

# DeepGB-TB: A Risk-Balanced Cross-Attention Gradient-Boosted Convolutional Network for Rapid, Interpretable Tuberculosis Screening

Zhixiang Lu<sup>1,2\*</sup>, Yulong Li<sup>1,2,3\*</sup>, Feilong Tang<sup>3</sup>, Zhengyong Jiang<sup>1</sup>, Chong Li<sup>1</sup>, Mian Zhou<sup>1†</sup>, Tenglong Li<sup>4†</sup>, Jionglong Su<sup>1†</sup>

<sup>1</sup>School of Artificial Intelligence and Advanced Computing, Xi'an Jiaotong-Liverpool University

<sup>2</sup>Department of Computer Science, University of Liverpool

<sup>3</sup>Mohamed bin Zayed University of Artificial Intelligence

<sup>4</sup>Wisdom Lake Academy of Pharmacy, Xi'an Jiaotong-Liverpool University

{Mian.Zhou, Tenglong.Li, Jionglong.Su}@xjtlu.edu.cn

## Abstract

Large-scale tuberculosis (TB) screening is limited by the high cost and operational complexity of traditional diagnostics, creating a need for artificial-intelligence solutions. We propose DeepGB-TB, a non-invasive system that instantly assigns TB risk scores using only cough audio and basic demographic data. The model couples a lightweight one-dimensional convolutional neural network for audio processing with a gradient-boosted decision tree for tabular features. Its principal innovation is a Cross-Modal Bidirectional Cross-Attention module (CM-BCA) that iteratively exchanges salient cues between modalities, emulating the way clinicians integrate symptoms and risk factors. To meet the clinical priority of minimizing missed cases, we design a Tuberculosis Risk-Balanced Loss (TRBL) that places stronger penalties on false-negative predictions, thereby reducing high-risk misclassifications. DeepGB-TB is evaluated on a diverse dataset of 1,105 patients collected across seven countries, achieving an AUROC of 0.903 and an F1-score of 0.851, representing a new state of the art. Its computational efficiency enables real-time, offline inference directly on common mobile devices, making it ideal for low-resource settings. Importantly, the system produces clinically validated explanations that promote trust and adoption by front-line health workers. By coupling AI innovation with public-health requirements for speed, affordability, and reliability, DeepGB-TB offers a tool for advancing global TB control.

## Introduction

Tuberculosis (TB) inflicts a staggering toll on global health, remaining a leading infectious cause of death worldwide (World Health Organization 2023, 2024). The cornerstone of TB control, early diagnosis and treatment, is severely undermined in many regions by the limitations of conventional methods. Diagnostic tools like sputum smear microscopy and nucleic acid amplification tests (NAATs) either suffer from low sensitivity or are prohibitively expensive and require centralized laboratories and skilled technicians, creating significant barriers to access in low-resource

settings (Lawn and Zumla 2011; Boehme et al. 2010). This diagnostic gap leads to delayed treatment, increased transmission, and preventable mortality (World Health Organization 2015), underscoring an urgent need for accessible, affordable, and scalable screening solutions.

The ubiquity of mobile phones presents a unique opportunity to bridge this gap. Cough, a cardinal symptom of pulmonary TB, contains a wealth of acoustic information that, if properly analyzed, could serve as a non-invasive digital biomarker (Imran, Posokhow et al. 2019). While prior research has explored AI for cough-based analysis, existing models often face critical limitations. Many either rely solely on audio, ignoring crucial demographic and clinical risk factors (e.g., age, sex, exposure history), or they struggle to effectively integrate these heterogeneous data types. Simple concatenation or late-stage fusion of features often fails to capture the complex, non-linear interplay between a patient's background risk and their real-time acoustic symptoms, a process central to a clinician's diagnostic reasoning (Ramachandram and Taylor 2017).

To address these shortcomings, we propose DeepGB-TB, a novel, multimodal deep learning system designed for end-to-end, instantaneous TB risk stratification. Our architecture is explicitly designed to model the synergy between who the patient is and how they cough. The framework processes two parallel data streams: a lightweight 1D Convolutional Neural Network (CNN) extracts discriminative features from raw cough audio, while demographic data is handled by our first key innovation. (1) The **Cross-Validated Probability Embedding Module (CVPEM)** transforms raw tabular data into a robust, high-dimensional feature vector, a technique designed to mitigate overfitting and enhance generalization. These distinct data pathways are then unified by (2) the **Integrated Multimodal Diagnostic Module (IMDM)**. The centerpiece of this module is (3) the **Cross-Modal Bidirectional Cross-Attention (CM-BCA)** mechanism. This module moves beyond feature fusion and emulates clinical ratiocination by allowing the audio and tabular embeddings to iteratively query each other, mutually refining their representations to focus on the most salient diagnostic indicators.

Furthermore, recognizing that failing to detect a true TB

\*These authors contributed equally.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

case (a false negative) has far more severe consequences than a false alarm, we introduce a Tuberculosis Risk-Balanced Loss (TRBL) function. This clinically-aligned loss function imposes a greater penalty on false negatives, systematically steering the model towards the high sensitivity required for an effective screening tool.

Our main contributions are summarized as follows:

1) **A Novel Hybrid Architecture:** We introduce DeepGB-TB, a carefully designed deep learning framework that synergistically combines a 1D-CNN for audio analysis with a novel CVPEM-enhanced gradient boosting model for tabular data, enabling a holistic patient assessment.

2) **Advanced Multimodal Fusion:** We propose the CM-BCA cross-attention module that achieves integration of heterogeneous data by modeling the bidirectional dependencies between acoustic and demographic features.

3) **Clinically-Informed Optimization:** We design the TRBL function, a custom loss function that addresses the clinical priority of minimizing false negatives in TB screening, enhancing the model’s real-world utility and safety.

4) **State-of-the-Art Performance and Deployability:** We demonstrate that DeepGB-TB achieves state-of-the-art results on multi-national dataset and is computationally efficient for real-time, offline deployment on mobile devices, paving the way for equitable access to TB screening.

