Explainable AI for early detection of Antimicrobial Resistance: An Approach to Reduce Clinical Risk

Mariana Timbó de Oliveira[1], Israel Martins Valcareggi[1], Jean Paes Landim de Lucena[1], Marise Reis de Freitas[2], Ícaro Boszczowski[3], Patrick Terrematte[1]

[1] Metropolis Digital Institute (IMD)
Federal University of Rio Grande do Norte (UFRN)
Natal, RN, Brazil

[2] Department of Infectious Diseases, Institute of Tropical Medicine (IMT), UFRN
Natal, RN, Brazil

[3] Hospital das Clínicas da Faculdade de Medicina, University of São Paulo (USP)
São Paulo, SP, Brazil

{mariana.timbo.114, israel.valcareggi.130}@ufrn.edu.br,
jeanlucena@gmail.com, icaro.b@hc.fm.usp.br
{marise.freitas, patrick.terrematte}@ufrn.br

Abstract. *The appropriate prescription of antimicrobials is critical for the clinical outcome of patients with infectious diseases. However, the advancement of Antimicrobial Resistance (AMR) contributes to therapeutic ineffectiveness and worse prognoses. In this scenario, machine learning models have emerged as a promising approach to support the early assessment of AMR risk and assist clinical decision-making. Based on hospital admission records from the MIMIC-IV database, this work evaluates supervised machine learning models for predicting antimicrobial resistance. The features used in this study include demographic, clinical, microbiological, and derived variables. The target variable represents whether the bacterial isolate was resistant or susceptible to the tested antimicrobial. Random Forest, Logistic Regression, and LightGBM were evaluated. To improve model performance, class imbalance strategies and hyperparameter optimization with Optuna were applied. In addition, explainable AI methods were used to interpret the model outputs through SHapley Additive exPlanations (SHAP). The Random Forest achieved the best result, with an AUC of 0.88, approximately. The results suggest that the proposed approach has potential to support AMR risk assessment within the MIMIC-IV context, while further external and clinical validation remains necessary.*

Keywords. *Antimicrobial Resistance. Machine Learning. Explainable AI. SHAP. MIMIC-IV. Random Forest.*

1. Introduction

Antimicrobial Resistance (AMR) represents one of the greatest health challenges worldwide, ranking among the top 10 causes of death globally [Ministry of Health (Brazil) 2026]. The frequent and inappropriate use of antibiotics, as well as improper prescription in relation to the clinical condition, are the main factors contributing to the development of drug resistance. This phenomenon

directly impacts treatment effectiveness, contributing to poorer patient prognosis and an increased risk of negative outcomes, and is responsible for 1.27 million deaths annually [World Health Organization 2023]. In addition, there is a significant economic impact, as it contributes to longer treatment durations, making them more complex, costly, and demanding.

Resistance detection can be performed through microbiological culture testing; however, this process is time-consuming and impractical for emergency treatments and everyday prescriptions. Therefore, it is important to develop solutions that can assist in medical decision-making before antimicrobial susceptibility results become fully available. In this context, the use of supervised machine learning techniques shows promise, allowing for the prediction of the risk of bacterial resistance to certain drugs and the analysis of patterns in clinical data. This approach has the potential to contribute to more accurate decisions, reducing therapeutic failures and assisting in the more appropriate use of antimicrobials.

The remainder of this paper is organized as follows. Section 2 presents the related work. Section 3 presents the methodology, including the computational environment, dataset description, machine learning models, hyperparameter optimization, imbalance handling strategies, and explainability techniques. Section 4 presents the experimental results. Section 5 discusses the findings, including model performance, explainability analysis, limitations, and future research directions. Finally, Section 6 concludes the paper.

