Evolutionary Optimization of Hospital Bed Allocation to Reduce Mortality from Severe Acute Respiratory Syndrome

Beatriz da Silva Vieira*, Murillo Guimarães Carneiro*, Luiz Gustavo Almeida Martins*
*Federal University of Uberlândia, Brazil
Email: {beatriz.silva,mgcarneiro,lgamartins}@ufu.br

***Abstract*—Hospital bed availability is a critical factor in reducing mortality from Severe Acute Respiratory Syndrome (SARI), particularly under high-demand conditions. This study presents a framework that integrates machine learning-based mortality prediction with a Genetic Algorithm (GA) to optimize hospital bed allocation using a multi-criteria scalar objective function.**

Tree-based models, particularly the Random Forest Regressor, capture variability in SARI mortality under overdispersion and zero inflation. The GA explores trade-offs between mortality reduction, equity, and redistribution cost.

Results show that performance depends on hospital occupancy. In high-occupancy scenarios (98%), small redistributions reduce mortality, whereas in lower occupancy scenarios (60% and 75%), gains are limited.

These results indicate higher effectiveness under high occupancy and limited impact in lower-demand scenarios.

***Index Terms*—Hospital bed allocation, Severe Acute Respiratory Syndrome (SARI), Random Forest, Genetic Algorithm, Multi-criteria optimization, Healthcare resource management.**

I. Introduction

The limited availability of hospital beds is a central challenge for the effectiveness of healthcare delivery in Brazil, especially in high-demand situations such as Severe Acute Respiratory Syndrome (SARI). Demographic factors such as population aging, combined with urbanization and the impacts of the COVID-19 pandemic, which resulted in an increase in severe cases reported as SARI, aggravated hospital overload, highlighting persistent deficiencies in the public healthcare system [4], [5].

Although the COVID-19 pandemic motivated the analysis of high-demand scenarios, this study considers a recent temporal window allowing the evaluation of the relationship between hospital capacity and mortality under contemporary conditions of the healthcare system. The analysis is conducted in the context of SARI, adopted as a case study, with extension to other scenarios being a possibility to be investigated in future work.

The shortage of hospital beds, especially in intensive care units (ICUs), may delay treatments, overload emergency services, and increase mortality, mainly among severe patients. Although SARI has well-defined clinical characteristics, its outcome depends strongly on the healthcare system response capacity [4], [5].

In scenarios of high hospital occupancy, high levels of bed utilization may generate delays in care, patient transfers, and compromise clinical outcomes [14], making the use of quantitative tools relevant to support efficient allocation of resources. In real contexts, these decisions may involve the transfer of patients between municipalities, the prioritization of regions with higher risk and the definition of strategies under operational and capacity constraints.

Understanding the relationship between hospital bed availability and mortality due to SARI is essential to support resource allocation decisions. This study investigates how different levels of bed supply and occupancy influence mortality in distinct epidemiological and demographic scenarios considering variables such as number of cases, population size, seasonality, and municipal characteristics.

For this purpose, supervised machine learning algorithms (*Random Forest*, *XGBoost*, and *Gradient Boosting*) are applied to model mortality based on structural variables. The obtained estimates are then integrated into a Genetic Algorithm (GA) used to optimize the redistribution of hospital beds between municipalities considering different levels of pressure on the system.

The main differential of this work lies in the integration between predictive modeling and evolutionary optimization in a single decision-making process. While traditional approaches treat the prediction of clinical outcomes and the allocation of resources separately, this study directly incorporates mortality predictions into the fitness function of the genetic algorithm. This integration allows the evaluation of bed redistribution strategies, explicitly considering their impacts on clinical outcomes, characterizing a data-driven approach to support hospital management.

II. Objectives

The main objective of this work is to analyze the relationship between the availability of hospital beds and mortality, integrating mortality predictive models into an evolutionary optimization process applied to bed allocation.

- Develop and evaluate hospital mortality predictive models using epidemiological, structural, and seasonal variables, employing temporal validation and bias calibration;
- Select variables through temporal cross-validation;
- Integrate predictions into a multi-criteria fitness function of a GA, considering mortality, equity, and cost;

- Assess the GA performance under different hospital occupancy scenarios.

III. Related Work

The literature presents several studies focused on the prediction of hospital mortality through statistical methods and machine learning with a predominant focus on intensive care units (ICUs). Among them, the study that analyzes all-cause mortality within 28 days in patients with chronic heart failure admitted to ICUs stands out by using the stress hyperglycemia ratio as an explanatory variable [3]. The analysis was conducted based on retrospective data from the MIMIC database, involving 913 patients, and indicated an association between high levels of this indicator and higher mortality in the analyzed period.