## Related Works

Early work in AI-driven TB diagnosis predominantly focused on analyzing Chest X-rays (CXRs), where Convolutional Neural Networks (CNNs) have achieved radiologist-level performance (Lakhani and Sundaram 2017). However, the reliance on specialized hardware for CXR imaging limits its applicability for widespread, community-level screening. This has motivated a research shift towards more accessible biomarkers, with cough sounds emerging as a leading alternative. The analysis of cough sounds via AI has emerged as a promising, low-cost screening paradigm. Initial approaches relied on traditional audio features paired with classical machine learning models like Logistic Regression (Cox 1958). More recently, deep learning has become the standard, with architectures such as 1D-CNNs (Kiranyaz, Ince, and Gabbouj 2016) and deeper models like ResNet (He et al. 2016) applied to spectrograms. The state-of-the-art (SOTA) includes pre-trained audio foundation models like Google’s HeAR (Google Research 2024), which demonstrated superior performance on 33 health-related acoustic tasks, including TB recognition. Nevertheless, HeAR is not open-source and accessible only via a restricted API for online inference, limiting adoption in resource-constrained settings and creating a critical gap for an efficient, open-source TB screening tool. Clinical TB diagnosis is inherently multimodal, integrating symptoms, history, and risk factors. Common multimodal fusion baselines, such as late-fusion CNN-LightGBM ensembles (Lu 2023), often fail to capture deep inter-modal synergies (Li et al. 2025a). While advanced tabular models like TabTransformer (Huang et al. 2020) exist, their effective integration remains challenging. Frontier models include Large Multimodal Models (LMMs) like

Qwen-Omni (Xu et al. 2025), offering general-purpose understanding (Li et al. 2025b; Zhang et al. 2025). However, LMMs are computationally prohibitive, and their superiority over specialized models for this diagnostic task is unestablished. Furthermore, the poor interpretability of many deep learning models (Zhang et al. 2021; Smith, Doe et al. 2022) remains an obstacle to reliable clinical deployment.

## Methodology

### Dataset

The dataset used in this study (Jaganath et al. 2024) was derived from a retrospective, multicenter case-control investigation encompassing 1,105 adult participants across seven countries (India, the Philippines, South Africa, Uganda, Vietnam, Tanzania, and Madagascar). Each participant presented with a new or worsening cough lasting at least two weeks. Comprehensive clinical, demographic, and cough acoustic data were collected for each participant. Participants were classified as TB positive if testing positive on Mycobacterial culture (Kent and Kubica 1985), Xpert MT-B/RIF (Boehme et al. 2010), or Xpert Ultra (Chakravorty et al. 2017), and TB negative if negative on all assays. The study received ethical approval from all participating institutional review boards. Exploratory analysis revealed significant group differences: TB prevalence was higher in males (33.2%) than females (19.7%). Furthermore, hemoptysis (43.2% vs. 24.3%), night sweats (39.0% vs. 17.3%), fever (40.0% vs. 16.1%), and weight loss (36.5% vs. 11.9%) were all strongly associated with TB.

### Data Pre-processing and Feature Extraction

Raw audio recordings first undergo standard pre-processing, including spectral subtraction for noise removal (Boll 1979) and peak volume normalization. From the cleaned audio segments, we extract a comprehensive feature set. The primary features are Mel-frequency cepstral coefficients (MFCCs) (Davis and Mermelstein 1980), which robustly capture spectro-temporal characteristics. The  $n$ -th coefficient  $c_n$  is computed as:

$$c_n = \sum_{k=1}^M (\log S_k) \cos \left[ n \left( k - \frac{1}{2} \right) \frac{\pi}{M} \right], \quad n = 1, \dots, L \quad (1)$$

where  $S_k$  are the log-energies from  $M$  Mel-scaled filterbanks and  $L$  is the number of coefficients. To create a richer acoustic representation, we supplement MFCCs with auxiliary features including spectral centroid, chroma, zero-crossing rate (ZCR), and fundamental frequency (F0) (Giannakopoulos and Pikrakis 2014).

### Statistical Analysis

To provide an interpretable analysis of the high-dimensional and often collinear acoustic features (Hastie, Tibshirani, and Friedman 2009; Kuhn and Johnson 2013), we employ a classical statistical pipeline. First, we use t-distributed Stochastic Neighbor Embedding (t-SNE) (Maaten and Hinton 2008) for dimensionality reduction, which minimizes the

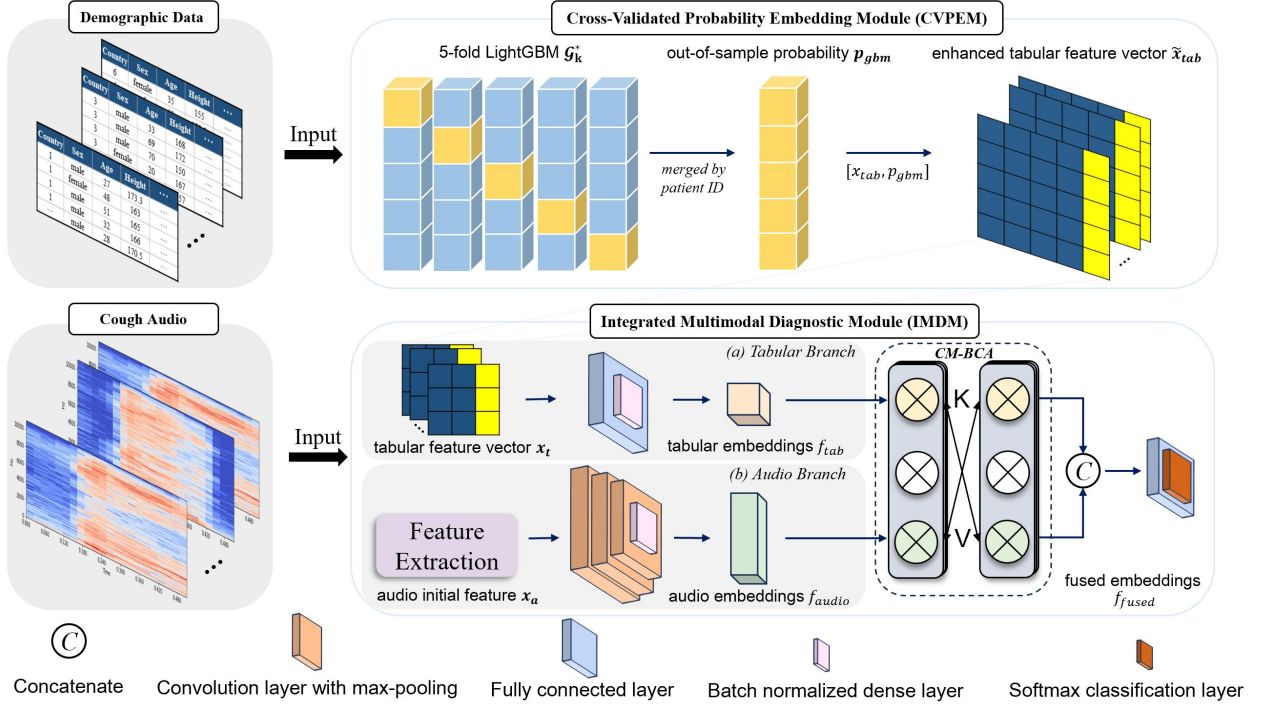


Figure 1: The architecture of DeepGB-TB.

