Scarce, scarcer, scarcest: performance-flexible AI-based planning of elective surgeries for efficient and effective intensive care capacity management

Milena Grieger, Jens O. Brunner, Axel R. Heller, Christina C. Bartenschlager

Angaben zur Veröffentlichung / Publication details:

Grieger, Milena, Jens O. Brunner, Axel R. Heller, and Christina C. Bartenschlager. 2025. "Scarce, scarcer, scarcest: performance-flexible AI-based planning of elective surgeries for efficient and effective intensive care capacity management." *OR Spectrum*. https://doi.org/10.1007/s00291-025-00830-1.

OR Spectrum
https://doi.org/10.1007/s00291-025-00830-1

ORIGINAL ARTICLE

Scarce, scarcer, scarcest: performance-flexible AI-based planning of elective surgeries for efficient and effective intensive care capacity management

Milena Grieger[1] · **Jens O. Brunner**[1,2,5] · **Axel R. Heller**[3] · **Christina C. Bartenschlager**[3,4]

Received: 30 August 2023 / Accepted: 15 July 2025

Abstract
Operating room and intensive care unit (ICU) capacities belong to the scarcest resources in hospitals and strongly depend on each other. When planning elective surgeries, it is therefore important to consider both resources in an integrated way and to guarantee a certain flexibility in planning to avoid under- and overutilization, e.g., in the form of cancellations. In this work, we introduce a performance-flexible artificial intelligence (AI)-based planning approach for predicting whether an elective patient will be transferred to the ICU after elective surgery. This approach includes a performance-flexible loss function in a machine learning (ML) model and a subsequent simulation about ICU occupancy. The algorithm is evaluated by a large data set of the University Hospital of Augsburg, Germany, consisting of more than 26,600 elective surgeries between 2017 and 2021, and extensive simulation studies. This approach is generalizable as it uses data typically available during surgery planning in the outpatient clinic. Our findings demonstrate that, unlike state-of-the-art ML algorithms, our performance-flexible AI-based planning approach can prioritize a specific label in binary classification (i.e., ICU or non-ICU) subject to capacity considerations while maintaining high accuracy. This ensures a stable ratio of realized demand to planned ICU capacity that is close to 1 across different scenarios. Our performance-flexible AI-based planning algorithm outperforms state-of-the-art ML algorithms and supports hospital decision-makers with a flexible planning tool.

Keywords Medical decision-making · Machine learning · Binary classification · Integrated capacity planning · Loss function

Extended author information available on the last page of the article

Published online: 24 September 2025

1 Introduction

Operating room (OR) and intensive care unit (ICU) capacities belong to the scarcest resources in hospitals and strongly depend on each other: Approximately 10% of patients are transferred to the ICU after elective surgery, with considerable fluctuations over time. This finding aligns with Pearse et al. (2012) and is supported by similar patterns observed in our data. Effective planning of elective surgeries requires an integrated consideration of both OR and ICU capacities while maintaining planning flexibility to manage scarcity in both units (van Oostrum et al. 2008). Scarcity may be caused by staff shortages, mass casualties, or pandemics, for example (Heimerl and Kolisch 2010; Rodríguez-Espíndola 2023).

In recent years, particularly driven by the COVID-19 pandemic, artificial intelligence (AI) and machine learning (ML) algorithms have started to revolutionize healthcare planning and decision-making (e.g., Reig et al. 2020, Sheng et al. 2022). ML-based predictions can support the critical decision of whether a patient is transferred to the ICU (i.e., ICU or non-ICU[1]) after elective surgery. This decision involves especially two stakeholders (i.e., the anesthetist and the surgeon) along the patient's elective surgery pathway (see Fig. 1). In the current hospital process, clarity about how many elective patients will require ICU admission only emerges after all three decision points (D^2, D^1, and D^0) have been passed. The first appointment a patient receives before elective surgery is in the surgical outpatient clinic. At this initial visit, the surgeon provides a first assessment (D^2) of whether the patient will need ICU treatment after surgery (i.e., ICU or non-ICU). This early decision is crucial as it influences the appointment of the surgery and sets the stage for subsequent planning. Typically, this first assessment occurs several weeks before the planned surgery date, often around 20 to 30 days in advance. However, in practice, this planning window may be adjusted to a shorter period, depending on data availability and institutional workflows, to enable more reliable and up-to-date assessments closer to the surgery date. This early focus is particularly important when ICU beds are a scarce resource. By concentrating on D^2,

