Data-Driven Bed Assignment for Emergency Patients Using Supervised Learning

Hela Jedidi[1], Hajer Ben-Romdhane[1], Issam Nouaouri[2], Saoussen Krichen[1]

Abstract— **In this paper, we propose a hybrid approach for the optimal allocation of hospital beds to emergency patients, tackling the common challenge of unlabeled clinical data. Instead of relying on predefined labels, we use a rule-based labels step to assign initial risk levels based on key medical features. These labels are then used to train several supervised learning models, helping to improve risk prediction and explainability. The predicted levels: Minimum, Moderate, or Maximum, serve as input to an optimization model that assigns patients to beds while reducing transfer cost and supporting infection control. Results show promising improvements in fairness and allocation efficiency, with a comparative study highlighting the benefits of our method.**

*K*eywords: **Supervised learning, Patient classification, Linear programming, Rules-based labels.**

I. Introduction

Hospitals must continuously balance limited bed capacity against fluctuating demand, a task that becomes exceptionally difficult during emergencies. In such scenarios, every admission decision must reconcile two intertwined concerns: (i) *clinical triage*, which assesses each patient's severity and contagiousness to prevent cross-infection, and (ii) *operational allocation* [5], which assigns that patient to the most appropriate bed while respecting capacity, speciality, and isolation rules [15]. When waves of infectious patients arrive in rapid succession, as seen in the early stages of COVID-19, traditional 'first come, first served' heuristics quickly break down, causing delayed care, especially for severe patients.

To support triage, many studies employ supervised machine-learning (ML) models that transform raw clinical, laboratory, or demographic features into severity scores or risk categories. Gradient-boosted trees, support-vector machines (SVM), logistic regression (LR), and random forests (RF) dominate this literature. Yan *et al.* [7] achieved 0.95 precision in predicting COVID-19 mortality with XGBoost, while on a small cohort. Wong and So [8] analysed 93 variables from the UK Biobank, while SVM-based studies by Yao *et al.* [10] and Sun *et al.* [9] obtained accuracies around 0.80 on medium-sized datasets. Larger investigations using LR or RF reported AUCs between 0.85 and 0.95 on cohorts of several hundred to thousands of cases [13], [11], [12], [6]. These works demonstrate that ML can identify highly contagious or clinically critical patients early, guiding isolation and treatment priorities. Yet they stop short of turning those predictions into concrete bed-allocation plans. Also, these methods often rely on labeled datasets to train predictive models, which may not always be available in real-time emergency settings, limiting their direct applicability.

A parallel research stream addresses the Hospital Patient Bed Assignment Problem (HPBAP) from an optimization perspective. Because HPBAP is NP-hard [20], exact mixed-integer programming handles only small instances; larger hospitals rely on heuristics such as tabu search, simulated annealing, or hyper-heuristics to minimize transport distance [14], balance ward workloads, or enforce gender segregated wards [3], [2], [1]. These models, however, usually assume that each patient's acuity and contagious level are known, an assumption rarely valid in an emergency context.

Recognising that prediction and prescription are two sides of the same coin, a handful of studies have begun to couple ML outputs with optimization engines. FSchäfer *et al.* [4] embedded regression predictions in an integer model; Ahmed *et al.* [18] combined tree ensembles with explainable rules; and Jedidi *et al.* [16] linked KNN/SVM classifiers to a linear programming. While promising, these hybrid approaches still depend on fully labeled training data, even though labels are often missing, delayed, or noisy during surges.

Table I summarises representative contributions across three strands: pure ML triage, standalone optimization, and early ML-optimization hybrids. The comparison highlights a persistent gap: no existing framework simultaneously (i) generates patient severity labels when none are available, (ii) refines those labels with state-of-the-art classifiers, and (iii) plugs the resulting predictions into an optimization model that enforces infection-control constraints and minimizes patient transfer.

This paper closes that gap. We introduce an end-to-end decision-support system that begins with rule-based labels to create provisional contagion-severity labels, learns from those labels via RF, XGBoost, and LR models, and finally feeds the predicted categories into a linear programme that assigns patients to beds while satisfying assignment and capacity rules and minimising transfer costs.

[1]LARODEC laboratory, Université de Tunis, Institut Superieur de Gestion de Tunis, Avenue la liberté, bouchoucha, Le bardo, Tunisia hela.jedidi@isg.rnu.tn

[2]LGI2A laboratory, Université de l'Artois, UR 3926, Bethune, 62400, France

The rest of this paper is organized as follows: Section 2 describes the problem; Section 3 describes the hybrid solution methodologies; Section 4 presents computational experiments and analyses; Section 5 concludes with future directions.

