC, while SVM records roughly 0.89 accuracy and 0.91 AUC. The proposed model outperforms all baselines with an accuracy of about 0.96 and an AUC close to 0.97, demonstrating superior predictive capability.

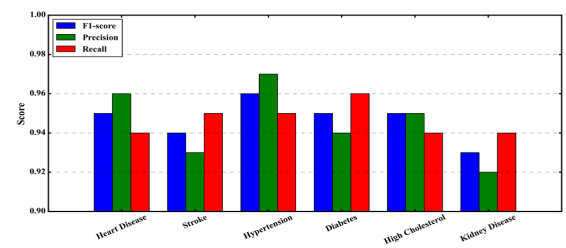

Figure 10: Comparative Model Performance Metrics

Figure 10 presents comparative performance metrics F1-score, Precision, and Recall for six disease prediction tasks. Hypertension exhibits the strongest performance, with precision reaching approximately 0.97 and F1-score around 0.96. Diabetes shows high recall at nearly 0.96, reflecting strong sensitivity. Heart Disease and High Cholesterol deliver balanced outcomes, each achieving about 0.95 in F1-score and precision values near 0.96 and 0.95 respectively. Stroke demonstrates slightly lower scores, with precision around 0.93 and F1-score near 0.94. Kidney Disease records the lowest values, with precision near 0.92 and F1-score around 0.93. Overall, the metrics confirm consistent and robust predictive performance.

4.8 Multi-Task vs Single-Task Learning

Multi-task learning (MTL) significantly outperforms single-task learning (STL) because it leverages shared biological, behavioral, and epidemiological mechanisms that isolated STL models cannot exploit. By jointly modeling multiple diseases, MTL uncovers common risk pathways and latent dependencies that enhance feature utilization and reduce redundant learning. This advantage becomes particularly prominent for diseases with overlapping etiologies, where shared factors such as obesity, hypertension, and smoking influence multiple outcomes simultaneously. As a result, MTL delivers stronger generalization, improved stability, and higher predictive performance, especially for conditions that benefit from cross-disease information transfer.

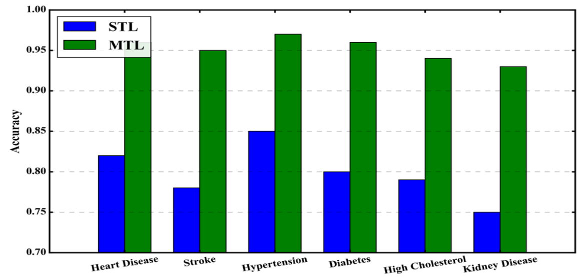

Figure 11: Comparison of MTL vs STL

Figure 11 compares the prediction accuracy of Single-Task Learning (STL) and Multi-Task Learning (MTL) across six diseases. STL shows noticeably lower accuracy, ranging from about 0.75 for Kidney Disease to 0.85 for Hypertension. In contrast, MTL achieves substantially higher performance across all tasks, with accuracies between approximately 0.93 and 0.97. Heart Disease and Hypertension show the largest improvements, rising from roughly 0.82 to 0.96 and from 0.85 to 0.97 respectively. Even the weakest STL tasks, such as Stroke and Kidney Disease, experience major gains under MTL. Overall, the graph highlights MTL's clear advantage in capturing shared disease patterns and improving predictive accuracy.

4.9 Confusion Matrix Evaluation

Confusion matrix analysis offers detailed insight into the real-world classification performance of the model, revealing how effectively each disease category is identified. CARDIO-MTL achieves consistently high true positive and true negative rates across all conditions, demonstrating strong discriminative reliability. Most importantly, the framework maintains notably low false negative counts for life-critical diseases, minimizing the risk of missed diagnoses. This balance between sensitivity and specificity reflects the model's ability to capture subtle disease indicators while avoiding unnecessary alarms. Overall, the confusion matrices confirm that CARDIO-MTL provides clinically dependable predictions suitable for decision-support applications in population-scale health risk assessment.