Kullback-Leibler (KL) divergence between the joint probability distributions of the original data,  $p_{ij}$ , and the low-dimensional embeddings,  $q_{ij}$ :

$$C = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}} \quad (2)$$

A logistic regression model is then trained on the resulting embeddings to assess predictive power. We validate the structural integrity of this non-linear mapping using a Mantel correlation test (Mantel 1967; Wattenberg, Viégas, and Johnson 2016) and confirm the significance of the most discriminative features using Wald and independent two-sample t-tests.

### Proposed Framework

The architecture of DeepGB-TB, shown in Figure 1, integrates two key modules: the Cross-Validated Probability Embedding Module (CVPEM) and the Integrated Multimodal Diagnostic Module (IMDM). CVPEM employs LightGBM to generate robust probability embeddings from tabular demographic data, while IMDM processes cough audio signals via a 1D-CNN to extract temporal features. DeepGB-TB effectively fuses these complementary modalities, offering a comprehensive, robust, and accurate framework for TB diagnosis.

**Cross-Validated Probability Embedding Module (CVPEM).** In our proposed pipeline, LightGBM is trained with a 5-fold cross-validation scheme on demographic data to generate a stable, out-of-sample probability estimate for each patient. These probabilities are then merged by

patient ID and embedded as an additional feature, effectively capturing cross-validated insights into TB risk. Let  $\{(x_{\text{tab},i}, y_i)\}_{i=1}^n$  denote a dataset with tabular demographic features  $x_{\text{tab},i} \in \mathbb{R}^{d_{\text{tab}}}$  and binary labels  $y_i \in \{0, 1\}$ . We partition the index set  $i \in \{1, 2, \dots, n\}$  into  $K = 5$  disjoint subsets  $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_5\}$  for 5-fold cross-validation. For each fold  $k \in \{1, \dots, 5\}$  to train a LightGBM model  $\mathcal{G}_k^*$  on  $\bigcup_{j \neq k} \mathcal{I}_j$ . For each  $i \in \mathcal{I}_k$ , compute the out-of-sample probability:

$$p_{\text{gbm},i} = \mathcal{G}_k^*(x_{\text{tab},i}). \quad (3)$$

Collecting these out-of-sample predictions  $\{p_{\text{gbm},i}\}_{i=1}^n$  yields  $p_{\text{gbm}} \in \mathbb{R}^n$ . Each instance  $i$  is thus assigned an enhanced tabular feature vector:

$$\tilde{x}_{\text{tab},i} = [x_{\text{tab},i}, p_{\text{gbm},i}] \in \mathbb{R}^{d_{\text{tab}}+1}. \quad (4)$$

By embedding the cross-validated probability  $p_{\text{gbm},i}$  into the tabular features, CVPEM leverages robust out-of-sample estimates from LightGBM, which are subsequently fed into the downstream network to provide an enriched input representation that boosts both predictive performance and interpretability.

**Integrated Multimodal Diagnostic Module (IMDM).** Our proposed DeepGB-TB framework incorporates the IMDM to seamlessly combine heterogeneous data modalities for TB diagnosis. In this module, tabular (demographic) data are processed using a boosting-based method (LightGBM) to yield preliminary diagnostic probabilities, while cough audio features are extracted via a 1D-CNN. These two branches are then fused in a unified architecture, effectively bridging classical statistical methods with advanced

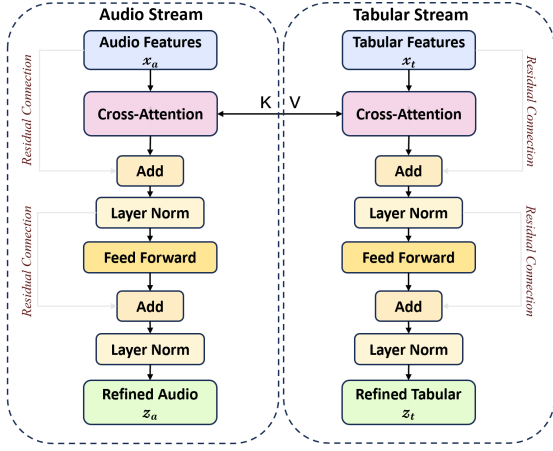


Figure 2: The process of CM-BCA.

deep learning techniques to achieve both high predictive performance and improved interpretability. Let  $\{(T_i, s_i)\}_{i=1}^n$  denote a dataset where  $T_i \in \mathbb{R}^{d_{\text{tab}}+1}$  is the enhanced tabular feature vector (demographic data plus the cross-validated probability embedding  $p_{\text{gbm}}$ ), and  $s_i \in \mathbb{R}^L$  is the raw 1D cough audio signal for the  $i$ -th patient. The IMDM consists of two primary branches (tabular and audio) that converge in a final fusion layer to produce the TB diagnosis probability.

**(1) Tabular Branch.** A batch normalization (BN) is applied to stabilize tabular inputs, followed by a dense layer:

$$T'_i = \text{BN}(T_i), \quad f_{\text{tab},i} = \text{ReLU}(W^{(t)} T'_i + b^{(t)}). \quad (5)$$

**(2) Audio Branch.** A preprocessing function (e.g., MFCC extraction) transforms  $s_i$  into an initial feature map:

$$F_i^{(0)} = \text{Preprocess}(s_i). \quad (6)$$

Next, three consecutive convolutional blocks are applied (each comprising a 1D convolution, batch normalization, ReLU activation, and max pooling), yielding:

$$F_i^{(l)} = \text{MaxPool}\left(\text{ReLU}\left(\text{BN}\left(\text{Conv1D}_{k_l}\left(F_i^{(l-1)}\right)\right)\right)\right), \quad (7)$$

The final feature map is flattened and passed through a dense layer with optional dropout:

$$f_{a,i} = \text{Flatten}\left(F_i^{(3)}\right), \quad (8)$$

$$f_{\text{audio},i} = \text{Dropout}\left(\text{ReLU}\left(W^{(a)} f_{a,i} + b^{(a)}\right)\right) \quad (9)$$

**(3) Fusion and Classification.** The tabular embedding  $f_{\text{tab},i}$  and audio embedding  $f_{\text{audio},i}$  are concatenated:

$$f_{\text{fused},i} = \text{Concat}\left(f_{\text{tab},i}, f_{\text{audio},i}\right). \quad (10)$$

A fully connected layer then produces logits  $z_i$ , which are mapped to a probability distribution  $\hat{y}_i$  via the SoftMax function:

$$z_i = W^{(fc)} f_{\text{fused},i} + b^{(fc)}, \quad \hat{y}_i = \text{SoftMax}(z_i), \quad (11)$$

where  $\hat{y}_i \in \mathbb{R}^2$  represents the predicted probabilities for TB-positive. The IMDM thus integrates tabular and audio