Another set of studies explores the use of machine learning to support hospital resource management. In this context, the study presented in [1] proposes the prediction of length of stay of pediatric patients based on Brazilian hospital data using different supervised learning algorithms. The results indicated better performance of tree-based models for estimating the average length of stay.

Despite employing techniques similar to those used in this study, the cited work differs regarding the analyzed outcome since it focuses on predicting length of stay and not on modeling hospital mortality nor on the analysis of scenarios of bed constraints.

Additionally, studies focused on dynamic prediction of mortality in ICUs are also reported in the literature. In [6] a machine learning model is proposed for predicting hospital mortality based on repeated measurements incorporating clinical and demographic variables. The performance of the model is compared with traditional clinical scores used in hospital environments.

In addition to approaches based exclusively on machine learning the literature also presents applications of evolutionary algorithms in problems of optimization of hospital processes. In [12] the use of genetic algorithms is proposed to optimize patient scheduling in an oncology center based on simulation models aiming to reduce waiting times and improve the use of hospital resources.

More closely related to the problem addressed in this work [13] proposes a hybrid model that combines queueing theory and genetic algorithms to simultaneously optimize bed occupancy and associated costs explicitly considering patient flow in hospital environments. This approach highlights the potential of evolutionary methods in hospital capacity management under operational constraints.

More recently, approaches based on artificial intelligence integrated with operational practices such as Lean have been used to optimize hospital capacity and patient flow demonstrating significant gains in efficiency and reduction of waiting times [14].

In summary, the related works generally address two main directions (i) the prediction of hospital outcomes through machine learning models and (ii) the optimization of hospital processes and resource allocation through mathematical models simulation or evolutionary algorithms. However, it is observed that these approaches are largely treated separately.

Differently from these studies, the present work proposes an integrated approach in which a machine learning-based predictive model of mortality is directly incorporated into the fitness function of a genetic algorithm. This integration allows the evaluation of the impact of different bed allocation strategies explicitly considering clinical outcomes characterizing a data-driven approach for optimizing hospital management in scenarios of system overload.

IV. Methodology

This section presents the procedures adopted in the study, including data collection and integration, pre-processing, exploratory analysis, predictive modeling, and model tuning.

A. Data Collection and Integration

The main database used was provided by the Ministry of Health through the Health Surveillance Secretariat (SVS) responsible for the national surveillance program of SARI. This database consolidates records since the Influenza A(H1N1)pdm09 pandemic initiated in 2009 until 2025, including detailed information on municipality of occurrence, demographic characteristics, symptoms, comorbidities, clinical evolution, and patient outcomes. For this study, only variables referring to municipalities in Minas Gerais with estimated population equal to or greater than 35000 inhabitants were considered, covering number of cases, hospitalizations, deaths, and date of notification. The data were aggregated by municipality and by year and month, allowing the obtaining of monthly totals of cases, hospitalizations, and deaths [7].

With the objective of enriching the epidemiological dataset, additional information from the Brazilian Institute of Geography and Statistics (IBGE) with municipal population estimates and from the Unified Health System (SUS) through the OpenDataSUS portal, containing data on hospital infrastructure of municipalities, were integrated. Although the SUS database discriminates different types of beds in this work only the total number of available beds per municipality aggregated by municipal code was considered allowing the analysis of its relationship with SARI mortality indicators [2], [8], [9].

The population estimates of IBGE have annual periodicity. To make them compatible with the monthly SARI data linear interpolation was applied to estimate average monthly population values for each municipality. For 2023, a year without data publication by IBGE due to the postponement of the Demographic Census caused by the COVID-19 pandemic, the estimated values for 2022 were maintained ensuring continuity and consistency of the time series [11].

Given the population heterogeneity of Minas Gerais and the structural dependence of smaller municipalities on urban centers for hospital care the analysis was restricted to municipalities with population equal to or greater than 35000 inhabitants aiming to reduce distortions caused by intermunicipal dependence on hospital infrastructure. In addition only records until December 2024 were considered ensuring temporal consistency of the historical series used.

B. Pre-Processing

After the integration of epidemiological, demographic, and structural databases, the variables were subjected to preprocessing procedures aiming at standardization and temporal consistency. The temporal variable was converted to date format allowing the extraction of the month of occurrence and chronological ordering by municipality and period.