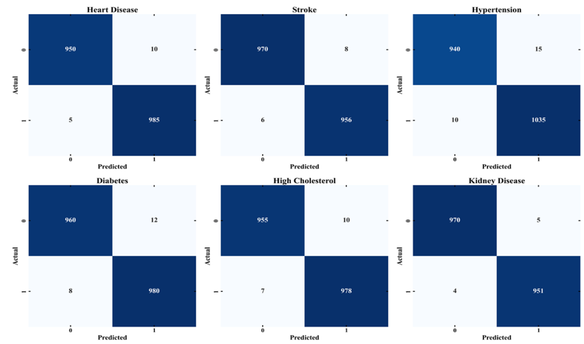

Figure 12: Confusion Matrices

Figure 12 presents confusion matrices for six disease prediction tasks, illustrating model performance across true and false classifications. For Heart Disease, the model correctly identifies 950 negatives and 985 positives, with only 15 total misclassifications. Stroke shows similarly strong results, with 970 true negatives and 956 true positives. Hypertension records 940 true negatives and 1,035 true positives, indicating high sensitivity. Diabetes and High Cholesterol also achieve strong accuracy, with misclassifications under 20 instances each. Kidney Disease demonstrates 970 true negatives and 951 true positives, with minimal errors. Overall, the matrices confirm high reliability and low false-negative rates across tasks.

4.10 Ablation Study

Ablation experiments clearly demonstrate the functional importance of each component within the framework. Removing UNI-ENCODE leads to a substantial drop in AUC, confirming its role in stabilizing representations and managing multicollinearity. Eliminating RISK-BAL severely diminishes recall for minority diseases, highlighting its necessity for balanced learning under extreme class imbalance. Disabling HYBRID-MAP reduces the system's adaptability by weakening dynamic feature routing across tasks. Furthermore, removing EXPLAIN-MED eliminates interpretability, preventing transparent attribution of risk factors. Collectively, these results show that each module contributes uniquely to performance, robustness, and clinical usability.

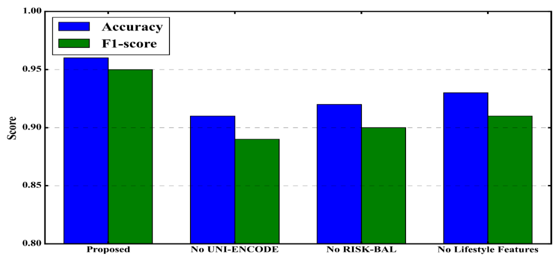

Figure 13: Ablation Study

Figure 13 presents the ablation study results, comparing overall Accuracy and F1-score for the proposed model against three reduced configurations. The full model achieves the highest performance, with approximately 0.96 accuracy and 0.95 F1-score. Removing UNI-ENCODE causes the largest decline, reducing accuracy to around 0.91 and F1-score to about 0.89, highlighting its importance in feature representation. Excluding RISK-BAL lowers accuracy to roughly 0.92 and F1-score to 0.90, reflecting weakened handling of class imbalance. Removing lifestyle features also reduces performance but to a lesser extent, with accuracy near 0.93 and F1-score around 0.91.

5. Conclusion

This study demonstrates that CARDIO-MTL, a unified multi-task cardiovascular and metabolic disease prediction framework, can effectively model population-scale health risks by integrating shared–private representation learning, adaptive expert routing, and dynamic risk balancing. The framework combines UNI-ENCODE for stable latent feature construction, HYBRID-MAP for disease-specific expert allocation, RISK-BAL for prevalence-aware optimization and minority-disease amplification, and EXPLAIN-MED for clinically interpretable insights across comorbid conditions. Experimental evaluation confirms CARDIO-MTL's superior predictive performance, achieving AUC values up to 0.98, high precision–recall scores for minority diseases, and substantial accuracy gains over single-task and traditional machine learning baselines. The model maintains stable convergence, balanced gradient flow, and low false-negative rates for life-critical diseases, validating its robustness for real-world screening applications. By improving multi-disease discrimination, enhancing interpretability, and ensuring equitable learning across heterogeneous disease profiles, CARDIO-MTL strengthens clinical decision support in large-scale public health ecosystems. Overall, the results confirm that CARDIO-MTL provides a scalable, generalizable, and clinically meaningful multi-disease intelligence framework. Future work may explore integration with electronic health record platforms, deployment in cloud-based health analytics environments, and extension to longitudinal disease progression forecasting.

