

Clinically Constrained Multi-Stage Architecture for Integrated Tuberculosis Risk, Drug-Resistance and Severity Assessment

Kavitha K. ^{*}, Subrahmanyam R., Asritha Sai P., Zaheer P., and Santhosh Kumar M.

Department of Artificial Intelligence and Machine Learning, School of Computing, Mohan Babu University, Tirupati, India

^{*}Corresponding author: kavi.mca.mba@gmail.com

Abstract. Tuberculosis (TB) remains a major global health burden because of high infection and mortality rates, delays in diagnosis, the emergence of drug-resistant strains, and the wide variation in how the disease presents. Diagnosis of TB, estimation of the risk of multidrug resistance (MDR), and early identification of patients with severe illness are usually carried out as separate tasks. This study proposes a clinical decision support system that detects TB, estimates MDR-TB risk, and grades disease severity through a single multi-stage pipeline, with each stage's output checked against rules set by clinicians. The system was evaluated with stratified 5-fold cross-validation on harmonized public clinical datasets. It achieved an area under the receiver operating characteristic curve (AUROC) of 0.91 for TB detection and 0.83 for MDR-TB risk estimation, and a quadratic weighted kappa (QWK) of 0.74 for severity grading, outperforming logistic regression, random forest, and multi-task neural network baselines on all three tasks. The results indicate that a unified, explainable decision support system is a feasible and practical tool for the combined assessment of patients with TB.

Keywords. Tuberculosis; clinical decision support system; machine learning; MDR-TB risk; disease severity prediction; healthcare AI

I. INTRODUCTION

Tuberculosis (TB) is a significant public health problem that affects millions of people each year, mainly in developing nations. Multidrug-resistant tuberculosis (MDR-TB) has made the diagnosis and treatment of the disease more complex, so timely diagnosis and sound risk assessment are central to disease management.

Recent work has shown considerable promise for artificial intelligence (AI) and machine learning (ML) in TB diagnosis and clinical decision support [1], [2]. Several models have been proposed to estimate mortality risk and severity in TB patients [3], and clinical scoring systems have been developed to support patient management [4]. Most existing approaches, however, address a single prediction task, which leaves no unified framework for TB detection, drug resistance, and severity.

This study addresses that gap with a clinically constrained multi-stage architecture for integrated TB assessment. TB risk prediction, MDR-TB risk estimation, and severity assessment are combined in one clinical decision support pipeline. The system pairs ML models with clinically guided constraints to give more accurate and reliable support for TB diagnosis and management.

II. RELATED WORK

Recent studies have applied ML and AI to TB diagnosis, drug-resistance prediction, and clinical outcome assessment. Explainable AI models have improved the interpretability of MDR prediction [5], and deep learning mod-

els have been developed to capture nonlinear clinical patterns for resistance prediction in critical care settings [6]. Combining imaging with clinical parameters has improved prognostic prediction in TB [7], and ML applied to genomic data has been used to identify patterns of MDR-TB resistance [8]. Systematic reviews confirm that ML is increasingly adopted for predicting TB drug resistance [9], although most studies address individual tasks and use heterogeneous dataset types. Genome-based methods give a mutation-level view of resistance but require considerably more effort [10]. Existing systems therefore treat diagnosis, resistance, and severity separately, which points to the need for a unified decision support system. Table 1 summarizes these approaches.

III. DATASET DESCRIPTION

This study uses several publicly available clinical datasets. Features were harmonized systematically. Features with the same clinical meaning were mapped to a common name across sources, and TB and MDR labels were binarized. Distribution shift between sources was tested with the Kolmogorov-Smirnov or chi-squared test, and divergent features ($p < 0.05$) were z-score normalized within each source so that values were comparable. Redundant features with a Pearson correlation coefficient above 0.90 were removed.

The main dataset is the Kaggle Tuberculosis Patient Dataset, which contains approximately 3,200 patient records with 22 structured clinical features covering demographics, symptom duration, treatment duration, anthropo-

metric measures, and clinical outcomes. It is the primary source for training the TB risk prediction model.