2. Related Work

The application of machine learning (ML) techniques to predict antimicrobial resistance has gained attention to support early clinical decision-making. For instance, [Lewin-Epstein et al. 2021] proposed an EHR-based ensemble approach using clinical variables such as prior resistance and antibiotic exposure. Their models achieved AUC-ROC values of 0.73–0.79 without bacterial species information, and 0.80–0.88 when species identification was included, highlighting previous resistance history as a key predictor.

In a Brazilian cohort, [Silveira 2023] applied ML models using routinely available clinical variables, excluding prior resistance or bacterial species identification. Random Forest achieved the best performance, with AUC values varying widely across antimicrobials (0.585 to 0.986), suggesting sensitivity to antimicrobial-specific patterns. Similarly, [Rosa et al. 2024] evaluated MIMIC-III data for ICU patients using a minimal set of routine variables. Among five models tested, XGBoost achieved the best performance, reaching an AUC of 0.845.

Although these studies report promising results, direct comparisons should be interpreted with caution due to substantial differences in datasets, target definitions, and evaluated antimicrobials. To address these variations, the present study evaluates multiple ML models for AMR prediction using MIMIC-IV data, combining a comprehensive set of demographic, clinical, and derived variables. Furthermore, it incorporates explainability analysis through SHapley Additive exPlanations (SHAP) to investigate the factors associated with AMR risk, supporting transparent clinical decision-making.

3. Methodology

The methodological pipeline adopted in this study is illustrated in Figure 1 and comprises three main stages: data preparation, model training, and performance evaluation. The procedures performed in each phase are described in the following subsections.

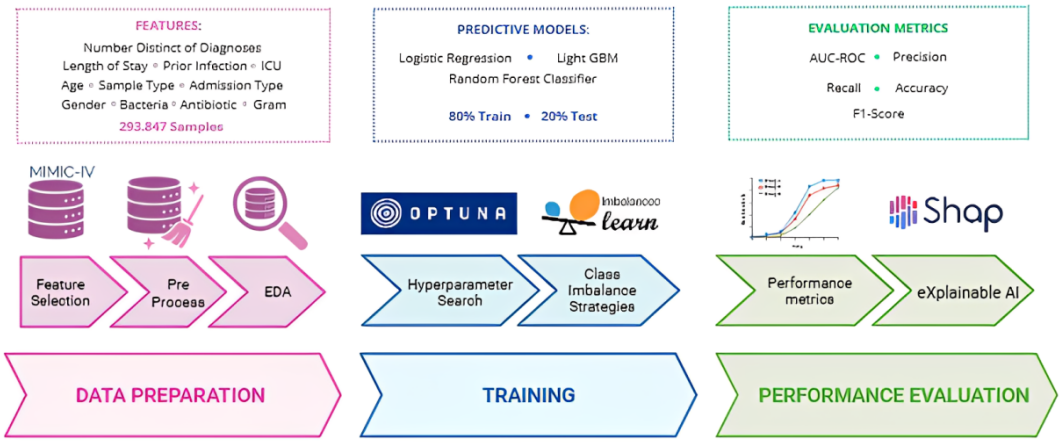


Figure 1. Workflow of the proposed methodology.

3.1. Computational Environment

This study was conducted using the computational resources of the High-Performance Computing Center (NPAD/UFRN). The execution environment consisted of two 64-thread Intel Xeon processors and a NVIDIA Tesla T4 GPU with 15 GB of dedicated memory. Additionally, Scikit-learn, SHAP, LightGBM, and Imbalanced-learn libraries were used in the implementation.

3.2. Data Preparation

3.2.1. Dataset and Feature Selection

The models were trained using data from the MIMIC-IV database [Johnson et al. 2024]. Developed by the Massachusetts Institute of Technology in partnership with Beth Israel Deaconess Medical Center (BIDMC) in the United States, MIMIC-IV contains detailed clinical information on patients treated in Intensive Care Units (ICUs) and hospital wards, including demographic records, laboratory test results, and medical prescriptions. The final dataset used in this study comprises 293,847 microbiological susceptibility records obtained after applying the inclusion and exclusion criteria described in this section. In total, these records represent 23,333 unique patients.