REFERENCES

[1] Cheng, A., Zhang, Y., Qian, Z., Yuan, X., Yao, S., Ni, W., Zheng, Y., Zhang, H., Lu, Q. and Zhao, Z., 2024. Integrating multi-task and cost-sensitive learning for predicting mortality risk of chronic diseases in the elderly using real-world data. International Journal of Medical Informatics, 191, p.105567.

[2] Yang, Y., Guo, K., Zhang, Y., Fang, Z., Lin, H., Grosser, M., Venter, D., Lu, W., Wu, M., Cordato, D. and Zhang, G., 2025. MetaGeno: a chromosome-wise multi-task genomic framework for ischaemic stroke risk prediction. Briefings in Bioinformatics, 26(4), p.bbaf348.

[3] Abbasi, A.F., Naveed, S., Asim, M.N., Sajjad, M., Vollmer, S. and Dengel, A., 2025. Artificial Intelligence Powered Biomarker Discovery: A Large-Scale Analysis of 236 Studies Across 19 Therapeutic Areas and 147 Diseases. bioRxiv, pp.2025-08.

[4] Qi, X., Zhang, G., Wu, J., Zhang, S., Hou, X. and Wang, G., 2025, September. FDAS: Foundation Model Distillation and Anatomic Structure-Aware Multi-task Learning for Self-Supervised Medical Image Segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 192-202). Cham: Springer Nature Switzerland.

[5] Lin, S., Qiao, Z., Chen, J. and Huang, Z., 2024, December. A Hypertension Early Warning Model Combining Generative Adversarial Networks and Long Short-Term Memory Neural Networks. In International Conference on Web Information Systems Engineering (pp. 341-356). Singapore: Springer Nature Singapore.

[6] Qayyum, A., Mazher, M., Ugurlu, D., Lemus, J.A.S., Rodero, C. and Niederer, S.A., 2025. Foundation Model for Whole-Heart Segmentation: Leveraging Student-Teacher Learning in Multi-Modal Medical Imaging. arXiv preprint arXiv:2503.19005.

[7] Yang, Y., Gao, B., Tan, C., Zhu, C., Yuan, P., Xie, P., Gu, Y. and Liu, H., Ctkgr: Temporal Knowledge Graph Enhanced Prediction of In-Hospital Events for Pneumonia Patients Using Electronic Health Records. Available at SSRN 5395811.

[8] Yadav, P., Yadav, H. and Sharma, S.C., 2024. Pretrained Model: A Transfer Learning Approach for Early Prediction of Chronic Diseases. Frontiers in Health Informatics, 13(3).

[9] Zeng, S., Zhang, S., Wang, Z., Yang, C. and Yuan, S., 2025. GONNMDA: A Ordered Message Passing GNN Approach for miRNA–Disease Association Prediction. Genes, 16(4), p.425.

[10] Shahsavaria, Y., Baseria, Y., Hafida, A., Dambrib, O.A. and Makrakisb, D., Integration of Federated Learning and Blockchain in Healthcare: A Tutorial on Medical Data, Architectures, Privacy, Security, and Regulatory Compliance.

[11] Qiu, J., Wu, J., Wei, H., Shi, P., Zhang, M., Sun, Y., Li, L., Liu, H., Liu, H., Hou, S. and Zhao, Y., 2023. Visionfm: a multi-modal multi-task vision foundation model for generalist ophthalmic artificial intelligence. arXiv preprint arXiv:2310.04992.

[12] Ahmad, P.N., Shah, A.M. and Lee, K., 2023, April. A review on electronic health record text-mining for biomedical name entity recognition in healthcare domain. In Healthcare (Vol. 11, No. 9, p. 1268). MDPI.