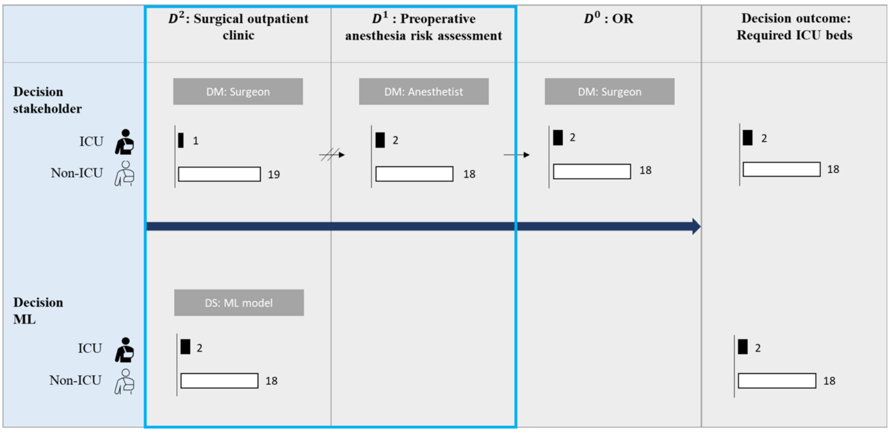

Fig. 1 The pathway of elective surgery in the hospital (top) vs.ML algorithm (bottom); DM: decision-making; DS: decision support

[1] Labels: ICU = 1, non-ICU = 0.

the ML model helps minimize the discrepancy between the surgeon's initial assessment and the anesthetist's later evaluation (D^1), leading to more accurate predictions and better alignment with actual ICU needs. For example, when free ICU capacity is limited, achieving high sensitivity is critical to avoid missing patients who truly require ICU admission. Following the initial assessment, the anesthetist performs a preoperative evaluation as soon as possible (D^1), which may modify the initial assessment (D^2). Ideally, this evaluation occurs on the same day as the surgeon's assessment, but due to practical constraints, it may take place up to one day before surgery. This timing carries the risk of cancellations if the revised assessment indicates a need for ICU that was not previously anticipated. During the surgery (D^0), unforeseen complications might necessitate ICU admission even if no ICU capacity was reserved. It is possible to apply ML models at each decision point (D^2, D^1 and D^0) to predict actual ICU needs, thereby improving the planning and management of resources. However, our focus is on D^2, where the surgeon's initial assessment significantly impacts the subsequent steps. By providing decision support at this stage, we aim to minimize the discrepancy between the surgeon's decision (D^2) and the anesthetist's later decision (D^1), and ultimately align these decisions with the actual number of ICU beds required. Focusing on D^2 allows for maximizing the lead time for decision-making and resource allocation, ensuring more effective planning and reducing the risk of surgery cancellations due to ICU capacity issues.

While we focus on long-term planning based on D^2, our approach also enables updated predictions closer to surgery (e.g., at D^1) by applying the model to newly available patient data. Although we do not explicitly integrate short-term adjustments into the initial prediction, the model implicitly captures variability through training on historical data, allowing for adaptive and flexible ICU capacity planning under variable conditions.

According to Heider et al. (2022), up to 35% of elective surgeries may be cancelled due to ICU capacity constraints, highlighting a significant area for improvement. This capacity issue necessitates a proactive approach to fully utilize OR capacity while effectively managing ICU capacity to avoid cancellations. Therefore, capacity is defined as the potential number of patients the OR and ICU can handle. Scarcity occurs when the feasible patient number is below typical capacity in either unit, while high capacity indicates exceeding typical numbers. The most critical scenario is when OR capacity is high, but ICU capacity is low, potentially leading to postponed elective surgeries or the bypassing of postoperative ICU monitoring. Pearse et al. (2012) indicate that many postoperative deaths could have been prevented with adequate ICU care. While state-of-the-art ML applications concentrate on overall prediction accuracy, our approach integrates OR and ICU capacity considerations, allowing for planning flexibility with a dynamic focus on sensitivity or specificity. For example, when ICU capacity is scarce, high sensitivity in predictions is crucial, whereas if ICU capacity is sufficient but OR capacity is scarce, specificity becomes more important.