Data extraction was performed using Google BigQuery. Microbiological records were obtained from the `microbiologyevents` table and combined with demographic information from `patients`, admission data from `admissions`, prescription records from `prescriptions`, and ICU stay information from `icustays`. Only records containing identified microorganisms, tested antimicrobials, and valid susceptibility interpretations were retained.

The attributes selected for this work, including demographic, clinical, and microbiological variables, are summarized in Table 1. The target variable was formulated as a binary classification problem (1 for resistant, 0 for susceptible), derived directly from the susceptibility interpretation in the *microbiologyevents* table. Furthermore, considering the diversity of information available in the dataset, several variables were derived from the original records to provide better clinical context. For instance, Gram classification was not directly available in MIMIC-IV and was obtained through a manual mapping between microorganisms and their corresponding Gram-positive or Gram-negative classification based on standard microbiological references.

Table 1. Description of the attributes used in the predictive models.

Variable	Type	Description
Demographic & Admission Context		
age	Integer	Patient age (18–90 years).
gender	Categorical	Patient gender.
admission_type	Categorical	Type of hospital admission.
is_icu	Binary	Admitted to ICU at culture collection.
length_of_stay	Integer	Days from admission to culture.
Clinical History		
prior_infection	Binary	Previous microbiological event in past admissions.
num_distinct_diagnoses	Integer	Total distinct diagnoses at admission.
Microbiological Characteristics		
sample_type	Categorical	Culture collected (e.g., blood, urine).
bacteria	Categorical	Identified microorganism (255 species).
gram	Categorical	Gram-positive or Gram-negative.
antibiotic	Categorical	Tested antimicrobial (16 drugs).
Target Variable		
target	Binary	Susceptibility (1 = resistant, 0 = susceptible).

The proposed model was developed to avoid training a separate model for each microorganism. Instead, a single model was trained using a dataset comprising 255 microorganisms, combined with the 16 most frequently occurring antimicrobials. In this way, the model learns patterns associated with antimicrobial resistance across different microorganism–antimicrobial combinations. During model use, the user provides the microorganism and antimicrobial of interest, together with the patient's other clinical characteristics, and the model predicts the antimicrobial susceptibility for that specific combination.

3.2.2. Data Preprocessing

The dataset was restricted to adult patients (18-90 years old). Patients older than 90 years were excluded due to the de-identification protocols of the MIMIC-IV database [Johnson et al. 2024]. To comply with HIPAA privacy standards, the database alters the chronological timelines and caps the ages of patients above 89 years, which could in-

troduce inaccuracies and bias into the demographic features of the predictive models. Records without microorganism identification, antimicrobial information, or susceptibility interpretation were excluded. In addition, screening and swab samples were removed, and only susceptibility tests involving 16 commonly prescribed antimicrobials were considered.

After preprocessing, the dataset was divided into training and testing subsets using a stratified hold-out strategy. Approximately 80% of the samples were allocated to the training set and 20% to the testing set. Stratification was applied to preserve the original class distribution in both subsets, and a fixed random seed (`random_state=42`) was used to ensure reproducibility.

3.2.3. Exploratory Data Analysis

In the exploratory data analysis, it was observed that the target variable is imbalanced, consisting of approximately 79% antimicrobial-sensitive microorganisms and about 21% resistant microorganisms. Regarding demographic characteristics, the dataset includes patients aged between 18 and 90 years, with a balanced distribution between female and male patients.

With respect to clinical variables, a predominance of Gram-negative bacteria was observed, accounting for approximately 75% of the dataset. Furthermore, the dataset includes 255 distinct microorganisms, evaluated across different combinations with 16 antimicrobials, including Ceftriaxone, Meropenem, and Vancomycin.