TABLE I
COMPARISON OF HOSPITAL PATIENT BED ASSIGNMENT SOLUTION APPROACHES

Reference	Labeled Data	ML mthod	Optimization approach	ML-optimization hybrids
Wong and So, 2020 [8]	✓	XGBoost	-	-
Sun et al., 2020 [9]	✓	SVM	-	-
Sanchez et al., 2020 [13]	✓	Logistic Regression	-	-
Sumayh et al., 2020 [6]	✓	RF, others	-	-
Ahmed et al., 2022 [18]	✓	T-XGB, T-ADAB, and T-MLP. T-XGB	-	✓
Schäfer et al., 2023 [4]	✓	ML regression approaches	ILP	✓
Jedidi et al., 2024 [16]	✓	KNN, SVM	LP	✓
Proposed Method	-	Rule-Based + RF/XGBoost/LR	LP	✓

II. Problem Description

During public health emergencies, such as pandemics or disaster scenarios, hospitals are often overwhelmed by the influx of patients, particularly those in critical condition. In such contexts, healthcare systems are pressured to allocate limited resources, especially hospital beds, as efficiently and ethically as possible. This situation gives rise to what is known as the (HPBA) problem, a complex decision-making task that involves determining how to assign patients to appropriate beds based on medical urgency, contagiousness, and resource constraints.

A key challenge in this setting is that patient data upon arrival is often unlabeled, particularly in real-world scenarios where triage decisions must be made rapidly and with incomplete information. In the absence of explicit severity labels, the system must infer each patient's health status from available clinical features such as vital signs, symptoms, and test results. This step is crucial, as effective patient prioritization depends heavily on correctly identifying risk levels, typically categorized as high, moderate, or low risk: (MAR, MER, MIR).

Once risk levels are estimated, the next challenge involves assigning each patient to an available hospital bed, taking into account the suitability of the bed type (e.g., critical care, respiratory support, or general ward) and the constraints on capacity. In this context, hospitals prioritize patients with severe or highly contagious conditions, aiming to reduce mortality and limit the spread of infectious diseases. Due to the urgency of care delivery in emergencies, secondary factors such as gender-based room separation are often disregarded, with the primary goal being the rapid and effective use of available beds.

The HPBA problem thus emerges as a two-phase decision process under uncertainty: (1) classifying patients based on severity when explicit labels are unavailable, and (2) optimally allocating limited hospital beds based on the inferred patient needs.

III. A Hybrid Solution Approach for HPBAP

This paper presents a new approach to solving the HPBAP in emergency settings that deal with the clinical urgency and bed-type requirements to reflect the real challenges faced in critical care scenarios. During health crises, such as pandemics, hospital systems must rapidly and accurately assess incoming patients, classify their risk levels, and allocate them to appropriate and available beds. Addressing such complexity requires integrating medical expertise, data-driven decision-making, and optimization.

To tackle this multifaceted problem, we propose a two-phase solution framework that starts with classifying patients based on their clinical risk, followed by an optimization process to assign them to suitable hospital beds [17]. Our methodology begins by labeling patients using a rule-based process grounded in clinical knowledge and expert-defined thresholds. This rule-based triage system was applied to real-world data collected from the University Hospital in Sousse, Tunisia, between September and December 2020, consisting of demographic, clinical, and comorbidity information of COVID-19-positive patients.

The classification into three severity levels: Maximum Risk (MAR), Medium Risk (MER), and Minimum Risk (MIR), was guided by a well-defined set of clinical rules, visualized in Figure (1). For example, patients with oxygen saturation below 90% or systolic blood pressure lower than 90 mmHg, combined with critical symptoms such as severe shortness of breath, were automatically labeled as MAR. Those with moderate vitals and one or two underlying chronic conditions, such as diabetes or hypertension, were categorized as MER. MIR included patients with stable vitals and no or minor symptoms. This rule-based approach ensured consistent and reproducible triage, serving as a basis for automated classification through machine learning.