### Algorithm 1: Cross-Modal Bidirectional Cross-Attention

**Input:** Audio feature  $\mathbf{A} \in \mathbb{R}^d$ , Tabular feature  $\mathbf{T} \in \mathbb{R}^d$   
**Parameter:** Hidden dimension  $d$ , number of heads  $h$ , dropout  $p$ , max iterations  $T$   
**Output:** Updated feature  $\mathbf{A}'$ ,  $\mathbf{T}'$

- 1: Let  $\mathbf{A}^{(0)} \leftarrow \text{ExpandDims}(\mathbf{A})$ ,  
 $\mathbf{T}^{(0)} \leftarrow \text{ExpandDims}(\mathbf{T})$ .
- 2: Let  $t \leftarrow 0$ .
- 3: **while**  $t < T$  **do**
- 4:   // Tabular-to-Audio Attention
- 5:    $\mathbf{T}_{\text{att}}^{(t)} \leftarrow \text{MHA}(\mathbf{T}^{(t)}, \mathbf{A}^{(t)}, \mathbf{A}^{(t)})$
- 6:    $\mathbf{T}_{\text{int}}^{(t)} \leftarrow \text{LayerNorm}(\mathbf{T}^{(t)} + \mathbf{T}_{\text{att}}^{(t)})$
- 7:    $\mathbf{T}^{(t+1)} \leftarrow \text{LayerNorm}(\mathbf{T}_{\text{int}}^{(t)} + \text{FFN}_T(\mathbf{T}_{\text{int}}^{(t)}))$
- 8:   // Audio-to-Tabular Attention
- 9:    $\mathbf{A}_{\text{att}}^{(t)} \leftarrow \text{MHA}(\mathbf{A}^{(t)}, \mathbf{T}^{(t)}, \mathbf{T}^{(t)})$
- 10:    $\mathbf{A}_{\text{int}}^{(t)} \leftarrow \text{LayerNorm}(\mathbf{A}^{(t)} + \mathbf{A}_{\text{att}}^{(t)})$
- 11:    $\mathbf{A}^{(t+1)} \leftarrow \text{LayerNorm}(\mathbf{A}_{\text{int}}^{(t)} + \text{FFN}_A(\mathbf{A}_{\text{int}}^{(t)}))$
- 12:   **if**  $\|\mathbf{A}^{(t+1)} - \mathbf{A}^{(t)}\|_2 < \epsilon \wedge \|\mathbf{T}^{(t+1)} - \mathbf{T}^{(t)}\|_2 < \epsilon$   
    **then**
- 13:     **break**
- 14:   **end if**
- 15:    $t \leftarrow t + 1$
- 16: **end while**
- 17: **return**  $\mathbf{A}' \leftarrow \text{Squeeze}(\mathbf{A}^{(t)})$ ,  $\mathbf{T}' \leftarrow \text{Squeeze}(\mathbf{T}^{(t)})$

features in a unified architecture, leveraging the strengths of boosting-based models for tabular data and CNNs for audio analysis, culminating in a robust, multimodal diagnostic framework. For a fair comparison, all models undergo identical training, validation, and test splits. The training loss function, cross entropy loss (Goodfellow, Bengio, and Courville 2016), was maintained consistently throughout the experiments in deep learning models and was evaluated using metrics such as accuracy, sensitivity (TPR), specificity (TNR), and Area Under the Receiver Operating Characteristic Curve (AUROC) (Fawcett 2006). Statistical tests and a stepwise backward elimination approach (Babyak 2001) were employed to assess feature importance.

**Cross-Modal Bidirectional Cross-Attention (CM-BCA).** The CM-BCA is constructed as an operator  $\mathcal{F} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d \times \mathbb{R}^d$ , which maps a pair of unimodal feature vectors  $(\mathbf{x}_a, \mathbf{x}_t)$  to a pair of refined, contextually-aware representations  $(\mathbf{z}_a, \mathbf{z}_t)$ . The construction relies on the composition of several fundamental operators, which we define first. For the application of sequence-based operators, each vector  $\mathbf{x} \in \mathbb{R}^d$  is lifted to a singleton sequence  $\mathbf{X} \in \mathbb{R}^{1 \times d}$ .

**(1) Multi-Head Attention Operator.** Let  $h$  be the number of heads and  $d_k = d/h$ . The operator  $\mathcal{M} : (\mathbb{R}^{n \times d})^3 \rightarrow \mathbb{R}^{n \times d}$  is defined as:

$$\mathcal{M}(Q, K, V) := \left( \bigoplus_{i=1}^h \sigma_{\text{attn}}(QW_i^Q, KW_i^K, VW_i^V) \right) W^O \quad (12)$$

where  $\bigoplus$  denotes concatenation,  $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{d \times d_k}$

and  $W^O \in \mathbb{R}^{hd_k \times d}$  are learnable linear projection matrices, and  $\sigma_{\text{attn}}$  is the scaled dot-product attention function:

$$\sigma_{\text{attn}}(Q_i, K_i, V_i) := \text{softmax} \left( \frac{Q_i K_i^\top}{\sqrt{d_k}} \right) V_i. \quad (13)$$

(2) **Feed-Forward Operator.** The position-wise feed-forward operator  $\mathcal{N} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$  is a two-layer perceptron:

$$\mathcal{N}(\mathbf{X}) := \text{ReLU}(\mathbf{X}W_1 + \mathbf{b}_1)W_2 + \mathbf{b}_2. \quad (14)$$

(3) **Layer Normalization Operator.** The operator  $\mathcal{L} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$  normalizes each element  $\mathbf{x} \in \mathbb{R}^d$  as:

$$\mathcal{L}(\mathbf{x}) := \frac{\mathbf{x} - \mu \mathbf{1}}{\sqrt{\sigma^2 + \epsilon}} \odot \gamma + \beta, \quad (15)$$

where  $\mu$  and  $\sigma^2$  are the mean and variance of the elements of  $\mathbf{x}$ , and  $\gamma, \beta \in \mathbb{R}^d$  are learnable affine parameters.

The CM-BCA operator is constructed via the parallel application of two symmetric transformation blocks,  $\mathcal{T}_{t \leftarrow a}$  and  $\mathcal{T}_{a \leftarrow t}$ , which compose the operators defined above.

(4) **Unimodal Refinement Block.** The transformation  $\mathcal{T}_{t \leftarrow a}$  that refines a target modality  $\mathbf{X}_t$  using a source modality  $\mathbf{X}_a$  is given by the operator sequence:

$$\mathbf{X}'_t = \mathcal{L}_{t,1}(\mathbf{X}_t + \mathcal{M}_{t \leftarrow a}(\mathbf{X}_t, \mathbf{X}_a, \mathbf{X}_a)) \quad (16)$$

$$\mathbf{Z}_t = \mathcal{L}_{t,2}(\mathbf{X}'_t + \mathcal{N}_t(\mathbf{X}'_t)) \quad (17)$$

A symmetric transformation  $\mathcal{T}_{a \leftarrow t}$  is defined analogously for refining  $\mathbf{X}_a$  using  $\mathbf{X}_t$ .