Missing values in the hospital beds variable were replaced with zero assuming local unavailability. Undefined or infinite values resulting from mathematical operations were also treated by replacement with zero ensuring numerical stability.

To capture seasonal patterns, each observation was associated with the season of the year (summer: December to February, autumn: March to May, winter: June to August, spring: September to November). Additionally, the month was transformed into cyclic components using sine and cosine functions while the seasons were encoded as indicator variables (dummies) with exclusion of the first category to avoid multicollinearity.

To allow comparison between municipalities of different sizes, the variable beds per thousand inhabitants was constructed. The total number of cases was maintained in absolute values while the monthly growth rate was calculated as percentage variation in relation to the previous month with handling of division by zero.

The annual population estimates from IBGE were linearly interpolated to generate monthly values compatible with the epidemiological records [10]. The final set of variables included total cases, growth rate, beds per thousand inhabitants, cyclic components of the month and seasonal dummies. The data were organized chronologically and partitioned into training validation and test sets ensuring consistent predictive evaluation.

C. Exploratory Analysis and Statistical Diagnostics of Mortality Data

Prior to the definition of predictive models, an exploratory analysis was conducted with the aim of characterizing the distribution of hospital mortality data due to SARI and evaluating the suitability of classical statistical approaches. The response variable consists of integer counts of deaths, with the presence of zero values and strong skewness, characteristics that violate the assumptions of normality and homoscedasticity required by linear regression models.

The ratio between the variance and the mean of the response variable was evaluated as a criterion for identifying overdispersion, a condition indicating the inadequacy of the classical Poisson model. Subsequently, count models based on the Poisson, Negative Binomial, and Zero-Inflated Negative Binomial (ZINB) distributions were fitted, with the objective of assessing the simultaneous presence of overdispersion and excess structural zeros. The comparison between models considered the AIC and BIC information criteria, as well as the ability to reproduce the observed proportion of zeros in the sample.

From this analysis, the adequacy of classical count models to represent nonlinear relationships and interaction effects was evaluated.

D. Predictive Modeling of SARI Mortality

The predictive modeling stage aimed to estimate the monthly number of deaths due to SARI based on epidemiological, structural, and seasonal variables, providing a quantitative basis for integration with the evolutionary optimization process. The response variable corresponds to the total number of deaths recorded per municipality and month, characterizing a non-negative, skewed outcome subject to abrupt variations over time.

To address these characteristics, supervised machine learning algorithms based on decision trees were adopted, nonparametric methods capable of capturing nonlinear relationships and interaction effects among explanatory variables, with greater robustness to outliers and structural heterogeneity in the data.

Six tree-based models were evaluated: *Random Forest Regressor*, *Extra Trees Regressor*, *Gradient Boosting Regressor*, *HistGradientBoosting Regressor*, *XGBoost Regressor*, and *CatBoost Regressor*. This approach allowed the comparison of different ensemble strategies, including *bagging* and *boosting*, widely used in regression problems with structured data.

The initial set of explanatory variables included epidemiological indicators (total cases and monthly case growth rate), a structural variable of hospital capacity (beds per thousand inhabitants), and seasonal components, represented by cyclic encoding of the month (sine and cosine) and indicator variables for the seasons of the year. All observations were chronologically ordered by municipality and period, respecting the temporal nature of the problem.

E. Feature Selection, Training and Hyperparameter Tuning

Feature selection was performed using *Recursive Feature Elimination with Cross-Validation* (RFECV), using a reference model based on *XGBoost Regressor* to capture nonlinear relationships and interactions between variables. The procedure adopted temporal cross-validation (*TimeSeriesSplit*) preserving the chronological order of the data and avoiding information leakage between periods. The metric used was the root mean squared error (RMSE) evaluated in negative form for compatibility with the optimization criterion.

The data were partitioned into three distinct temporal sets training validation and test. The validation set was used to estimate a multiplicative bias calibration factor defined as the ratio between the mean of the predictions and the mean of the observed values applied subsequently to the predictions of the test set. Negative values were truncated to zero ensuring consistency with the response variable.

The hyperparameter tuning of the evaluated models was performed using *Grid Search* combined with temporal cross-validation in the training set considering parameters such as number of estimators maximum depth of the trees learning rate and regularization in a specific way for each algorithm.

The final performance was evaluated on the test set after bias calibration using the metrics RMSE MAE and R^2. The model with the lowest RMSE was selected as the final predictive model being incorporated into the genetic algorithm together with the selected variables and the calibration factor composing the multi-criteria fitness function for optimization of hospital bed allocation.