In this work, we introduce a performance-flexible AI-based planning approach for predicting whether an elective patient will be transferred to the ICU after elective surgery (i.e., ICU or non-ICU). This approach includes a performance-flexible loss function in an ML model and a subsequent simulation about ICU occupancy. Here, performance-flexible refers to the ability of the user to adjust a weighting parameter within the model to prioritize sensitivity or specificity according to current hospital capacity or planning priorities. The approach is evaluated by a large data set of the University Hospital of Augsburg, Germany, consisting of more than $26{,}600$ elective surgeries between

2017 and 2021 and extensive simulation studies. This approach is generalizable as it uses data typically available during surgery planning in the outpatient clinic. Our findings demonstrate that, unlike state-of-the-art ML algorithms, our performance-flexible AI-based planning approach can prioritize a specific label in binary classification (i.e., ICU or non-ICU) subject to capacity considerations while maintaining high accuracy. This ensures a stable ratio of realized demand to planned ICU capacity that is close to 1 across different scenarios. Our approach reduces overbooking for normal situations dramatically, i.e., by 30% compared to standard (ML) approaches, and further supports decision-makers in hospitals while guaranteeing planning flexibility.

Our work is structured as follows. Section 2 provides an overview of related literature. In Sect. 3, we describe our performance-flexible AI-based planning approach, including loss functions, thresholds, ML models, the data set, and our simulation study. An overview of the findings is provided in Sect. 4. Section 5 summarizes our work and outlines future research directions.

2 Related literature and methodological background

The planning of surgeries, including ICU capacities, is crucial due to the scarcity of resources. In recent years, several publications have explored operations research methods, including optimization models and applications of ML models, to support this planning (see Sect. 2.1). Another way to enhance this planning is through the methodological adaption of ML models. Possible adjustments include modifying loss functions and thresholds, commonly called cost-sensitive learning (see Sect. 2.2). An overview of the related literature is given in Table 1. In this paper, we narrow our focus to binary classification, because we are interested in the binary decision of ICU treatment after elective surgery (i.e., ICU or non-ICU). We investigate two popular ML algorithms, namely logistic regression (LR) and deep neural networks (DNN), to support the planning of OR and ICU capacities using ML models.

2.1 OR and ICU capacity planning

The literature on OR scheduling and ICU capacity management highlights their close interconnection. Efficient OR scheduling is crucial as it directly impacts patient flow, resource utilization, and healthcare cost, and even the survival of the patient. Research has explored various methods to optimize OR schedules while considering ICU demands. Approaches include cyclical scheduling to enhance OR utilization and reduce ICU refusals (van Houdenhoven et al. 2008), stochastic programming models to handle uncertainties in surgery duration and patient stay (Jebali and Diabat 2015, Neyshabouri and Berg 2017), robust optimization techniques for reliable OR planning (Jebali and Diabat 2017, Neyshabouri and Berg 2017), and guidelines on a tactical level (Bai et al. 2016, Rachuba et al. 2022).

Strategic decisions based on contribution margin and growth potential (Wachtel and Dexter 2008), as well as the coordination of ORs, ICUs, and wards (Demeulemeester et al. 2013), are emphasized for efficient capacity management. Heuristic methods, like quota systems, help balance ICU bed occupancy (Heider 2022). These strategies aim to optimize cost, resource use, and patient flow amidst OR scheduling complexities (Vargas et al. 2009). Integrating ML models for ICU capacity could enhance efficiency and patient-centered resource management. Various ML techniques, such as random forests, neural networks, and decision trees, have been used to forecast ICU bed occupancy and patient admission risks (e.g., Lorenzen et al. 2021; Qian et al. 2021; Schiele et al. 2021). These tools have proven effective in predicting ICU requirements days in advance, aiding capacity planning and resource allocation (Goic et al. 2021; Lorenzen et al. 2021). However, existing literature focuses on standard prediction methods without modifications to loss functions or other methodologies for a flexible focus on specific classes.