3.3. Training

3.3.1. Machine Learning Models

The supervised models used in this study include Random Forest, Logistic Regression, and Light Gradient Boosting Machine (LightGBM). In binary classification problems, such as the one addressed in this study, logistic regression is widely used because it directly models the probability of class membership through a sigmoid function.

Additionally, more complex models based on decision trees were explored. Random Forest is built from multiple trees trained independently, using data sampling and random feature selection at each split. This strategy controls model variance, making it less susceptible to overfitting and more robust. In turn, LightGBM implements boosting methods, which consist of training multiple models sequentially, so that each new model aims to correct the errors made by the previous ones. This method has a high capacity to capture non-linear relationships and interactions between variables.

3.3.2. Hyperparameter Search with Optuna

Hyperparameters were optimized using the open-source library Optuna, which applies probabilistic models to search the space more efficiently than traditional grid or random search. Optimization was performed on the training set using five-fold cross-validation, with the mean AUC-ROC across folds as the objective, over 100 trials for each model.

Table 2 summarizes the AUC-ROC before and after optimization, along with the best configuration found for each model. Random Forest showed the largest gain (0.86 to 0.88), followed by LightGBM (0.84 to 0.86) and Logistic Regression (0.78 to 0.80).

Table 2. Best hyperparameters found via Optuna for each model.

Model	AUC-ROC (before→after)	Best configuration
Random Forest	0.86 → 0.88	`n_estimators=311, max_depth=30, min_samples_split=3, min_samples_leaf=1, max_features=None`
Logistic Regression	0.78 → 0.80	`C=3.77, penalty=l1, solver=liblinear, max_iter=186`
LightGBM	0.84 → 0.86	`n_estimators=369, max_depth=8, num_leaves=47, learning_rate=0.10`

3.3.3. Class Imbalance Strategies

In class imbalance scenarios, such as the one observed in the target variable, models tend to favor the majority class, compromising the ability to correctly identify the minority class [He and Garcia 2009]. Thus, balancing techniques were applied to mitigate this issue and improve model performance.

Under-sampling and over-sampling strategies were employed. In the context of under-sampling, the Random Under-Sampling method was used, which consists of reducing the number of examples from the majority class through random removal. For over-sampling, the SMOTE (Synthetic Minority Over-sampling Technique) method was applied, which generates new synthetic examples of the minority class based on the similarity among existing data points. To rigorously prevent data leakage, SMOTE was applied exclusively to the training set after the initial train-test split, leaving the test set completely isolated with its original class distribution for evaluation. Both techniques generated balanced class distributions, allowing the impact of class imbalance mitigation on model performance to be evaluated.

3.4. Performance Evaluation

3.4.1. Evaluation Metrics

Model performance was evaluated using Accuracy, Precision, Recall, F1-score, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). Given the class imbalance observed in the dataset, AUC-ROC was adopted as the primary evaluation metric, as it assesses the model's discriminative ability across different classification thresholds and is less sensitive to class distribution than accuracy alone.

3.4.2. Explainable AI

In this study, the application of Explainable Artificial Intelligence (XAI) techniques aims to make machine learning models more interpretable, transparent, and reliable, especially in sensitive domains such as healthcare. Predictive models can be considered "black

boxes", as it is difficult to understand the factors that influence their decisions. In this context, interpretability techniques allow the identification of the contribution of each variable to the predictions made. Based on this, the SHapley Additive exPlanations (SHAP) method was applied, enabling a detailed analysis of the influence of variables at both the global and individual levels. In this application, a subset of 1,000 test samples was analyzed using the trained models, in order to reduce computational cost while maintaining representative explanations.

4. Results

This section presents the predictive performance of the evaluated machine learning models under different class imbalance scenarios, as well as the explainability analysis obtained through SHAP.