It is important to highlight that the predictive model is not used as a surrogate model in the classical sense of surrogate-assisted optimization. Differently from this approach, in which an approximate model replaces real evaluations typically costly (for example, complex simulations), in this work the machine learning model directly defines the evaluation function used in the optimization process.

Thus, the proposed approach is characterized as data-driven optimization in which the predictions of the model are used as the basis to evaluate candidate solutions in the genetic algorithm.

F. Evolutionary Optimization of Hospital Bed Allocation

After defining the predictive mortality model, a GA was applied to optimize the redistribution of hospital beds among municipalities, simultaneously considering the reduction of predicted deaths, equity in the distribution of impacts, and penalties associated with resource redistribution. The GA is suitable for multi-criteria optimization problems with operational constraints and high-dimensional continuous search spaces.

1) Representation and Base Scenario: Each individual represents a candidate solution encoded as a continuous vector

$$\mathbf{L} = (L_1, L_2, \ldots, L_N), \tag{1}$$

where L_i is the number of beds per thousand inhabitants in municipality i and N is the total number of municipalities. Values are constrained to

$$0.5\, L_i^{(0)} \leq L_i \leq 2.0\, L_i^{(0)}, \tag{2}$$

with $L_i^{(0)}$ representing the observed average value in the reference period. The predicted number of deaths in the base scenario, without redistribution, is obtained from the calibrated predictive model, being

$$Y^{(0)} = \sum_{j=1}^{M} \hat{y}_j^{(0)}, \tag{3}$$

where M is the total number of observations.

2) Linear Approximation of the Impact of Beds and Occupancy Scenarios: For efficient evaluation of solutions, a local linear approximation of the sensitivity of mortality to beds was adopted:

$$s_j = \frac{\hat{y}_j^{(\varepsilon)} - \hat{y}_j^{(0)}}{\varepsilon}, \quad \hat{y}_j = \hat{y}_j^{(0)} + \phi(\rho)\, s_j\, \Delta L_{c(j)}, \tag{4}$$

in which $\varepsilon > 0$ is a small increment used to numerically approximate the derivative by finite differences, and s_j represents the local sensitivity of mortality in relation to the variation in the number of beds for observation j. The term $\Delta L_{c(j)} = L_{c(j)} - L_{c(j)}^{(0)}$ corresponds to the variation in the number of beds of the municipality associated with observation j, while $\phi(\rho)$ represents a saturation factor that adjusts the impact of beds according to the level of system occupancy.

$$\phi(\rho) = \frac{1}{1 - \rho + \varepsilon}, \tag{5}$$

in which $\rho \in \{0.60, 0.75, 0.98\}$ represents scenarios of low, medium, and high hospital occupancy and $\varepsilon > 0$ is a small term introduced to avoid numerical instabilities when $\rho \to 1$.

3) Evaluation Metrics and Fitness Function: The adjusted predictions are aggregated by municipality, generating the vector $\mathbf{O} = (O_1, \ldots, O_N)$. The main metrics are:

- Weighted relative impact: $I = \frac{\sum_{i=1}^{N} O_i\, w_i}{\sum_{i=1}^{N} O_i^{(0)}\, w_i}$, with $w_i = 1/(L_i^{(0)} + \delta)$, in which $\delta > 0$ is a regularization term to avoid division by zero;
- Inequality in the distribution of deaths: Gini coefficient $G = \frac{2\sum_{i=1}^{N} i\, O_{(i)}}{N \sum_{i=1}^{N} O_i} - \frac{N+1}{N}$;
- Redistribution cost: $C = \left(\frac{\sum_{i=1}^{N} L_i - \sum_{i=1}^{N} L_i^{(0)}}{\sum_{i=1}^{N} L_i^{(0)}} \right)^2$.

The fitness function is

$$F = w_I I + w_G G + w_C C, \tag{6}$$

in which I represents the estimated mortality, G the measure of equity in the distribution of beds and C the redistribution cost. The weights w_I, w_G, w_C control the relative importance of each component.

The problem is formulated as a scalar multi-criteria optimization using a weighted sum. The variation of weight combinations allows exploring different trade-offs between the objectives, reflecting scenarios with distinct priorities between mortality equity and cost.