Despite advancements, a performance-flexible AI-based planning system for elective surgeries to manage intensive care capacity efficiently remains, to the best of our knowledge, unexplored. Such a system could significantly enhance decision-making and complement traditional OR scheduling, providing a robust tool for managing hospital resources and improving patient care. In the context of developing more advanced AI systems, cost-sensitive learning plays a crucial role.

2.2 Cost-sensitive learning

Cost-sensitive learning focuses on data balancing and emphasizing different performance measures, often crucial in healthcare. Cost can represent non-monetary factors like disease severity, making it more beneficial to predict the positive class even if the negative class has a higher probability (Elkan 2001) (see Table 2 for notation).

2.2.1 Weighting

A key approach in cost-sensitive learning is manipulating the loss function to minimize errors between label and prediction. In binary classification, the loss function aims to reduce false negatives (FN) and false positives (FP) by introducing weights α for the positive class and β for the negative class, which is represented by Eq. (1).

$$J = -\frac{1}{M}\sum_{m=1}^{M}\left[\alpha \cdot y_m \cdot \log\left(h_{\boldsymbol{\theta}}\left(\boldsymbol{x}_m\right)\right) + \beta \cdot \left(1 - y_m\right) \cdot \log\left(1 - h_{\boldsymbol{\theta}}\left(\boldsymbol{x}_m\right)\right)\right] \quad (1)$$

The weights α and β vary across different loss functions (see Table 3 For instance, *binary cross-entropy (BCE)* is a state-of-the-art loss function that penalizes misclassification equally, i.e., $\alpha = \beta = 1$ (Murphy 2022). Other *BCE*-based loss functions focus on data balancing, essential for healthcare applications. Examples are the *weighted binary cross-entropy (WBCE)* (Jadon 2020) or the *cost-sensitive cross-entropy (CSCE)* (Aurelio et al. 2019). Further loss functions can be found in Tables 1 and 3.

2.2.2 Thresholding

In addition to influencing learning through the loss function, the outcome (i.e., the performance measurements) can also be changed through thresholding. The threshold parameter τ converts predicted probabilities p into specific labels based on a predetermined value (see Eq. 2).

Table 1 Literature overview subject to OR and ICU capacity planning and cost-sensitive learning; CSCE: cost-sensitive cross-entropy; RWWCE: real-world-weight cross-entropy; WBCE: weighted binary-cross-entropy; WBCE: weighted binary cross-entropy; BCE: binary cross-entropy; (N)PFBCE: (normalized) performance-flexible binary cross-entropy

Part	Authors	Review	Healthcare	OR capacity planning	ICU capacity planning	Binary classification	Loss function	Data balancing	Flexible perfor-mance
OR and ICU capacity planning	Bai et al. (2016)	x	x	x	x				
	Demeulemeester et al. (2013)	x	x	x					
	Goic et al. (2021)		x		x				
	Heider (2022)		x	x	x				
	Jebali and Diabat (2015)		x	x	x				
	Jebali and Diabat (2017)		x	x	x				
	Lorenzen et al. (2021)		x		x				
	Neyshabouri and Berg (2017)		x	x	x				
	Qian et al. (2021)		x		x	x			
	Rachuba et al. (2022)		x	x	x				
	Schiele et al. (2021)		x	x	x				
	van Houdenhoven et al. (2008)		x	x	x				
	Vargas et al. (2009)		x	x					
	Wachtel and Dexter (2008)	x	x	x					
Cost-sensitive learning	Aurelio et al. (2019)					x	*CSCE*	x	
	Cui et al. (2019)						*Class-balanced*	x	
	Drummond and Holte (2003)							x	
	He and Garcia (2009)	x						x	
	Ho and Wookey (2020)					x	*RWWCE*	x	
	Jadon (2020)	x				x	*WBCE*	x	
	Lin et al. (2017)					x	*Focal*	x	
	Lipton et al. (2014)					x			x
	Maalouf and Siddiqi (2014)					x		x	

Table 1 (continued)