Following the expert-driven labeling, we proceeded to the classification phase (see figure 2). The data was preprocessed to improve model performance and robustness. Notably, since the dataset was imbalanced, we employed the Synthetic Minority Over-sampling Technique (SMOTE)

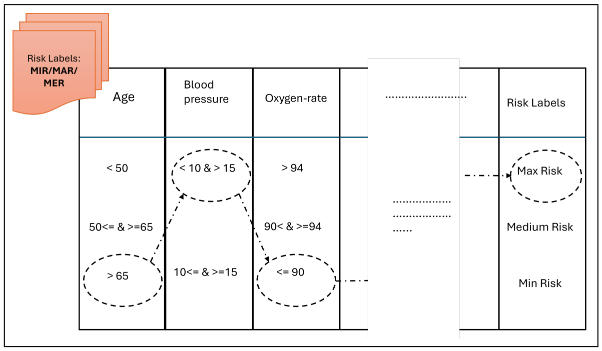

Fig. 1. An example of rule-based labeling construction

to generate synthetic samples for underrepresented classes, ensuring that all severity categories were adequately represented during training [19].

We then tested four machine learning algorithms: K-Nearest Neighbors (KNN), Random Forest (RF), Logistic Regression (LR), and Extreme Gradient Boosting (XGBoost), each evaluated using various hyperparameter configurations. KNN was selected for its simplicity and effectiveness in non-linear and small-sample settings. Random Forest provided robustness and interpretability by aggregating decision trees and reducing overfitting. Logistic Regression was used as a strong linear baseline and for its ease of interpretation. XGBoost, a powerful boosting technique, was included due to its ability to model complex feature interactions and handle missing data efficiently [8], [6], [9].

Once patients were classified into clinical risk groups, namely high-risk (P_{M_A}), medium-risk (P_{M_E}), and low-risk (P_{M_I})—the optimization phase was activated to assign each patient $p \in P$ to an appropriate hospital bed $r \in R$ for each planning day $d \in D$, while ensuring efficient and scalable resource usage across multiple hospitals. This phase was formulated as a linear programming model that aimed to minimize a total cost function composed of two components: the cost of assigning a patient p to a bed of type r, denoted by $C_{p,r}$, and the penalty cost incurred for transferring a patient from one bed type to another on day d, captured by a transfer indicator variable $t_{p,r,d}$, weighted by a parameter W_{Tr}.

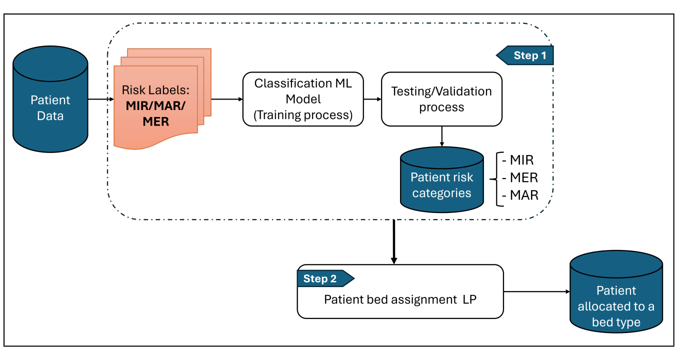

Fig. 2. Schematic overview of the solution methodology

The model used binary decision variables $x_{p,r,d}$, where $x_{p,r,d} = 1$ indicates that patient p is assigned to bed type r on day d, and $x_{p,r,d} = 0$ otherwise. Similarly, $t_{p,r,d} = 1$ implies that patient p is transferred from bed type r on day d. The objective function thus minimized the sum of all assignment costs $\sum_{p\in P, r\in R, d\in D} C_{p,r} \cdot x_{p,r,d}$ and transfer penalties $\sum_{p\in P, r\in R, d\in D} W_{Tr} \cdot t_{p,r,d}$.

Patient assignment followed medical prioritization rules:

- High-risk patients ($p \in P_{M_A}$) could be assigned to any bed type: critical ($r = 1$), respiratory ($r = 2$), or normal ($r = 3$);
- Medium-risk patients ($p \in P_{M_E}$) were limited to respiratory and normal beds ($r \in \{2, 3\}$);
- Low-risk patients ($p \in P_{M_I}$) were restricted to normal beds only ($r = 3$)

These conditions were enforced for each day $d \in D_p$, where D_p is the set of days patient p is hospitalized. Capacity constraints were also imposed: for every bed type r and each day d, the total number of patients assigned must not exceed the number of available beds Q_r, ensuring $\sum_{p\in P} x_{p,r,d} \leq Q_r$. To promote stability in patient care, the model penalized unnecessary transfers by enforcing that any change in bed assignment between consecutive days was reflected in the transfer variable, i.e., if $x_{p,r,d} \neq x_{p,r,d+1}$, then $t_{p,r,d} = 1$. This optimization formulation followed a "soft constraint" approach: while all efforts were made to satisfy constraints, the model allowed minor violations with associated penalties to simulate real-world healthcare variations under pressure.