4) Genetic Operators and Execution: The initial population is generated from the observed configuration with Gaussian perturbations of standard deviation 0.15 and the individuals are limited to $[0.5, 2.0]\, L_i^{(0)}$. Arithmetic crossover is applied with random coefficients $\alpha \in [0.3, 0.7]$:

$$\mathbf{L}^{(f)} = \boldsymbol{\alpha} \odot \mathbf{L}^{(p_1)} + (1 - \boldsymbol{\alpha}) \odot \mathbf{L}^{(p_2)}, \tag{7}$$

with $\odot$ element-wise product. Probabilistic mutation adds Gaussian noise to the selected genes and selection is performed by tournament. A fraction of elitism preserves the best solutions between consecutive generations.

The GA parameters were defined based on configurations commonly adopted in the literature, aiming to balance population diversity and stable convergence. The adopted configuration proved to be adequate for the problem under study showing stabilization of the fitness function along the generations as evidenced by the convergence analysis.

The overall workflow of the proposed approach, including data processing, predictive modeling, and evolutionary optimization, is illustrated in Fig. 1.

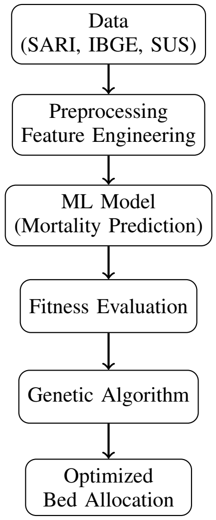

Fig. 1: Overall framework of the proposed methodology, integrating predictive modeling and evolutionary optimization.

V. RESULTS

This section presents the results of the analyses conducted to evaluate SARI death prediction and hospital bed optimization. First, the performance of classical statistical count models is assessed to highlight their limitations in capturing overdispersion and zero-inflation. Next, the predictive performance of tree-based machine learning models is evaluated, including variable selection, calibration, and test set evaluation. Finally, the results of the GA are presented, showing how different weightings and hospital occupancy scenarios affect bed redistribution and predicted deaths.

A. Statistical Model

The response variable, number of deaths due to SARI, presented a mean of 1.573 and a variance of 38.272, resulting in a variance-to-mean ratio of 24.33, which indicates strong overdispersion. To evaluate parametric approaches, classical count models were fitted: Poisson, Negative Binomial (NB), and Zero-Inflated Negative Binomial (ZINB).

The Poisson model presented a Pearson dispersion statistic of 1443.38 ($df = 670$, $p < 0.001$), confirming the overdispersion. The NB and ZINB models did not converge stably and provided inconsistent estimates. The observed proportion of zeros was 55%, while the Poisson, NB, and ZINB models predicted 38%, 63%, and 100%, respectively.

These findings indicate that classical count models do not adequately capture the variability and dispersion of SARI mortality, supporting the adoption of tree-based predictive models for death forecasting.

B. Predictive Model Results

For forecasting SARI deaths, the dataset was temporally split into training (January-December 2023), validation (January to April 2024), and testing (May to December 2024). The initial set of variables included total number of cases, case growth rate, beds per thousand inhabitants, cyclic month components, and seasonal indicators.

Variable selection was performed with RFECV, in a time series scheme, resulting in five optimal predictors: total cases, case growth rate, beds per thousand inhabitants, sine of the month, and winter indicator.

The models were trained on the training set and calibrated on the validation set using a multiplicative correction factor, ensuring non-negative predictions and mitigating systematic biases on the test set.

Table I: Predictive Model Performance on the Test Set

Model	*RMSE*	*MAE*	*R²*	*Bias*
Random Forest	**2.595**	**1.399**	**0.903**	0.602
Extra Trees	2.656	1.503	0.898	0.626
Gradient Boosting	3.033	1.512	0.867	0.634
XGBoost	3.391	1.618	0.834	0.620
CatBoost	3.859	1.719	0.785	0.569
HistGradientBoosting	4.809	2.222	0.667	0.590

In the test set, the *Random Forest Regressor* presented the best overall performance, followed by the *Extra Trees*. The selected model was then integrated into the GA framework to estimate the impact of hospital bed redistribution on SARI mortality.

C. Genetic Algorithm Results

The GA was employed to optimize the redistribution of hospital beds among municipalities, considering three weighted objectives: impact on mortality, regional inequality, and redistribution cost. Thirty independent runs were performed, each with 50 generations and a population of 150 individuals, adopting crossover, mutation, and elitism rates of 80%, 20%, and 10%, respectively. The main fitness function used weights (0.7, 0.2, 0.1), with alternative combinations explored in a sensitivity analysis.