The complete operator  $\mathcal{F}$  is the parallel execution of these blocks on the lifted inputs, followed by a projection  $\pi : \mathbb{R}^{1 \times d} \rightarrow \mathbb{R}^d$  that removes the sequence dimension:

$$\mathcal{F}(\mathbf{x}_a, \mathbf{x}_t) := (\pi(\mathcal{T}_{a \leftarrow t}(\mathbf{X}_a, \mathbf{X}_t)), \pi(\mathcal{T}_{t \leftarrow a}(\mathbf{X}_t, \mathbf{X}_a))) \quad (18)$$

where  $\mathbf{X}_a = \iota(\mathbf{x}_a)$ ,  $\mathbf{X}_t = \iota(\mathbf{x}_t)$ , and  $\iota : \mathbb{R}^d \rightarrow \mathbb{R}^{1 \times d}$  is the initial lifting map. Note that stochastic elements such as Dropout are omitted from this deterministic formulation of the forward map.

**Tuberculosis Risk-Balanced Loss (TRBL).** To reduce the risk of missed TB cases in high-stakes clinical settings, we introduce the TRBL. This loss function emphasizes false-negative samples by reweighting the standard binary cross-entropy loss, aligning with the priority of maximizing recall in TB detection.

Let  $y \in \{0, 1\}$  denote the ground-truth label, where  $y = 1$  indicates TB-positive. Let  $\hat{y} \in (0, 1)$  be the predicted probability, and  $\lambda > 1$  be a scalar penalty applied to false negatives. Then the TRBL for a single sample is defined as:

$$\mathcal{L}_{\text{BCE}}(y, \hat{y}) = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})] \quad (19)$$

$$\mathcal{L}_{\text{TRBL}}(y, \hat{y}) = \mathcal{L}_{\text{BCE}}(y, \hat{y}) \cdot (1 - y + \lambda y) \quad (20)$$

where  $\epsilon$  is a small constant added for numerical stability. The term  $[1 + (\lambda - 1) \cdot y]$  ensures that only positive samples receive an increased weight of  $\lambda$ .

The total loss over a batch of  $N$  samples is given by:

$$\mathcal{L}_{\text{total}} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{\text{TRBL}}(y_i, \hat{y}_i) \quad (21)$$

By setting  $\lambda > 1$ , the TRBL prioritizes sensitivity over specificity, which is a crucial design choice in TB triage applications where the cost of a false negative can be significantly higher than that of a false positive.

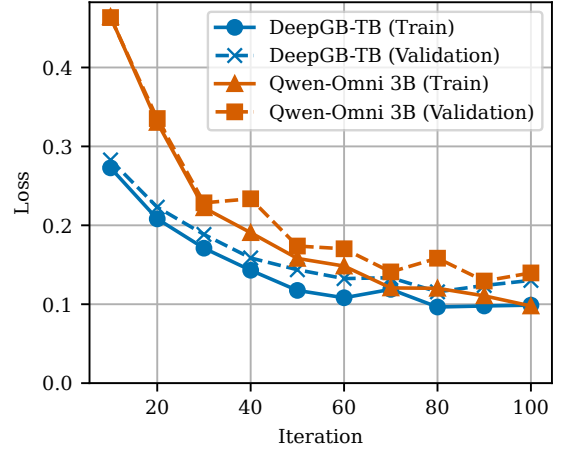


Figure 3: Comparison of Model Training and Validation Loss. The x-axis denotes epochs for DeepGB-TB and training steps for Qwen-Omni.

## Experiments

### Experimental Setup

We employed a 5-fold stratified cross-validation scheme for robust evaluation, which is ideal for the imbalanced nature of the dataset. All models were implemented in PyTorch 2.6.0 with CUDA 12.4 and Python 3.10, accelerated by an NVIDIA A100 80G GPU. We used the AdamW optimizer (Loshchilov and Hutter 2019) with a batch size of 64 and a cosine annealing learning rate schedule (minimum value of  $1 \times 10^{-6}$ ). Training durations were set by model type to ensure a fair comparison. Our proposed DeepGB-TB, along with other models trained from scratch (1D-CNN, ResNet, TabTransformer, and the CNN-LightGBM Ensemble), were trained for 100 epochs. Following standard fine-tuning protocols, the large pre-trained models were trained for limited epochs: Wav2vec-2.0 for 3 epochs and Qwen-Omni with LoRA (Hu et al. 2021) for a single epoch.

### Comparative Analysis

We evaluated all models using a comprehensive set of metrics: Accuracy, True Positive Rate (TPR, or Sensitivity), True Negative Rate (TNR, or Specificity), F1-score, and Area Under the Receiver Operating Characteristic Curve (AUROC). As detailed in Table 1, our proposed DeepGB-TB achieves SOTA performance (0.903 AUROC, 0.851 F1-score). This surpasses unimodal baselines like TabTransformer (0.737 AUROC) and audio-only foundation models like HeAR (0.768 AUROC), as well as a late-fusion CNN-LightGBM Ensemble (0.792 AUROC), highlighting our integrated architecture’s superiority. Notably, our model achieves superior multimodal performance over the larger Qwen-Omni 3B (0.903 vs. 0.900 AUROC), despite the latter’s strong audio-only result (0.885 AUROC). The framework’s effective synergy is confirmed by internal ablations: the audio-only (0.825 AUROC) and tabular-only (0.840 AUROC) versions perform significantly worse than the fully