The Kaggle Meningitis Clinical Dataset adds diversity to the clinical features. It contains about 750 patient records with 18 clinical and cerebrospinal fluid (CSF) attributes, including neurological signs (headache, vomiting, and neck stiffness) and laboratory values (glucose, protein, total cell count, and lymphocyte percentage).

Regional TB data from the World Health Organization (WHO) and India's National Tuberculosis Elimination Programme (NTEP), about 1,200 records, provide epidemio-

logical context. These data report TB cases, MDR-TB counts, treatment success, and relapse rates by region.

To strengthen severity and mortality risk modeling, clinical data from the UCI Machine Learning Repository were added, comprising more than 1,500 patient records with vital signs and laboratory values such as hemoglobin, blood pressure, blood glucose, and pulse rate. Together, these datasets describe the clinical, epidemiological, and physiological characteristics that the multi-stage TB prediction framework requires.

TABLE 1. SUMMARY OF EXISTING MACHINE LEARNING APPROACHES FOR TB AND DRUG-RESISTANCE PREDICTION

Ref.	Method	Focus area	Key contribution
[5]	ML with explainable AI	MDR infection prediction	Better prediction and interpretability of resistance risk
[6]	Deep learning	MDR infection in ICU patients	Captures complex clinical patterns for resistance prediction
[7]	Imaging and clinical data	TB prognosis	Combines radiological and clinical data for outcome prediction
[8]	ML on genomic and clinical data	MDR-TB prediction	Identifies resistance patterns in genomic datasets
[9]	Systematic review of ML models	TB drug resistance	Sets out strengths and limits of ML approaches
[10]	Genome-based ML	Mutation detection	Predicts resistance-associated genetic mutations

IV. METHODOLOGY

The proposed system follows a clinically constrained multi-stage architecture for integrated TB assessment. Its objective is to predict TB probability, MDR-TB risk, and disease severity from structured clinical data within a unified clinical decision support framework.

As Fig. 1 shows, the process begins with patient data covering clinical history, demographics, clinical signs, and laboratory results. In the preprocessing and feature engineering layer, the data are standardized, missing values are handled, and clinical features are engineered. A rule-based

clinical logic layer then applies physiological consistency checks and removes biologically implausible values.

Three predictive models follow. The TB clinical decision support model uses clinical symptoms and features to predict the likelihood of TB and its clinical category. The MDR risk estimation model uses this output, together with clinical variables and regional epidemiological factors, to estimate the probability of MDR-TB. The disease severity estimation model then grades severity from vital signs, laboratory parameters, and the outputs of the earlier models.

These outputs are combined through the clinical logic layer into a single clinical decision report that gives the TB probability, the MDR risk level, the severity class, the factors contributing to severity, and suggested next steps.




Fig. 1. Proposed integrated architecture for TB risk prediction, MDR risk estimation, and severity assessment.

For reproducibility, the model settings were as follows. TB detection used XGBoost (300 estimators, maximum depth 6, learning rate 0.05, 28 features). MDR estimation used L2-regularized logistic regression ($C = 0.1$, 19 features). Severity estimation used an ordinal random forest (200 trees, minimum leaf size 5). The implementation used Python 3.10 with scikit-learn 1.3 and XGBoost 2.0. Z-score normalization and median or mode imputation were computed on the training folds only, to avoid data leakage.

V. EXPERIMENTAL RESULTS

The system was evaluated on all three tasks with stratified 5-fold cross-validation, and the results are given in Table 2. The proposed approach achieved an AUROC of 0.91 for TB detection (0.03 above the next-best model), an AUROC of 0.83 for MDR estimation (0.05 above random forest), and a QWK of 0.74 for severity (0.08 above random forest). The consistent gain across all tasks supports the value of staged modeling with clinical rule constraints. The tree-based models also outperform logistic regression because clinical features interact nonlinearly.