IV. Experimental Results

This section presents the results obtained from our classification models applied to the patient severity classification task. The experiments were conducted following the methodology described earlier. The dataset was collected from the University Hospital in Sousse, Tunisia, and includes 20 instances across four configurations (Small1, Small2, Large1, Large2), with patient numbers ranging from 76 to 1500 and bed distributions of (BED1: 12–16 for critical care, BED2: 13–22 for respiratory care, BED3: 38–78 for normal care) over a 7-day period. In the largest configuration (Large2), there are 1500 patients divided into 893 MaxRisk, 570 MediumRisk, and 27 MinRisk cases, based on rule-based labeling strategies applied to clinical and demographic indicators. The dataset comprises clinical signs (e.g., temperature, SpO2, SOB), demographics (e.g., age, gender), and chronic illnesses.

Figure 3(a) illustrates the original distribution of patients across the three severity classes. The dataset is clearly imbalanced, with class 0 (high severity) comprising **60%**, class 1 (moderate severity) **38.2%**, and class 2 (low severity) only **1.8%** of the data. Such imbalance can lead to classifiers being biased toward predicting more frequent classes, potentially overlooking low-risk cases that are important for effective hospital resource management.

To mitigate this issue, we applied the *SMOTE (Synthetic Minority Over-sampling Technique)*, which synthetically generates samples for the minority classes. As shown in Figure 3(b), SMOTE effectively rebalanced the dataset, leading to a uniform distribution of **33.3%** across all three classes. This preprocessing step is essential to ensure fair model training and robust evaluation, especially when relying on metrics such as **F1-score** and **recall**, which are particularly informative under class imbalance conditions.

We evaluated the classification performance using four widely adopted machine learning models: **K-Nearest Neighbors (KNN)**, **Random Forest**, **XGBoost**, and **Logistic Regression**. For each model, we performed **hyperparameter tuning using cross-validation** to identify optimal configurations and ensure robust and generalized model performance. In table 4 (a), metrics are reported from a single evaluation (likely on a hold-out test set), whereas table 4 (b) presents macro-averaged scores using 5-fold cross-validation, which provides a more reliable estimate of generalization.

- **KNN**: Tested values for the number of neighbors $k \in \{3, 5, 7, 11, 19\}$. The best performance was achieved with $k = 5$, balancing bias and variance. Despite its simplicity, KNN showed moderate performance (Accuracy: 0.7533, F1-macro: 0.6078) before SMOTE, which improved post-balancing (Accuracy: 0.7778).

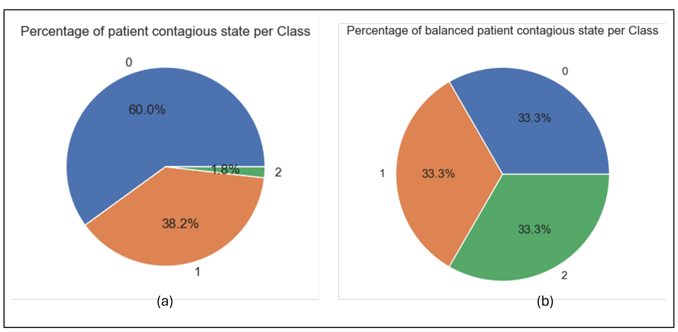

Fig. 3. (a) % of patient risk levels per class, (b) % of balanced patient risk levels per class

(a)

	Model	Accuracy	Recall (macro)	F1-score (macro)
0	KNN	0.7533	0.7803	0.6078
1	Random Forest	0.9867	0.9522	0.9693
2	XGBoost	0.9933	0.9963	0.9953
3	Logistic Regression	0.7600	0.6733	0.6389

(b)

	Modèle	F1-macro (CV)	Recall-macro (CV)	Accuracy (test)
0	Random Forest	0.9842	0.9936	0.9822
1	XGBoost	0.9896	0.9986	0.9911
2	Logistic Regression	0.5997	0.7019	0.7578
3	KNN	0.5704	0.6420	0.7778

Fig. 4. Model Performance Comparison: With (b) vs Without (a) Cross-Validation

- **Random Forest**: Evaluated across different $n\_estimators \in \{100, 200, 500\}$ and $max\_depth \in \{5, 10, 20\}$. Using cross-validation, the best result was obtained with $n\_estimators = 500$ and $max\_depth = 10$, achieving an impressive Accuracy of **0.9867** and F1-macro of **0.9693**.