Figure 2 shows the evolution of the mean fitness over generations for hospital occupancy scenarios of 60%, 75%, and 98%, considering the weight combination (0.7, 0.2, 0.1). The curves represent the average of the 30 runs, and the shaded areas indicate the 95% confidence interval, highlighting the convergence of the algorithm and the stability of the optimization process across runs.

Table II presents the average results of the genetic algorithm for different hospital occupancy scenarios considering weights (0.7, 0.2, 0.1). It is observed that the increase in occupancy reduces the average value of the fitness indicating better overall performance.

Additionally, the relative impact increases and inequality decreases with occupancy. In the 98% scenario, it is observed

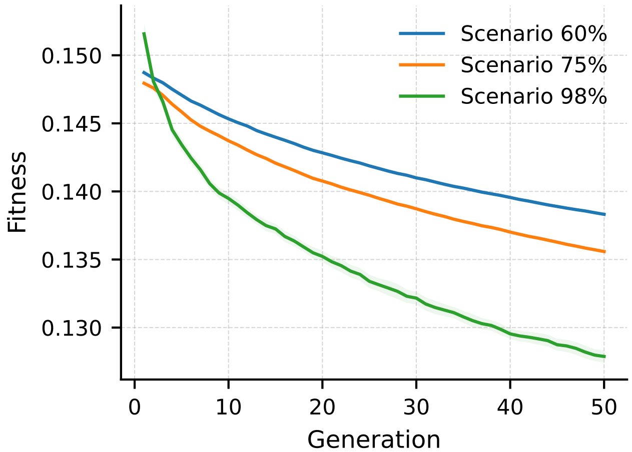

Fig. 2: Genetic Algorithm convergence across hospital occupancy scenarios (60%, 75%, and 98%), considering weights (0.7, 0.2, 0.1) in the fitness function. The curves represent the mean fitness over 30 independent runs, and the shaded areas indicate the 95% confidence interval.

Table II: Average Results of the GA by Occupancy Scenario

Scenario	Genetic Algorithm Metrics		
	Fitness	***Impact***	***Inequality / Cost***
Low (60%)	0.138 ± 0.0003	0.0077	0.664 / 0.00053
Medium (75%)	0.136 ± 0.0003	0.0081	0.649 / 0.00053
High (98%)	0.128 ± 0.0012	0.0102	0.604 / 0.00015

Fitness presented as mean $\pm$ standard deviation.

simultaneously greater reduction in mortality and lower inequality indicating that, under saturation, bed redistribution becomes more effective, with a better trade-off between impact equity and cost.

Table III details the results of the GA considering different weight combinations in the fitness function for each occupancy scenario. The observed patterns confirm that combinations with a higher weight on impact (0.7) favor solutions with better performance in mortality reduction, while maintaining balance with inequality and cost. The variation in fitness across runs is low, reinforcing the robustness of the algorithm.

Table III: Genetic Algorithm Results by Weight Combination.

Scenario 60%			
Weights	Fitness	Impact	Inequality / Cost
(0.7, 0.2, 0.1)	0.1383 ± 0.0003	0.00771	0.6644 / 0.00053
(0.6, 0.3, 0.1)	0.2039 ± 0.0004	0.00775	0.6641 / 0.00069
(0.5, 0.4, 0.1)	0.2694 ± 0.0006	0.00780	0.6635 / 0.00100
Scenario 75%			
(0.7, 0.2, 0.1)	0.1356 ± 0.0003	0.00809	0.6494 / 0.00053
(0.6, 0.3, 0.1)	0.1996 ± 0.0005	0.00815	0.6488 / 0.00068
(0.5, 0.4, 0.1)	0.2635 ± 0.0007	0.00817	0.6483 / 0.00059
Scenario 98%			
(0.7, 0.2, 0.1)	0.1279 ± 0.0012	0.01022	0.6035 / 0.00015
(0.6, 0.3, 0.1)	0.1858 ± 0.0017	0.01048	0.5984 / 0.00017
(0.5, 0.4, 0.1)	0.2446 ± 0.0028	0.01068	0.5981 / 0.00017

Finally, Table IV ompares the results of the GA with the baseline scenario. This baseline represents the observed allocation of hospital beds without any optimization process. The expected number of deaths in this scenario is obtained from the calibrated predictive model applied to the original dataset reflecting the current state of the healthcare system. In this way, all comparisons with the solutions generated by the GA are performed within the same predictive framework ensuring consistency in the evaluation.