Table 3. Comparison of models after class balancing strategies

Strategy	Acc.	Prec.	Recall	F1	ROC AUC
Random Forest					
Imbalanced	**0.85**	**0.77**	0.40	0.53	**0.88**
Random Under-Sampling	0.79	0.50	**0.80**	**0.62**	0.87
SMOTE	**0.85**	0.68	0.52	0.59	0.87
Logistic Regression					
Imbalanced	0.82	0.65	0.31	0.42	0.80
Random Under-Sampling	0.73	0.42	**0.74**	**0.54**	**0.80**
SMOTE	0.73	0.42	0.73	0.53	0.80
LightGBM					
Imbalanced	**0.84**	**0.71**	0.44	0.54	**0.86**
Random Under-Sampling	0.78	0.48	**0.79**	**0.60**	0.86
SMOTE	**0.84**	0.66	0.50	0.57	0.86

As shown in Table 3, the Random Forest achieved the best overall performance, reaching an AUC-ROC of 0.88 on the original imbalanced dataset. Furthermore, it also achieved the highest accuracy (0.85) and precision (0.77).

The balancing strategies impacted the models in different ways. In general, Random Under-Sampling significantly increased recall, indicating a greater ability to identify resistant samples, though this gain was accompanied by a reduction in accuracy and precision. From a clinical perspective, higher recall may be desirable because it reduces the number of resistant cases incorrectly classified as susceptible. However, the substantial decrease in precision increases the number of false positive predictions, which may lead to unnecessary use of broad-spectrum antimicrobials. Therefore, the selection of the final model involved balancing sensitivity to resistant cases with overall predictive reliability.

Although the models trained with Random Under-Sampling achieved the highest recall values, this approach was not considered the most suitable as the main model. This is because the method randomly reduces the number of examples in the majority class, potentially discarding clinically relevant information present in the original data. Similarly, while SMOTE produced competitive results, it was not selected as the primary

approach, because generating synthetic clinical profiles can introduce artificial noise and distort the true epidemiological distribution of antimicrobial resistance. Therefore, priority was given to the Random Forest trained on the original dataset, as retaining the original imbalanced dataset preserves the real-world clinical prevalence, while the Random Forest algorithm demonstrated a strong inherent capability to handle this class imbalance. Overall, this approach provided a better balance among AUC-ROC, accuracy, precision, and predictive robustness.

The interpretability analysis using SHAP, presented in Figures 2, 3, and 4, revealed consistent patterns across the evaluated models. Among the most influential variables were prior infection history, the number of distinct diagnoses, the Gram classification of the microorganism, and the length of hospital stay until sample collection. Figure 2 shows that prior infection history was the most influential feature for the Random Forest, followed by hospitalization-related variables and microbiological characteristics. In Figure 3, a similar pattern is observed, although the relative contribution of demographic and clinical variables differs due to the linear nature of Logistic Regression. Finally, Figure 4 confirms the importance of prior infection history and microbiological characteristics, reinforcing the consistency of the explainability results across different models.

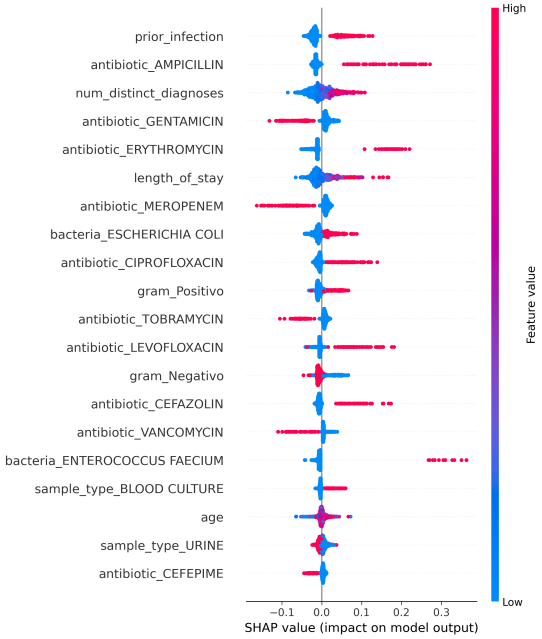


Figure 2. Feature Importance Plot for the Random Forest Model.