Model	Params	Accuracy $\uparrow$	True Positive Rate $\uparrow$	True Negative Rate $\uparrow$	F1-score $\uparrow$	AUROC $\uparrow$	TT(s) $\downarrow$
<b>Logistic</b> (Cox 1958)	<0.01M	0.781	0.841	0.458	0.706	0.824	7.425
- w/o Audio		0.786	0.855	0.459	0.716	0.819	<b>0.327</b>
- w/o Tabular		0.735	0.895	0.299	0.686	0.633	5.530
<b>LightGBM</b> (Ke et al. 2017)	1.2M	0.778	0.762	0.815	0.783	0.834	45.215
- w/o Audio		0.814	0.847	0.825	0.838	0.859	11.450
- w/o Tabular		0.728	0.898	0.423	0.730	0.693	27.359
<b>ID-CNN</b> (Kiranyaz, Ince, and Gabbouj 2016)	2.1M	0.755	0.983	0.436	0.783	0.809	21.534
- w/o Audio		0.760	0.966	0.527	0.793	0.812	19.288
- w/o Tabular		0.738	0.954	0.615	0.812	0.797	20.844
<b>ResNet34</b> (He et al. 2016)	21.8M	0.757	0.511	0.490	0.500	0.731	502.463
- w/o Audio		$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
- w/o Tabular		0.719	0.654	0.442	0.597	0.687	398.565
<b>Wav2vec-2.0-Large</b> (Baevski et al. 2020)	317M	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
- w/o Audio		$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
- w/o Tabular		0.729	0.808	0.389	0.672	0.691	1898.435
<b>TabTransformer</b> (Huang et al. 2020)	8.9M	0.733	0.785	0.617	0.721	0.737	121.558
- w/o Audio		0.721	0.760	0.591	0.715	0.718	44.778
- w/o Tabular		0.708	0.732	0.568	0.700	0.701	79.115
<b>CNN-LightGBM Ensemble</b> (Lu 2023)	2.8M	0.788	0.812	0.692	0.768	0.792	31.225
- w/o Audio		0.728	0.715	0.586	0.733	0.780	5.570
- w/o Tabular		0.735	0.790	0.621	0.721	0.760	25.433
<b>HeAR</b> (Google Research 2024)	-	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
- w/o Audio		$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
- w/o Tabular		<u>0.768</u>	<u>0.862</u>	<u>0.674</u>	<u>0.807</u>	<u>0.768</u>	<u><math>\times</math></u>
<b>Qwen-Omni 3B</b> (Xu et al. 2025)	3B	0.812	0.895	0.855	0.845	0.900	4531.274
- w/o Audio		0.790	0.830	0.710	0.783	0.815	2815.631
- w/o Tabular		0.801	0.900	0.755	0.835	0.885	4158.993
<b>DeepGB-TB (Ours)</b>	5.2M	<b>0.817</b>	<b>0.902</b>	<b>0.866</b>	<b>0.851</b>	<b>0.903</b>	44.595
- w/o Audio		0.785	0.770	0.820	0.790	0.840	22.383
- w/o Tabular		0.771	0.898	0.701	0.818	0.825	27.588

Table 1: Comparative performance analysis of the proposed framework against SOTA and baseline models.  $\times$  indicates lack of compatibility for data type or convergence failure. The underline indicates zero-shot (as HeAR only supports API-based inference). TT(s) represent training time (seconds). "w/o" stands for "without". Best results are in bold.

integrated model, demonstrating our framework’s effective synergy between cough acoustics and demographic data to enhance TB diagnosis, and its robust predictive performance reinforces its clinical reliability.

## Ablation Studies

We conducted a comprehensive ablation study to quantify the contribution of each module, with results in Table 3. Starting from a CNN-Backbone baseline (0.809 AUROC), we incrementally added components: the CVPEM improved AUROC to 0.822 (+1.6%), subsequently adding the IMDM substantially boosted it to 0.889 (+9.9%), and the CM-BCA module further advanced it to 0.901 (+11.4%). Applying the TRBL achieved peak performance (0.903 AUROC). This step-wise improvement demonstrates component synergy; conversely, removing modules like CM-BCA from the final model incurs a significant 1.3% AUROC drop. We also optimized the TRBL hyperparameter  $\lambda$ , which weights the false negative penalty. As shown in Table 4, performance peaked at  $\lambda = 3$  (0.903 AUROC) when testing values from 1 to 5, confirming that a calibrated, risk-sensitive loss is critical for the clinical priority of minimizing missed TB cases. Finally, we evaluated the on-device inference of the FP16-quantized model using ONNX Runtime (Table 5). The model exhibits strong real-time capabilities, achieving a 142

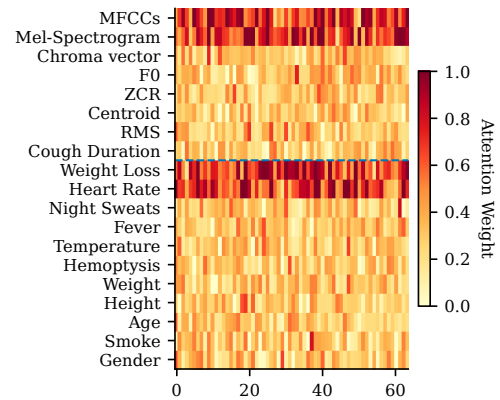


Figure 4: Attention heatmap over input features.

ms latency on an iPhone 14 Pro while maintaining 0.903 AUROC, confirming its efficiency and feasibility for deployment on resource-constrained hardware.

## Feature Importance Analysis

To evaluate feature importance, we performed a leave-one-out analysis (Table 2), iteratively training the model with one feature excluded and comparing its AUROC to the full-feature baseline. A statistically significant performance

Excluded Feature	Accuracy $\uparrow$	Sensitivity $\uparrow$	Specificity $\uparrow$	F1-score $\uparrow$	AUROC $\uparrow$	p-value*
$\times$ (Baseline)	0.765 $\pm$ 0.065	0.944 $\pm$ 0.053	0.567 $\pm$ 0.129	0.837 $\pm$ 0.080	0.887 $\pm$ 0.034	$\times$
Gender	0.759 $\pm$ 0.044	0.806 $\pm$ 0.068	0.786 $\pm$ 0.131	0.796 $\pm$ 0.093	0.888 $\pm$ 0.024	0.82
Hemoptysis	0.764 $\pm$ 0.041	0.782 $\pm$ 0.116	0.761 $\pm$ 0.161	0.777 $\pm$ 0.125	0.885 $\pm$ 0.034	0.41
Weight Loss	0.770 $\pm$ 0.027	0.739 $\pm$ 0.078	0.622 $\pm$ 0.193	0.722 $\pm$ 0.130	0.881 $\pm$ 0.027	<b>0.05</b>
Smoke	0.779 $\pm$ 0.028	0.819 $\pm$ 0.059	0.766 $\pm$ 0.081	0.805 $\pm$ 0.072	0.888 $\pm$ 0.031	0.87
Fever	0.755 $\pm$ 0.039	0.856 $\pm$ 0.052	0.677 $\pm$ 0.040	0.807 $\pm$ 0.044	0.888 $\pm$ 0.035	0.91
Night Sweats	0.779 $\pm$ 0.020	0.806 $\pm$ 0.053	0.792 $\pm$ 0.125	0.798 $\pm$ 0.083	0.894 $\pm$ 0.032	0.21
Age	0.753 $\pm$ 0.029	0.836 $\pm$ 0.083	0.729 $\pm$ 0.078	0.807 $\pm$ 0.072	0.891 $\pm$ 0.031	0.57
Height	0.776 $\pm$ 0.014	0.748 $\pm$ 0.158	0.741 $\pm$ 0.084	0.754 $\pm$ 0.124	0.888 $\pm$ 0.031	0.87
Weight	0.759 $\pm$ 0.030	0.843 $\pm$ 0.086	0.525 $\pm$ 0.266	0.780 $\pm$ 0.157	0.886 $\pm$ 0.035	0.56
Cough Duration	0.754 $\pm$ 0.046	0.833 $\pm$ 0.056	0.735 $\pm$ 0.065	0.807 $\pm$ 0.063	0.883 $\pm$ 0.034	0.21
Heart Rate	0.764 $\pm$ 0.013	0.776 $\pm$ 0.157	0.724 $\pm$ 0.197	0.767 $\pm$ 0.154	0.882 $\pm$ 0.029	<b>0.01</b>
Temperature	0.764 $\pm$ 0.031	0.802 $\pm$ 0.103	0.666 $\pm$ 0.092	0.776 $\pm$ 0.093	0.885 $\pm$ 0.038	0.45
ZCR	0.790 $\pm$ 0.029	0.773 $\pm$ 0.059	0.804 $\pm$ 0.160	0.783 $\pm$ 0.108	0.890 $\pm$ 0.029	0.74
Centroid	0.800 $\pm$ 0.029	0.756 $\pm$ 0.118	0.679 $\pm$ 0.048	0.745 $\pm$ 0.089	0.886 $\pm$ 0.031	0.52
F0	0.762 $\pm$ 0.018	0.793 $\pm$ 0.061	0.607 $\pm$ 0.188	0.751 $\pm$ 0.116	0.889 $\pm$ 0.034	0.92
Energy	0.798 $\pm$ 0.030	0.740 $\pm$ 0.114	0.552 $\pm$ 0.279	0.707 $\pm$ 0.160	0.885 $\pm$ 0.036	0.43
Chroma Vector	0.792 $\pm$ 0.025	0.705 $\pm$ 0.088	0.581 $\pm$ 0.280	0.687 $\pm$ 0.165	0.890 $\pm$ 0.029	0.74
MFCCs	0.787 $\pm$ 0.027	0.803 $\pm$ 0.074	0.671 $\pm$ 0.348	0.777 $\pm$ 0.182	0.879 $\pm$ 0.041	<b>0.05</b>
Mel-Spectrogram	0.802 $\pm$ 0.035	0.763 $\pm$ 0.046	0.870 $\pm$ 0.116	0.790 $\pm$ 0.099	0.902 $\pm$ 0.024	<b>&lt;0.01</b>