TABLE 2. PERFORMANCE OF THE PROPOSED SYSTEM AGAINST BASELINE MODELS

Model	TB AUROC	MDR AUROC	Severity QWK
Logistic regression	0.84	0.72	0.58
Random forest	0.86	0.78	0.66
Multi-task neural net	0.88	0.76	0.63
Proposed system	0.91	0.83	0.74

VI. EVALUATION METRICS

Model performance was evaluated with standard classification metrics, defined in (1) to (4), where TP , TN , FP , and FN denote true positives, true negatives, false positives, and false negatives.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (1)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (2)$$

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (3)$$

$$F_1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (4)$$

Overall discrimination was measured with the AUROC. Severity prediction was evaluated with the QWK, which measures agreement between predicted and actual ordinal classes.

VII. RESULTS AND ANALYSIS

The system was assessed on three tasks (TB detection, MDR risk estimation, and severity prediction) using the metrics above and visual analysis.

The confusion matrix of the TB clinical decision model (Fig. 2) shows that the model correctly classified 440 TB cases and 420 non-TB cases, while it misclassified 85 non-TB cases as TB and 55 TB cases as non-TB. This indicates good sensitivity and specificity for TB detection, with a small number of misclassifications.




Fig. 2. Confusion matrix of the TB clinical decision model.

The severity estimation model was evaluated with an ordinal confusion matrix (Fig. 3). It correctly predicted 210 mild, 180 moderate, and 160 severe cases. Most misclassifications fall between adjacent severity levels, such as mild predicted as moderate or moderate predicted as severe. This pattern is expected in ordinal classification and shows that the model grades severity consistently.

The calibration curve (Fig. 4) closely follows the ideal diagonal, which shows that the predicted probabilities are well calibrated and clinically usable. The multi-stage architecture provides reliable TB detection, effective severity classification, and well-calibrated probability outputs for clinical decision-making.




Fig. 3. Confusion matrix of the severity estimation model (ordinal classification).

VIII. CONCLUSION AND LIMITATIONS

The proposed multi-stage approach to integrated TB assessment achieved an AUROC of 0.91 for TB detection and 0.83 for MDR risk estimation, and a QWK of 0.74 for severity classification, outperforming all baselines. Staged modeling grounded in clinical knowledge, together with clinical rule constraints, gives more reliable predictions and supports transparent clinical decision-making.

The study has several limitations. Validation was limited to internal 5-fold cross-validation. External validation on an independent prospective cohort, with a target of at least 500 patients at regional TB clinics, is required before clinical deployment and is the main direction for future work. Public data sources may not represent all epidemiological

settings, and residual distribution differences may remain after harmonization. The system also lacks longitudinal follow-up and microbiological confirmation data, which could improve MDR stratification. A formal user study with clinicians to evaluate real-world acceptability is planned as the next step.




Fig. 4. Calibration curve of the TB probability model.

AUTHOR CONTRIBUTIONS

All authors contributed to the conception and preparation of this work and approved the final version of the manuscript.

FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

DECLARATION OF GENERATIVE AI USE

The authors declare that no generative artificial intelligence tools were used in the preparation of this manuscript. The authors take full responsibility for the content of the publication.

DATA AVAILABILITY

This study used publicly available data: the Kaggle Tuberculosis Patient Dataset and the Kaggle Meningitis Clinical Dataset (Kaggle), regional tuberculosis data from the World Health Organization (WHO) and India's National Tuberculosis Elimination Programme (NTEP), and clinical data from the UCI Machine Learning Repository.

ETHICS STATEMENT

This study used publicly available, de-identified data and did not involve direct interaction with human participants or animals; ethical approval was therefore not required.

CONFLICTS OF INTEREST

The authors declare no conflict of interest.