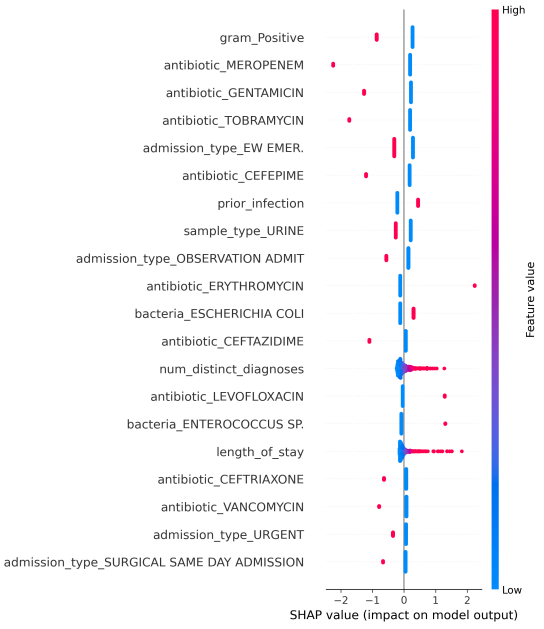


Figure 3. Feature Importance Plot for the Logistic Regression Model.

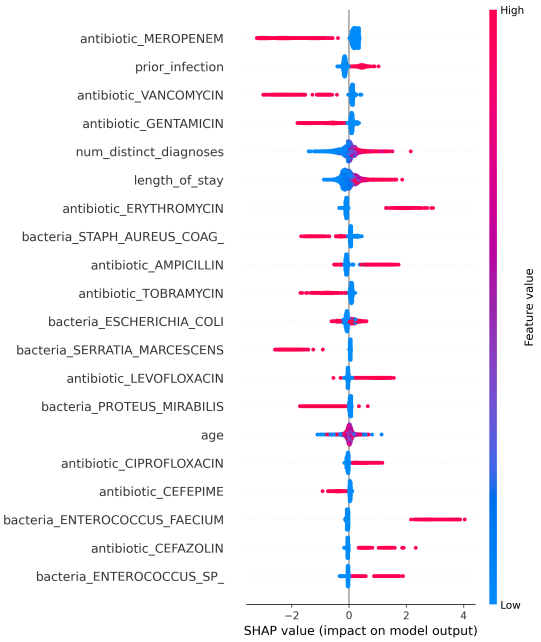


Figure 4. Feature Importance Plot for the LightGBM Model.

5. Discussion

5.1. Model Performance Analysis

As observed, the model that achieved the best performance in this study was Random Forest, with an AUC-ROC of approximately 0.88. When comparing this result with related studies, it is observed that higher values can be achieved in more specific scenarios. For example, in the study by [Silveira 2023], models were trained separately for each evaluated antibiotic. In this approach, Random Forest achieved an AUC-ROC of up to 0.98 for predicting vancomycin resistance, while for other antimicrobials the values ranged between 0.81 and 0.84. These results reflect a more constrained classification scenario, in which each model learns patterns associated with a single drug, reducing task complexity. In contrast, the present study proposes a unified approach capable of simultaneously predicting multiple combinations of microorganisms and antibiotics, significantly increasing the complexity of the problem.