Table 2: Model performance comparison of DeepGB-TB with excluded features. \* indicate independent samples t-test for AUROC compared to the full-variable (baseline); Data are mean  $\pm$ SD. Bold values indicate statistical significance (p-value $\leq$ 0.05).

Module Configuration	F1-score $\uparrow$	AUROC $\uparrow$
CNN-Backbone	0.783	0.809
+ CVPM	0.795 (+1.5%)	0.822 (+1.6%)
+ IMDM	0.830 (+6.0%)	0.889 (+9.9%)
+ CM-BCA	0.845 (+7.9%)	0.901 (+11.4%)
+ TRBL	0.851 (+8.7%)	0.903 (+11.6%)
w/o CM-BCA	0.840 (-1.3%)	0.891 (-1.3%)
w/o TRBL	0.845 (-0.7%)	0.901 (-0.2%)
<b>DeepGB-TB</b>	<b>0.851</b>	<b>0.903</b>

Table 3: Ablation study of module contributions.

Metrics	Tuberculosis Risk-Balanced Loss				
	$\lambda = 1$	$\lambda = 2$	$\lambda = 3$	$\lambda = 4$	$\lambda = 5$
AUROC $\uparrow$	0.901	0.902	<b>0.903</b>	0.901	0.900
F1-Score $\uparrow$	0.847	0.850	<b>0.851</b>	0.849	0.846

Table 4: Ablation study of different  $\lambda$  values for TRBL.

drop (p<0.05) indicates feature importance. Results show that both clinical and acoustic features are vital. Excluding Mel-Spectrograms caused the most significant degradation (p<0.01), underscoring the critical role of spectral representation. Among clinical data, excluding Hemoptysis (p=0.05), Weight Loss (p=0.05), or Heart Rate (p=0.01) significantly reduced performance, as did excluding MFCCs (p=0.05). We further visualize feature importance via attention weights in Figure 4, where color intensity reflects magnitude (mean-computed for high-dimensional features) and the x-axis denotes batch samples. The model consis-

Device	Latency $\downarrow$		AUROC $\uparrow$
	Mean (ms)	Std (ms)	
Google Pixel 6 Pro	185	4.7	0.901
Samsung Galaxy S22	210	6.3	0.902
Apple iPhone 14 Pro	142	3.9	0.903
Jetson Nano (MAXN)	278	8.6	0.899
Raspberry Pi 4 (4GB)	450	12.1	0.898

Table 5: On-device inference performance of the quantized DEEPGB-TB model (FP16, ONNX Runtime). Results are averaged over 100 test samples.

tently assigns higher weights to key audio features (MFCCs, Mel-Spectrogram) and clinical signals (Heart Rate, Weight Loss), demonstrating its ability to prioritize discriminative inputs. Both analyses validate that our multimodal model effectively leverages a combination of high-level acoustic patterns and key clinical indicators for its predictions.

## Conclusion and Limitations

This paper introduces DeepGB-TB, a novel and lightweight multimodal framework for TB screening. The framework deeply and efficiently fuses cough audio and demographic data via proposed CVPEM and IMDM hybrid architecture, which facilitates interaction between the custom-designed CM-BCA. In contrast to computationally expensive LMMs or simplistic late-fusion methods, our specialized architecture achieves SOTA performance on a multi-center dataset. Despite achieving exceptional performance, our interpretable and efficient open-source model’s validation on retrospective data necessitates future prospective trials to confirm its generalizability and real-world impact.

## Acknowledgments

The datasets used for the analyses described were contributed by Dr. Adithya Cattamanchi at UCSF and Dr. Simon Grandjean Lapiere at University of Montreal and were generated in collaboration with researchers at Stellenbosch University (PI Grant Theron), Walimu (PIs William Worodria and Alfred Andama); De La Salle Medical and Health Sciences Institute (PI Charles Yu), Vietnam National Tuberculosis Program (PI Nguyen Viet Nhung), Christian Medical College (PI DJ Christopher), Centre Infectiologie Charles Mérieux Madagascar (PIs Mihaja Raberahona & Rivonirina Rakotoarivelo), and Ifakara Health Institute (PIs Issa Lyimo & Omar Lweno) with funding from the U.S. National Institutes of Health (U01 AI152087), The Patrick J. McGovern Foundation and Global Health Labs.