It is observed that in low (60%) and medium occupancy (75%) scenarios the total number of predicted deaths increases slightly in relation to the baseline indicating that more intensive redistributions tend to prioritize equity and cost criteria with limited gains in mortality reduction. On the other hand, in high occupancy (98%) a reduction in the number of deaths is observed evidencing that under system saturation conditions small reallocations of beds can generate more direct impacts on mortality.

Additionally, it is observed that the redistribution of beds is more intense in less restrictive scenarios and more conservative in saturation situations with more localized adjustments. Among the evaluated weight combinations the configuration (0.7, 0.2, 0.1) presented more balanced solutions between impact equity and cost reflecting an adequate trade-off between the different criteria considered in the fitness function.

Table IV: Comparative evaluation between the real baseline and GA solutions.

Scenario	**Impact**	**Inequality**	**Cost**	**Deaths**	**Fitness**
Baseline	1.0000	0.7381	0.0000	1506	–
GA – Low (60%)	0.0073	0.7389	0.0009	1528	0.1377
GA – Medium (75%)	0.0074	0.7405	0.0002	1549	0.1348
GA – High (98%)	0.0069	0.7317	0.0000	1466	0.1254

D. Discussion

The results indicate that the performance of the genetic algorithm depends strongly on the level of occupancy of the hospital system. In high occupancy scenarios (98%), small redistributions of beds result in direct reductions in the number of deaths, suggesting that, under saturation conditions, the system becomes highly sensitive to variations in hospital capacity. In this context, the inclusion of the saturation factor in the fitness function amplifies the impact of allocation decisions, favoring solutions that prioritize mortality reduction.

On the other hand, in low and medium occupancy scenarios (60% and 75%), it is observed that the algorithm tends to propose more intensive redistributions, with greater focus on equity and cost, however, with limited gains or even an increase in the number of deaths in relation to the baseline. This behavior suggests that, in situations of lower pressure on the system, the relationship between bed availability and mortality is less direct, being influenced by other factors not modeled, such as clinical characteristics of patients and care dynamics.

From a comparative point of view, although the study uses the observed allocation of beds as a reference, simpler strategies, such as proportional redistributions based on population or number of cases, may be considered as a complementary reference in future analyses. Even so, the results indicate that the use of an evolutionary algorithm allows the exploration of non-trivial solutions in the search space, especially in high occupancy scenarios, in which small variations in allocation may generate relevant impacts on outcomes.

Another important aspect refers to the dependence of the optimization process on the predictive model. Although the selected model presented good performance (R^2 greater than 0.90), uncertainties in the predictions may influence the evaluation of candidate solutions, since mortality estimates are used directly in the fitness function. In this way, the obtained solutions must be interpreted considering this dependence, especially in scenarios in which small differences in mortality impact the final choice.

From a practical point of view, the application of the proposed approach in real healthcare systems involves additional challenges not explicitly modeled in this study. Aspects such as logistics of transferring patients between municipalities, availability of healthcare professionals, administrative constraints, and response capacity of hospital units may limit the direct implementation of the obtained solutions. Thus, the results must be interpreted as an exploratory analysis of the potential for resource redistribution, and not as a direct operational prescription.

In addition, the study considered exclusively SARI as the clinical outcome. In real scenarios, multiple conditions may simultaneously demand hospital resources, altering the dynamics of occupancy and the sensitivity of mortality to bed redistribution.

Although the analysis was conducted in the specific context of SARI, the structure of the approach does not depend on particular characteristics of this condition. However, its application in other contexts of healthcare resource allocation was not evaluated in this study and constitutes a direction for future investigations.

VI. Conclusion

This study presented an analytical approach to optimize hospital bed allocation by integrating predictive models of SARI mortality with a GA. Tree-based models, particularly the *Random Forest Regressor*, were able to capture the variability of mortality, outperforming classical count models in scenarios with overdispersion and zero-inflation.

The integration of predictive modeling with evolutionary optimization allowed the analysis of trade-offs between mortality reduction, equity, and redistribution cost. The results indicate that the performance of the GA depends on the level of hospital occupancy, with more noticeable effects under high-demand conditions.

These findings suggest that the proposed approach can be used to support the analysis of resource allocation strategies in constrained scenarios.

As limitations, the dependence on the quality of the predictive model and the lack of more detailed operational constraints, such as patient transfer logistics and healthcare workforce availability, are noted; in addition, the analysis considered only SARI, without accounting for competition for resources among different clinical conditions. Future work may include incorporating hospital occupancy rates into the predictive model, integrating external variables such as climatic factors, and evaluating the approach in broader contexts, including different bed specialties and multiple diseases.