Among the analyzed studies, the work by [Rosa et al. 2024] is the most methodologically similar to this research, as it also uses data derived from the MIMIC database, albeit from an earlier version, and proposes a general model for predicting antimicrobial resistance. In that study, the best performance was achieved by the Extreme Gradient Boosting (XGBoost) model, with an AUC of 0.84 on the original dataset and 0.84 after applying SMOTE. The Random Forest model, in turn, achieved AUC values between 0.78 and 0.79, which are lower than those found in this study. A possible explanation for this difference lies in the set of attributes used. However, differences in dataset versions, preprocessing procedures, and cohort composition may also have contributed to the observed performance gap, making it difficult to isolate the effect of the additional derived features alone. While Rosa et al. employed mainly basic demographic and microbiological variables, such as age, gender, sample type, tested antibiotic, and Gram classification, the present study also incorporated derived clinical variables, such as length of hospital stay

until sample collection, number of comorbidities within a one-year period, prior infection history, and ICU admission. These additional features tend to provide a more comprehensive representation of the patient's clinical context, contributing to greater discriminative capacity of the models and, consequently, improved predictive performance.

In summary, the comparative analysis shows that the Random Forest model developed in this study presents robust and competitive predictive performance, given the complexity of the prediction scenario. The obtained results reinforce the importance of incorporating contextual clinical variables in antimicrobial resistance modeling, indicating that the inclusion of such information can enhance the discriminative capacity of the models. Thus, a detailed interpretation of the contribution of variables to the predictions becomes essential for a deeper understanding of model behavior and the factors associated with resistance.

5.2. Model Explainability Analysis

In the interpretability analysis using SHAP (Figures 2, 3, and 4), consistency was observed among the most influential features across the three evaluated models, despite structural differences between the algorithms. Overall, convergence was observed in the most relevant predictors, suggesting that the identified factors represent robust clinical signals associated with the risk of antimicrobial resistance.

Among the most important variables, prior infection history stood out. In all models, higher values of this feature were associated with positive SHAP values, indicating that patients with prior infection history have a higher predicted probability of resistance. This result is consistent with findings in the literature [Lewin-Epstein et al. 2021], which also identify prior infections as a relevant factor in predicting bacterial resistance. Although the referenced study specifically emphasizes previously resistant infections, both studies indicate that recurrent infections constitute an important predictor of antimicrobial resistance.

Clinical indicators also showed strong predictive contribution. The number of distinct diagnoses (*num_distinct_diagnoses*) and the length of hospital stay until sample collection (*length_of_stay*) were among the most influential variables. Multimorbidity is associated with greater exposure to antibiotics and invasive procedures, a context classically linked to the development of antimicrobial resistance [Chahine et al. 2026]. Furthermore, longer hospital stays increase the likelihood of colonization and infection by multidrug-resistant pathogens associated with the hospital environment [Santos 2004].

In addition, in the microbiological context, the presence of *Enterococcus faecium* stood out as one of the main positive predictors of resistance across all models. This finding is clinically expected, given that this species exhibits high rates of intrinsic and acquired resistance, including resistance to vancomycin and multiple therapeutic classes [Hollenbeck and Rice 2012, Geraldes et al. 2022]. Moreover, *E. faecium* is classified by the World Health Organization as a high-priority pathogen [World Health Organization 2024], with extremely limited treatment options.

Finally, variables related to antibiotics should be interpreted with caution. Since not all bacteria are tested against all 16 available drugs, the presence of these features in the model does not necessarily indicate specific resistance to the corresponding antibiotic. In many cases, such variables reflect the clinical context in which these medications are

typically used, such as more severe infections, suspicion of multidrug-resistant pathogens, or broad-spectrum hospital use. Therefore, their relevance in SHAP may represent prescribing patterns and clinical severity captured by the model.

5.3. Limitations and Future Research

The main limitation of this study is the lack of external validation of the developed models. The entire training and evaluation process was conducted exclusively using data from the MIMIC-IV database, which may limit the generalizability of the findings to different clinical settings. The use of a single data source may introduce potential population, institutional, and operational biases. Although external validation is an important step toward assessing model robustness, obtaining an equivalent dataset remains challenging. To the best of the authors' knowledge, there are currently no publicly available datasets that provide a comparable combination of microbiological susceptibility records, antimicrobial prescriptions, hospitalization history, and patient-level clinical information at a scale similar to MIMIC-IV. Consequently, reproducing the complete experimental setting in an independent cohort is not straightforward.