- **XGBoost**: Fine-tuned using cross-validation across $learning\_rate \in \{0.01, 0.1, 0.2\}$, $max\_depth \in \{3, 6, 9\}$, and $n\_estimators \in \{100, 200\}$. The configuration $(learning\_rate = 0.1, max\_depth = 6, n\_estimators = 200)$ delivered the best results, with Accuracy: **0.9933** and F1-macro: **0.9953**, outperforming all other models.

- **Logistic Regression**: Tested with different regularization parameters $C \in \{0.01, 0.1, 1, 10\}$ and solvers (*liblinear*, *saga*). The best performance was achieved with $C = 1$, selected through cross-validation. However, this model consistently underperformed compared to others, with F1-macro around **0.6389**, indicating limited ability to capture complex nonlinear relationships in the data.

Across both pre- and post-balancing scenarios, **XGBoost and Random Forest** consistently outperformed the other models in terms of **macro F1-score** and **recall**. XGBoost shows a slight drop in accuracy (from 0.9933 to 0.9911) with cross-validation but retains exceptional F1 and recall scores, confirming its superiority. KNN and Logistic Regression, however, show more noticeable drops in recall and F1-macro under cross-validation, indicating sensitivity to data variation and less stability across folds.

Compared to existing works [16] in which patients are classified using the same set of features and only KNN and SVM models with various kernels are tested, showing that KNN with k=7 performs best with an accuracy of 0.94, our approach demonstrates superior performance. By introducing rule-based labels and applying data balancing techniques (SMOTE), our XGBoost-based model achieves even higher accuracy and better recall across all classes, confirming the added value of our method for real-world healthcare classification scenarios.

In summary, the application of SMOTE significantly improved classifier fairness and predictive power across all classes. Among the evaluated models, **XGBoost emerged as the most effective**, followed closely by **Random Forest**, making them well-suited for deployment in real-world hospital triage systems.

In the final stage of our methodology, we integrated the classification results into an optimization model aimed at assigning patients to appropriate bed types based on their severity levels. For large datasets, around 57

The assignment process was implemented using IBM ILOG CPLEX Optimization Studio v12.10, treating it as a resource-constrained optimization problem. In smaller data instances 1 and 2, the optimization model successfully admitted between 85% and 100% of critical patients, demonstrating strong performance under manageable demand. However, in larger datasets, the admission rates for high-risk patients dropped significantly to around 55% and 43% for Large sets 1 and 2, respectively. This decline reflects the increased competition for limited beds and the strain placed on hospital resources.

Additionally, the optimization cost increased substantially by around 10% compared to short and large sets, and computational time doubled, rising from 6.71 to 10.87 seconds. These results demonstrate a significant improvement of our model compared to existing approaches in the literature [16]. Specifically, our model increased the admission rate of critical patients from 20% in traditional methods to 43%, showcasing a greater capacity to handle high demand.

V. Conclusion

This study tackled the pressing issue of patient severity classification and bed allocation in emergency healthcare settings, where data is often unlabeled and highly imbalanced. We proposed a robust methodology that integrates rule-based labeling, SMOTE for dataset balancing, and supervised machine learning models to accurately classify patient risk levels.

Our experiments demonstrated the critical importance of cross-validation for both hyperparameter tuning and performance evaluation. Among the four models tested, K-Nearest Neighbors (KNN), Random Forest, XGBoost, and Logistic Regression. XGBoost consistently outperformed the others in terms of accuracy, recall, and F1-score, showing high robustness and generalization, particularly under 5-fold cross-validation.

In future developments, we plan to incorporate the classification output into a bi-objective optimization model that considers both patient severity and transportation costs. This extension will aim to optimize real-world hospital bed assignments by prioritizing critical patients while also minimizing delays and travel distances, ultimately improving emergency response efficiency.

References

[1] Ceschia, Sara, and Andrea Schaerf. *Local search and lower bounds for the patient admission scheduling problem.* Computers & Operations Research 38.10 (2011): 1452-1463.

[2] Liu, Haichao, Yang Wang, and Jin-Kao Hao. *Solving the patient admission scheduling problem using constraint aggregation.* European Journal of Operational Research (2024).