Acknowledgment

Authors thank the financial support given the Brazilian National Council for Scientific and Technological Development - CNPq (grant n. 420212/2023-0) and the Minas Gerais Research Foundation – FAPEMIG (grant n. BDT-00010-25).

References

[1] N. B. Medeiros, F. S. Fogliatto, M. K. Rocha, and G. L. Tortorella, ”Predicting length of stay of pediatric patients using machine learning algorithms,” *International Journal of Production Research*, vol. 63, no. 2, pp. 483–496, 2023, doi: 10.1080/00207543.2023.2235029.

[2] Brazilian Institute of Geography and Statistics (IBGE), ”Population Estimates,” Dec. 30, 2024. [Online]. Available: https://www.ibge.gov.br/estatisticas/sociais/populacao/9103-estimativas-de-populacao.html. [Accessed: Jun. 6, 2025].

[3] X. Li, X. Yang, B. Dong, *et al.*, ”Prediction of 28-day all-cause mortality in ICU patients with pre-existing chronic heart failure using the stress hyperglycemia ratio,” *Cardiovascular Diabetology*, vol. 24, p. 10, 2025, doi: 10.1186/s12933-025-02577-z.

[4] Ministry of Health (Brazil), ”Influenza Syndrome/SARI: Risk Classification and Patient Management,” Jun. 5, 2025. [Online]. Available: https://bvsms.saude.gov.br/bvs/cartazes/sindrome_gripal_classificacao_risco_manejo.pdf. [Accessed: Jun. 5, 2025].

[5] Ministry of Health (Brazil), ”Hospital Bed Management,” Jun. 8, 2025. [Online]. Available: https://hcfamema.sp.gov.br/wp-content/uploads/2022/07/03_GESTAO_DE_LEITOS_REV_01.pdf. [Accessed: Jun. 8, 2025].

[6] E. Pedarzani, A. Fogangolo, I. Baldi, *et al.*, ”Prioritizing patient selection in clinical trials: A machine learning algorithm for dynamic prediction of in-hospital mortality for ICU-admitted patients using repeated measurement data,” *Journal of Clinical Medicine*, vol. 14, no. 2, p. 612, 2025, doi: 10.3390/jcm14020612.

[7] Ministry of Health (Brazil), ”SARI 2021–2024 Dataset (including COVID-19 data),” Sep. 5, 2022. [Online]. Available: https://dados.gov.br/dados/conjuntos-dados/srag-2021-e-2022. [Accessed: Jun. 6, 2025].

[8] Brazilian Institute of Geography and Statistics (IBGE), ”SIDRA – Demographic Census,” [Online]. Available: https://sidra.ibge.gov.br/tabela/4714. [Accessed: Jun. 6, 2025].

[9] Ministry of Health (Brazil), ”Hospitals and Beds Dataset,” OpenDataSUS, 2024. [Online]. Available: https://opendatasus.saude.gov.br/dataset/hospitais-e-leitos. [Accessed: Jun. 6, 2025].

[10] T. Vieira, ”Interpolation,” [Online]. Available: https://ic.ufal.br/professor/thales/icg/Aula1.2.pdf. [Accessed: Jun. 5, 2025].

[11] Directorate of Research (IBGE), ”Clarifications on Municipal Population Estimates,” Oct. 8, 2024. [Online]. Available: https://www.ibge.gov.br/np_download/novoportal/documentos_institucionais/Esclarecimentos_prestados_pela_DPE_estimativas_das_populacoes_municipais_20241018.pdf. [Accessed: Jun. 5, 2025].

[12] M. Kühn, T. Baumann, and H. Salzwedel, ”Genetic algorithm for process optimization in hospitals,” in *Proc. 26th European Conference on Modelling and Simulation (ECMS)*, 2012, pp. 1–5.

[13] S. Belciug and F. Gorunescu, ”A hybrid genetic algorithm–queuing multi-compartment model for optimizing inpatient bed occupancy and associated costs,” *Artificial Intelligence in Medicine*, vol. 68, pp. 59–69, 2016, doi: 10.1016/j.artmed.2016.03.001.

[14] S. Anumolu and M. Samant, ”Optimizing hospital bed capacity: How lean-AI integration prevents capacity constraints from multiplying demand,” *Intelligent Hospital*, vol. 2, p. 100048, 2026, doi: 10.1016/j.inhs.2026.100048.