Part	Authors	Review	Healthcare	OR capacity planning	ICU capacity planning	Binary classification	Loss function	Data balancing	Flexible perfor-mance
	Murphy (2022)					x	*BCE*	x	
	Mushava and Murray (2024)					x			
	Rezeai-Dastjerehei et al. (2020)					x		x	
	Sheng and Ling (2006)		x			x			
	Sheng et al. (2022)		x			x			
	Shrivastava (2020)					x	*Adma*		
	Zare et al. (2013)		x						
	Zhou and Liu (2006)					x		x	
	Our approach		**x**	**x**	**x**	**x**	***(N)PFBCE***		**x**

Table 2 Notation of machine learning models

Mathematical formulations			
N	Number of features	α	Weight of positive label
M	Number of training samples	β	Weight of negative label
y_m	Target label for training sample m	p	Predicted probability
$\widehat{y}_m$	Predicted label for training sample m	τ	Threshold parameter
$\boldsymbol{x}_m$	Features for training sample m	γ	Learning rate
$h_{\boldsymbol{\theta}}$	Prediction model with weights $\boldsymbol{\theta}$	J	Loss function

Table 3 Overview of methods including abbreviations, terms, and objectives; M: method (including learning (L) vs. no learning (NL) in training, T: thresholding, W: Weighting, FLP: flexible performance, BD: balanced data, UD: unbalanced data

Abbreviation	Term	M	Parameter	Input	Objective
BCE	*Binary cross-entropy*	–	$\alpha, \beta = 1$	–	–
WBCE	*Weighted binary cross-entropy*	W (L)	$\alpha \in \mathbb{R}_0^+, \beta = 1$	UD	BD
CSCE	*Cost-sensitive cross-entropy*	W (L)	$\alpha, \beta \in \mathbb{R}_0^+$	UD	BD
RWWCE	*Real-world-weight cross-entropy*	W (L)	$\alpha, \beta \in \mathbb{R}_0^+$	UD	BD
BCET	*Binary cross-entropy with threshold*	T (NL)	$\tau \in [0, 1]$	UD, BD	BD, FLP
PFBCE	***Performance-flexible binary cross-entropy***	***W (L)***	$\alpha, \beta \in$ [0,1]	***BD***	***FLP***
NPFBCE	***Normalized performance-flexible binary cross-entropy***	***W (L)***	$\alpha, \beta \in$ [0,1]	***BD***	***FLP***

$$\hat{y}_m = \begin{cases} 1, & if\ p > \tau \\ 0, & if\ p \leq \tau \end{cases} \tag{2}$$

For binary classification, the state-of-the-art value is $\tau = 0.5$, meaning that the positive class is predicted ($\widehat{y}_m = 1$) if $p > 0.5$, and the negative class is predicted ($\widehat{y}_m = 0$) if $p \leq 0.5$. This threshold can be adjusted; for instance, setting $\tau = 0.2$ increase sensitivity, as all probabilities above 0.2 result in a positive prediction. A detailed explanation of thresholds is given in the Appendix.

Individual thresholds make algorithms cost-sensitive even after training (Sheng and Ling 2006). Thresholding addresses unbalanced data and maximizes certain performance measures, such as sensitivity and F1 score, which balances the trade-off between correctly identifying positive cases and overall detection capability (Lipton et al. 2014). Zhou and Liu (2006) demonstrate that sampling and thresholding yield good results for unbalanced data sets, particularly in binary classification. Moreover, thresholding is easy to implement and apply (Sheng and Ling 2006). However, a major limitation is that thresholding cannot be implemented during the training phase, meaning the model cannot learn existing preferences or influence feature importance.

The literature review underscores a significant gap: despite recent strides, an innovative, performance-flexible AI-based planning system tailored for managing intensive care capacity in elective surgeries remains, to the best of our knowledge, unexplored. Such a system holds promise for substantially enhancing decision-making in traditional OR scheduling, thereby optimizing hospital resources and improving patient care. In the context of advancing AI systems, the pivotal role of cost-sensitive learning is evident. Currently, while existing loss functions address data balancing, none effectively integrates flexible performance objectives into model training. Our contribution seeks to fill this gap by proposing a performance-flexible AI-based planning approach that uniquely integrates cost-sensitive learning within the clinical workflow. By employing a weighted loss function, the model enables the prioritization of ICU admissions based on patient need while maintaining high predictive performance. This allows decision-makers, such as surgeons, to make an informed first assessment (D^2) of whether the patient needs ICU treatment after elective surgery and provide nuanced insights for medical professionals involved in ICU capacity management.