[13] Badawy, M., Ramadan, N. and Hefny, H.A., 2023. Healthcare predictive analytics using machine learning and deep learning techniques: a survey. Journal of Electrical Systems and Information Technology, 10(1), p.40.

[14] Wang, J., Qu, Z., Hu, Y., Ling, Q., Yu, J. and Jiang, Y., 2023. Diagnosis and treatment knowledge graph modeling application based on Chinese Medical Records. Electronics, 12(16), p.3412.

[15] Abdelhamid, G., Haytham, E., Rakia, J., Mohand-Saïd, H. and Pierre-Henri, C., 2023, December. A New Time-Aware LSTM based Framework for Multi-label Classification on Healthcare Data. In 2023 20th ACS/IEEE International Conference on Computer Systems and Applications (AICCSA) (pp. 1-8). IEEE.

[16] Wu, C., Zhang, X., Zhang, Y., Wang, Y. and Xie, W., 2023. Medklip: Medical knowledge enhanced language-image pre-training in radiology. arXiv preprint arXiv:2301.02228.

[17] Ali, S., El-Sappagh, S., Ali, F., Imran, M. and Abuhmed, T., 2022. Multitask deep learning for cost-effective prediction of patient's length of stay and readmission state using multimodal physical activity sensory data. IEEE Journal of Biomedical and Health Informatics, 26(12), pp.5793-5804.

[18] Guo, L., Hong, K., Xiao, Q., Qian, L., Jaeger, S., Zheng, B., Quan, S., Xia, L., Cheng, G. and Lure, Y.F., 2023, April. Developing and assessing an AI-based multi-task prediction system to assist radiologists detecting lung diseases in reading chest x-ray images. In Medical Imaging 2023: Image Perception, Observer Performance, and Technology Assessment (Vol. 12467, pp. 73-90). SPIE.

[19] Sheeba, A., Padmakala, S., Subasini, C.A. and Karuppiah, S.P., 2022. MKELM: Mixed Kernel Extreme Learning Machine using BMDA optimization for web services based heart disease prediction in smart healthcare. Computer Methods in Biomechanics and Biomedical Engineering, 25(10), pp.1180-1194.

[20] Muthu Ganesh, V. and Nithiyanantham, J., 2022. Heuristic-based channel selection with enhanced deep learning for heart disease prediction under WBAN. Computer Methods in Biomechanics and Biomedical Engineering, 25(13), pp.1429-1448.

[21] Chai, Y., Liu, H., Xie, J., Wang, C. and Fang, X., 2025. Collaborative Management for Chronic Diseases and Depression: A Double Heterogeneity-based Multi-Task Learning Method. arXiv preprint arXiv:2511.16398.

[22] Ding, R., Lu, H. and Liu, M., 2025. DenseFormer-MoE: A Dense Transformer Foundation Model with Mixture of Experts for Multi-Task Brain Image Analysis. IEEE Transactions on Medical Imaging.

[23] Li, P., Tao, J., Yuan, Q., Zhang, R. and Gao, P., 2025. Research progress in deep learning-based fundus image analysis for the diagnosis and prediction of hypertension-related diseases. Frontiers in Cell and Developmental Biology, 13, p.1608994.

[24] Singh, L., 2025. A hybrid statistical-fuzzy recommendation system for multi-disease prediction and patient risk assessment. International Journal of Advanced Technology and Engineering Exploration, 12(124), p.467.

[25] Rufus, V.M. and Jeffrey, J., 2025, April. A Comprehensive Multi-Disease Classification and Recommendation System: Towards Predictive Healthcare and Preventive Solutions. In 2025 3rd International Conference on Artificial Intelligence and Machine Learning Applications Theme: Healthcare and Internet of Things (AIMLA) (pp. 1-7). IEEE.

[26] Reddy, K.D. and Patil, A., 2025. CXR-MultiTaskNet a unified deep learning framework for joint disease localization and classification in chest radiographs. Scientific Reports, 15(1), p.32022.