6. Conclusion

This study investigated the application of machine learning models for predicting antimicrobial resistance using clinical data from the MIMIC-IV database. The Random Forest model achieved the best overall performance, with an AUC-ROC of approximately 0.88, outperforming or remaining competitive with similar studies reported in the literature.

The interpretability analysis using SHAP enabled the identification of the main features associated with the model predictions. Variables such as prior infection history, number of distinct diagnoses, length of hospital stay until sample collection, and microbiological characteristics had the greatest influence on the decisions. These findings were consistent across different algorithms and are aligned with evidence reported in the literature, reinforcing the reliability and clinical coherence of the obtained results.

Thus, given the consistency of the results and the clinical grounding presented, the developed models show potential to support medical decision-making, particularly in aiding the early identification of the risk of antimicrobial resistance. Pending clinical validation, these findings may support more informed prescribing decisions and contribute to the more appropriate use of antibiotics, with potential downstream benefits for therapeutic outcomes.

Finally, as future work, external validation of the models on datasets from different institutions is intended, in order to assess their generalization capability across other clinical settings.

Supplementary Material. The supplementary resources for reproducibility are publicly available[1], including scripts and notebooks for hyperparameter tuning, sampling strategies, and SHAP analysis.

[1] `https://osf.io/t34wr/files/osfstorage`

Acknowledgments. The authors thank the Center for High Performance Computing (NPAD/UFRN), available at `https://npad.ufrn.br`, for providing computing resources for data processing.

The authors declare that Artificial Intelligence (AI) tools, such as large language models like Claude.ai, were used solely to translate and refine the manuscript's language, clarity, and structure. These tools did not generate any research data, results, or conclusions, which remain the exclusive product of the authors' validation and responsibility.

References

Chahine, E. B., Watanabe, J. H., Dong, M. T., and Andrade, R. M. (2026). Beyond resistance: Prioritizing safety in antimicrobial stewardship for older adults. *The Senior Care Pharmacist*, 41(3):88–97.

Geraldes, C., Tavares, L., Gil, S., and Oliveira, M. (2022). Enterococcus virulence and resistant traits associated with its permanence in the hospital environment. *Antibiotics*, 11(7):857.

He, H. and Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284.

Hollenbeck, B. L. and Rice, L. B. (2012). Intrinsic and acquired resistance mechanisms in enterococcus. *Virulence*, 3(5):421–433.

Johnson, A., Bulgarelli, L., Pollard, T., Gow, B., Moody, B., Horng, S., Celi, L. A., and Mark, R. (2024). MIMIC-IV. *PhysioNet*. Version 3.1.

Lewin-Epstein, O., Baruch, S., Hadany, L., Stein, G. Y., and Obolski, U. (2021). Predicting antibiotic resistance in hospitalized patients by applying machine learning to electronic medical records. *Clinical Infectious Diseases*, 72(11):e848–e855.

Ministry of Health (Brazil) (2026). Antimicrobial resistance (AMR). Gov.br Portal.

Rosa, S. d. C. L., Carvalho, N. I., and Duani, H. (2024). Machine learning for microbial resistance prediction. *Journal of Health Informatics*, 16(Especial).

Santos, N. d. Q. (2004). Bacterial resistance in the context of hospital infection. *Texto & Contexto – Enfermagem*, 13(spe):64–70.

Silveira, B. P. (2023). Use of machine learning to predict antibiotic susceptibility through hospitalized patient data. Master's thesis, Federal University of Rio Grande do Sul, Porto Alegre.

World Health Organization (2023). Global antimicrobial resistance and use surveillance system (GLASS) report 2023. Technical report, World Health Organization, Geneva.

World Health Organization (2024). WHO updates list of drug-resistant bacteria most threatening to human health. WHO News Release.