3 Methodology

In this section, we introduce our new performance-flexible AI-based planning approach. We begin by introducing our performance-flexible loss function for binary classification, where the *CSCE* is the basis for our loss function. Additionally, we use the three comparative loss functions *BCE* in combination with thresholding (*BCET)*, *WBCE,* and *real-world-weight cross-entropy* (*RWWCE)*. We then describe the ML models and the data set employed in our study. Finally, we conduct a simulation focusing on ICU capacity management to evaluate the new approach.

When predicting patients requiring critical care or a particular disease (i.e., ICU and non-ICU), it is essential for a healthcare decision-maker to prioritize a specific class. Therefore, we consider the performance measures sensitivity (true positive rate), specificity (true negative rate), and accuracy (overall correctness). These metrics can be influenced in different ways in training, e.g., by loss functions, and in testing, e.g., by thresholding (see Sect. 2).

3.1 Performance-flexible loss function

We address a supervised learning problem for binary classification of patients, evaluated by specific performance measures. While standard models aim to maximize accuracy, various clinical scenarios necessitate different focuses, such as a higher focus on potential ICU admissions. The notation for the following introduction of loss functions can be seen in Table 2.

Our performance-flexible AI-based planning approach introduces the *performance-flexible binary cross-entropy (PFBCE)* loss function, which builds on the *CSCE* but adjusts the weight parameter differently. The positive class weight α in *PFBCE* allows flexibility in performance prioritization but requires a balanced data set as foundation. To emphasize both classes equally and focus on accuracy, the positive class weight α can be set to 0.5, which can be used as state-of-the-art approach.

$$PFBCE : \alpha \in [0{,}1] \, , \beta = (1 - \alpha) \tag{3}$$

For higher sensitivity, which is important when ICU capacity is low, the positive class weight α can be set greater than 0.5 (e.g., $\alpha = 0.8$). This reduces the emphasis on correctly predicting the negative class, i.e., no focus on specificity. Conversely, to focus on the second class, i.e., non-ICU, when ICU capacity is normal, α can be set smaller than 0.5 (e.g., $\alpha = 0.2$). Extreme values in the lower range ($\alpha = 0.0$) and upper range ($\alpha = 1.0$) are generally not recommended as they skew the model's learning significantly.

The efficacy of the *PFBCE* approach can initially be justified through the following mathematical expressions. Like the *BCE*, the *PFBCE* is also a convex function, because an optimization approach like gradient descent should minimize the losses. For the *BCE*, the parameter update of the gradient descent would be calculated as in Eq. (4).

$$Gradient\ descent\ BCE : \theta{:=}\theta - \gamma \cdot \frac{1}{M} \sum_{m=1}^{M} \left[\left(h_{\boldsymbol{\theta}} \left(\boldsymbol{x}_m \right) - y_m \right) \cdot \boldsymbol{x}_m \right] \tag{4}$$

In this formulation, the parameter γ denotes the learning rate controlling the step size during optimization, while M indicates the number of samples. For the *PFBCE*, the convexity of the function is preserved by a linear convex combination of both terms (see Eq. 5).

$$Gradient\ descent\ PFBCE : \theta{:=}\theta - \gamma \cdot \frac{1}{M} \sum_{m=1}^{M} \left[\left(\alpha \cdot h_{\boldsymbol{\theta}} \left(\boldsymbol{x}_m \right) - (1-\alpha) \cdot y_m \right) \cdot \boldsymbol{x}_m \right] \tag{5}$$

Here, the parameter α allows for explicit control over the trade-off between sensitivity and specificity during training, enabling the model to prioritize positive or negative classes as needed while maintaining effective convergence properties during optimization.