[3] Demeester, Peter, et al. *A hybrid tabu search algorithm for automatically assigning patients to beds.* Artificial Intelligence in Medicine 48.1 (2010): 61-70.

[4] Schäfer, F., Walther, M., Grimm, D. G., and Hübner, A. *Combining machine learning and optimization for the operational patient-bed assignment problem*. Health Care Management Science, (2023), 26(4), 785-806.

[5] Hong, Woo Suk, Adrian Daniel Haimovich, and R. Andrew Taylor. *Predicting hospital admission at emergency department triage using machine learning.* PloS one 13.7 (2018): e0201016.

[6] Aljameel, Sumayh S., et al. *Machine learning-based model to predict the disease severity and outcome in COVID-19 patients.* Scientific Programming (2021): 1-10.

[7] L. Yan, H.-T. Zhang, and Y. Xiao. *Prediction of criticality in patients with severe Covid-19 infection using three clinical features: a machine learning-based prognostic model with clinical data in Wuhan,* medRxiv, 2020.

[8] K. C. Y. Wong and H.-C. So, *Uncovering clinical risk factors and prediction of severe COVID-19: a machine learning approach based on UK biobank data,* medRxiv, 2020.

[9] L. Sun, F. Song, and N. Shi. *Combination of four clinical indicators predicts the severe/critical symptom of patients infected COVID-19* Journal of Clinical Virology, (2020), vol. 128p. 104431.

[10] H. Yao, N. Zhang, and R. Zhang. *Severity detection for the coronavirus disease 2019 (COVID-19) patients using a machine learning model based on the blood and urine tests* Frontiers in Cell and Developmental Biology, (2020), vol. 8, pp. 1–10.

[11] C. Hu, Z. Liu, and Y. Jiang. *Early prediction of mortality risk among patients with severe COVID-19, using machine learning,* International Journal of Epidemiology, (2020), vol. 49, no. 6, pp. 1918–1929.

[12] D. Bertsimas, G. Lukin, L. Mingardi et al., *COVID-19 mortality risk assessment: an international multi-center study* PLoS One, (2020), vol. 15, no. 12, p. e0243262.

[13] M. Sanchez-Montañes, P. Rodrıguez-Belenguer, A. J. SerranoLopez, E. Soria-Olivas, and Y. Alakhdar-Mohmara, *Machine learning for mortality analysis in patients with COVID-19* International Journal of Environmental Research and Public Health,(2020), vol. 17, no. 22, pp. 8386–20.

[14] JEDIDI, Hela, BEN-ROMDHANE, Hajer, NOUAOURI, Issam, and KRICHEN, Saoussen. *Bi-objective Hospital Bed Assignment Problem in Emergencies.* International COnference on Decision Aid and Artificial Intelligence (ICODAI 2024). Atlantis Press, 2025. p. 136-146.

[15] JEDIDI, Hela, NOUAOURI, Issam, BEN-ROMDHANE, Hajer, and KRICHEN, Saoussen. *Min-Conflict Heuristic Approach for Elective Patient Bed Assignment Problem.* 2024 10th International Conference on Control, Decision and Information Technologies (CoDIT). IEEE, 2024. p. 1424-1429.

[16] JEDIDI, Hela, BEN-ROMDHANE, Hajer, NOUAOURI, Issam, and KRICHEN, Saoussen. *A two-stage approach combining machine learning and optimization for the hospital patient bed assignment problem in emergencies*. Procedia Computer Science, 2024, vol. 246, p. 4316-4324.

[17] Hossain, S. K. M., Ema, S. A., & Sohn, H.. *Rule-Based Classification Based on Ant Colony Optimization: A Comprehensive Review*. Applied Computational Intelligence and Soft Computing, 2022(1), 2232000.

[18] Ahmed, A., Ashour, O., Ali, H., Firouz, M.. *An integrated optimization and machine learning approach to predict the admission status of emergency patients*. Expert Systems with Applications, 2022, 202, 117314.

[19] Fernández, A., Garcia, S., Herrera, F., & Chawla, N. V. . *SMOTE for learning from imbalanced data: progress and challenges, marking the 15-year anniversary*. Journal of artificial intelligence research, (2018), 61, 863-905.

[20] Vancroonenburg, W., D. Goossens, and F. Spieksma. *On the complexity of the patient assignment problem.* Tech. rep., KAHO Sint-Lieven, Gebroeders De Smetstraat 1, Gent, Belgium (2011).