REFERENCES

- [1] Sameeullah Memon, Shabana Bibi, and Guozhong He, "Integration of AI and ML in tuberculosis (TB) management: From diagnosis to drug discovery," *Diseases*, vol. 13, no. 6, Art. no. 184, 2025, doi: 10.3390/diseases13060184.
- [2] Shuanglin Gu, Yichen Guo, Han Jiang, Chen Yang, Tengyue Li, Xiuchang Ma, Wenxian Qian, Zihao Liang, Zheng Hu, Apeng Chen, Yi Zeng, and Changhua Yi, "Machine learning of immune-related laboratory indicators enables discrimination between pulmonary and extrapulmonary tuberculosis," *Front. Med.*, vol. 12, Art. no. 1712073, 2025, doi: 10.3389/fmed.2025.1712073.
- [3] Joseph Baruch Baluku, Priscilla Sheilla Apolot, Brenda Namanda, Sharon Namiro, Shamim Katusabe, Diana Karungi, Reagan Nkonge, Mary Madalen Angut, Jasper Nidoi, Robinah Nalwanga, Charles Mondo, Emmanuel Seremba, and Charles Kabugo, "A predictive score for early in-patient tuberculosis mortality: A case-control study," *J. Clin. Tuberc. Other Mycobact. Dis.*, vol. 37, Art. no. 100487, 2024, doi: 10.1016/j.jctube.2024.100487.
- [4] Pattama Saisudjarit, Surasak Saokaew, Acharaporn Duangjai, Anurak Prasatkhetragarn, Sukrit Kanchanasurakit, and Pochamana Phisalprapa, "Development and validation of clinical prediction score for mortality in tuberculosis patients," *Narra J.*, vol. 5, no. 2, Art. no. e1701, 2025, doi: 10.52225/narra.v5i2.1701.
- [5] Wenting Zhao, Pei Sun, Wei Li, and Linping Shang, "Machine learning-based prediction model for multidrug-resistant organisms infections: Performance evaluation and interpretability analysis," *Infect. Drug Resist.*, vol. 18, pp. 2255-2269, 2025, doi: 10.2147/IDR.S459830.
- [6] Yaxi Wang, Gang Wang, Yuxiao Zhao, Cheng Wang, Chen Chen, Yaoyao Ding, Jing Lin, Jingjing You, Silong Gao, and Xufeng Pang, "A deep learning model for predicting multidrug-resistant organism infection in critically ill patients," *J. Intensive Care*, vol. 11, Art. no. 49, 2023, doi: 10.1186/s40560-023-00695-y.
- [7] Liane S. Canas, Trinh H. K. Dong, Daniel Beasley, Joseph Donovan, Jon O. Cleary, Richard Brown, Nguyen Thuy Thuong Thuong, Phu Hoan Nguyen, Ha Thi Nguyen, Reza Razavi, Sebastien Ourselin, Guy E. Thwaites, and Marc Modat, "Computer-aided prognosis of tuberculous meningitis combining imaging and non-imaging data," *Sci. Rep.*, vol. 14, Art. no. 17581, 2024, doi: 10.1038/s41598-024-68308-8.
- [8] Komal Saxena, S. Shyni Carmel Mary, Prolay Ghosh, Chinnem Rama Mohan, Mohd Naved, and Mukesh Madanan, "Predicting multi-drug resistant tuberculosis using machine learning on genomic and clinical data," *Indian J. Tuberc.*, vol. 72, suppl. 3, pp. S99-S105, 2025, doi: 10.1016/j.ijtb.2025.11.013.
- [9] Bright G. Akwaronwu, Folasade Y. Ayankoya, Shade O. Kuyoro, Innocent Uche Akwaronwu, and Oluwabamise Joseph Adeniyi, "Machine learning for predicting drug resistance in tuberculosis: A systematic review of model performance and clinical applications," *Big Data Min. Anal.*, vol. 9, no. 1, pp. 314-337, 2026, doi: 10.26599/BDMA.2025.9020063.
- [10] Ankita Pal and Debasisa Mohanty, "Machine learning-based approach for identification of new resistance associated mutations from whole genome sequences of *Mycobacterium tuberculosis*," *Bioinform. Adv.*, vol. 5, no. 1, Art. no. vbaf050, 2025, doi: 10.1093/bioadv/vbaf050.