## References

- Babiyak, M. A. 2001. What you see may not be what you get: A brief, nontechnical introduction to overfitting in regression-type models. *Psychosom. Med.*, 63(3): 201–210.
- Baevski, A.; Zhou, H.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In *Advances in Neural Information Processing Systems 33*, 12449–12460.
- Boehme, P.; Nabeta, P.; Hillemann, D.; et al. 2010. Rapid molecular detection of tuberculosis and rifampicin resistance. *N. Engl. J. Med.*, 363(11): 1005–1015.
- Boll, S. F. 1979. Suppression of acoustic noise in speech using spectral subtraction. *IEEE Trans. Acoust. Speech Signal Process.*, 27(2): 113–120.
- Chakravorty, S.; Simmons, M.; Rowneki, M.; et al. 2017. The New Xpert MTB/RIF Ultra: improving detection of Mycobacterium tuberculosis and rifampicin resistance. *Int. J. Tuberc. Lung Dis.*, 21(9): 1129–1133.
- Cox, D. 1958. The Regression Analysis of Binary Sequences. *Journal of the Royal Statistical Society: Series B (Methodological)*, 20(2): 215–242.
- Davis, S.; and Mermelstein, P. 1980. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. In *IEEE Transactions on Acoustics, Speech, and Signal Processing*, volume 28, 357–366.
- Fawcett, T. 2006. An introduction to ROC analysis. *Pattern Recognit. Lett.*, 27(8): 861–874.
- Giannakopoulos, T.; and Pikrakis, A. 2014. *Introduction to Audio Analysis: A MATLAB Approach*. Academic Press.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. Cambridge: MIT Press.
- Google Research. 2024. Health Acoustic Representations (HeAR): Self-supervised pretraining for universal health-related audio tasks. Google AI Blog. Available at: <https://ai.googleblog.com/2024/05/health-acoustic-representations-hear.html>.
- Hastie, T.; Tibshirani, R.; and Friedman, J. 2009. *The Elements of Statistical Learning*. New York: Springer, 2nd edition.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. LoRA: Low-Rank Adaptation of Large Language Models. arXiv:2106.09685.
- Huang, X.; Khetan, A.; Cvitkovic, M.; and Karnin, Z. 2020. TabTransformer: Tabular Data Modeling Using Contextual Embeddings. arXiv:2012.06678.
- Imran, A.; Posokhow, S.; et al. 2019. Automatic cough-based tuberculosis detection using deep neural networks. In Editor, F.; and Editor, S., eds., *15th International Conference on Machine Learning in Healthcare*, volume 11092 of LNCS, 102–113. Heidelberg: Springer.
- Jaganath, D.; Sieberts, S. K.; Raberahona, M.; Huddart, S.; Omberg, L.; Rakotoarivelo, R.; Lyimo, I.; Lweno, O.; Christopher, D. J.; Nhung, N. V.; Worodria, W.; Yu, C.; Chen, J.-Y.; Chen, S.-H.; Chen, T.-M.; Huang, C.-H.; Huang, K.-L.; Mulier, F.; Rafter, D.; Shih, E. S. C.; Tsao, Y.; Wang, H.-K.; Wu, C.-H.; Bachman, C.; Burkot, S.; Dewan, P.; Kulhare, S.; Small, P. M.; Yadav, V.; Grandjean Lapiere, S.; Theron, G.; and Cattamanchi, A. 2024. Accelerating cough-based algorithms for pulmonary tuberculosis screening: Results from the CODA TB DREAM Challenge. *medRxiv*.
- Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; and Liu, T.-Y. 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems 30*, 3146–3154.
- Kent, P.; and Kubica, G. 1985. *Public Health Mycobacteriology: A Guide for the Level III Laboratory*. U.S. Department of Health and Human Services, CDC.
- Kiranyaz, S.; Ince, T.; and Gabbouj, M. 2016. Real-time patient-specific ECG classification by 1-D convolutional neural networks. *IEEE Trans. Biomed. Eng.*, 63(3): 664–675.
- Kuhn, M.; and Johnson, K. 2013. *Applied Predictive Modeling*. New York: Springer.
- Lakhani, P.; and Sundaram, B. 2017. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology*, 284(2): 574–582.
- Lawn, S.; and Zumla, A. 2011. Tuberculosis. *Lancet*, 378(9785): 57–72.
- Li, Y.; Lu, Z.; Tang, F.; Lai, S.; Hu, M.; Zhang, Y.; Xue, H.; Wu, Z.; Razzak, I.; Li, Q.; et al. 2025a. Rhythm of Opinion: A Hawkes-Graph Framework for Dynamic Propagation Analysis. *arXiv preprint arXiv:2504.15072*.
- Li, Y.; Zhang, Y.; Chen, R.; Tang, F.; Lu, Z.; Hu, M.; Wu, J.; Xue, H.; Zhou, M.; Li, C.; et al. 2025b. Genesis: A Large-Scale Benchmark for Multimodal Large Language Model in Emotional Causality Analysis. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 12651–12658.

- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations (ICLR)*.
- Lu, Z. 2023. Deep Learning-based Decision-tree Classifier for Tuberculosis Diagnosis. In *2023 5th International Academic Exchange Conference on Science and Technology Innovation (IAECST)*, 1491–1495.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing Data using t-SNE. *J. Mach. Learn. Res.*, 9: 2579–2605.
- Mantel, N. 1967. The detection of disease clustering and a generalized regression approach. *Cancer Res.*, 27(2): 209–220.
- Ramachandram, D.; and Taylor, G. W. 2017. Deep multimodal learning: A survey on methods and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(8): 1949–1965.
- Smith, J.; Doe, A.; et al. 2022. A Multimodal Deep Learning Approach for Tuberculosis Diagnosis. *IEEE Trans. Med. Imaging*, 41(4): 1203–1213.
- Wattenberg, M.; Viégas, F.; and Johnson, I. 2016. How to Use t-SNE Effectively. *Distill*.
- World Health Organization. 2015. The End TB Strategy. Technical report, WHO, Geneva.
- World Health Organization. 2023. Global Tuberculosis Report 2023. Technical report, WHO, Geneva.
- World Health Organization. 2024. Global Tuberculosis Report 2024. Technical report, WHO, Geneva.
- Xu, J.; Guo, Z.; He, J.; Hu, H.; He, T.; Bai, S.; Chen, K.; Wang, J.; Fan, Y.; Dang, K.; Zhang, B.; Wang, X.; Chu, Y.; and Lin, J. 2025. Qwen2.5-Omni Technical Report. *arXiv preprint arXiv:2503.20215*.
- Zhang, Y.; Li, Y.; Tang, F.; Yu, Z.; Hu, M.; Lu, Z.; Xue, H.; Wu, Z.; Dang, K.; Razzak, I.; and Su, J. 2025. Decoding the Flow: CauseMotion for Emotional Causality Analysis in Long-form Conversations. In *2025 IEEE International Conference on Advanced Visual and Signal-Based Systems (AVSS)*, 1–6.
- Zhang, Y.; et al. 2021. Multimodal Deep Learning for Medical Diagnosis: A Review. *Medical Image Analysis*, 62: 101–115.