3.2 Normalization

Normalization of the loss function is essential for addressing data anomalies, such as different features or biased values, and ensuring comparability across data sets. By normalizing the loss function, large differences between features are prevented, which avoids problems with biased values (Ma et al. 2020) and ensures that each feature contributes proportionally to the model's learning process. This adjustment enhances the overall model stability and performance. Additionally, normalization standardizes the loss function across different data sets or experimental conditions, making it easier to interpret and compare model performance reliably, regardless of data variations. This approach supports more consistent evaluation metrics, such as accuracy and sensitivity, which are crucial for robust machine learning applications. However, it is important to note that normalization may also influence the sensitivity–specificity trade-off. While normalization helps stabilize training, it can reduce the relative weighting effect of the performance-flexible loss components, potentially limiting the model's ability to strongly prioritize sensitivity or specificity when required. Therefore, careful tuning of the normalization procedure and weighting parameters is necessary to balance stability and flexibility in optimizing for different clinical priorities.

In this work, normalization always refers to the normalization of the loss function. One way to achieve this is to divide the *PFBCE* by the sum of the weights for each class. This leads to the *normalized performance-flexible binary cross-entropy (NPFBCE)* loss function, which is represented by Eq. (6).

$$J_{NPFBCE} = -\frac{1}{M}\sum_{m=1}^{M}\frac{\left[\alpha \cdot y_m \cdot log\left(h_{\boldsymbol{\theta}}\left(\boldsymbol{x}_m\right)\right) + \beta \cdot \left(1-y_m\right) \cdot log\left(1-h_{\boldsymbol{\theta}}\left(\boldsymbol{x}_m\right)\right)\right]}{\alpha \cdot y_m + \beta \cdot \left(1-y_m\right)} \quad (6)$$

The user of the ML algorithm can decide whether to normalize the loss function by changing the corresponding hyperparameter. This grants the flexibility to adapt the loss function based on the situation and, consequently, the data set.

3.3 Machine learning models

Since loss functions are used in several supervised learning methods, we use LR and a DNN as examples. LR is often used together with weighted LR (Das et al. 2013), while DNNs are widely known for their strong performance across many applications (Cichy and Kaiser 2019). For sake of comparability, the same ML models are applied to all loss functions. For the DNN, hyperparameter tuning is conducted with the *NPFBCE* loss function with a positive class weight of $\alpha = 0.8$, reflecting our need for high sensitivity. This results in a DNN architecture with one input layer, six hidden layers, and one output layer. The tanh activation function is used in all layers except the output layer, which uses the sigmoid activation function (see Fig. 2).

For training, we use the loss functions *BCET*, *NPFBCE*, *PFBCE*, *RWWCE*, *WBCE* presented above and stochastic gradient descent as the optimizer. Additionally, we use 85 epochs and a batch size of 10. To avoid overfitting, we use fivefold cross validation. For a detailed evaluation of the results, we use the metrics accuracy, sensitivity, and specificity.

It is important to note that *WBCE* and *RWWCE* inherently have the objective of providing balanced data. Thus, using a pre-balanced data set affects their initial state and objective. The *(N)PFBCE* can be seen as a specific case of *RWWCE* (e.g., $\alpha = 0.6$ and $\beta = 0.4$). However, *RWWCE* typically uses real-world impact estimates as weights, which are usually real numbers.

3.4 Data

Both LR and DNN are applied to a real-world healthcare data set from the University Hospital of Augsburg, Germany, covering 26,677 patients from 2017 to 2021. The task is to predict whether a patient would require intensive care after elective surgery (i.e., ICU and non-ICU). All necessary data is available at the first decision period D^2 in the surgical outpatient clinic (see Fig. 1). In the data set, approx. 10% of elective patients require subsequent ICU treatment. The data set includes 14 features (see Table 4), and each planned surgery is labeled with a binary indicator of ICU admission (i.e., 1 for ICU and 0 for non-ICU).

Extensive data preparation is conducted, including removing non-relevant or erroneous data, and summarizing comorbidities using the charlson comorbidity index (CCI) (Charlson et al. 1987). We use the synthetic minority oversampling technique (SMOTE) to balance the data (Chawla et al. 2002). Feature scaling and imputation of missing values are performed using the iterative imputer with the random forest algorithm. The split of training and test data is 90–10%.

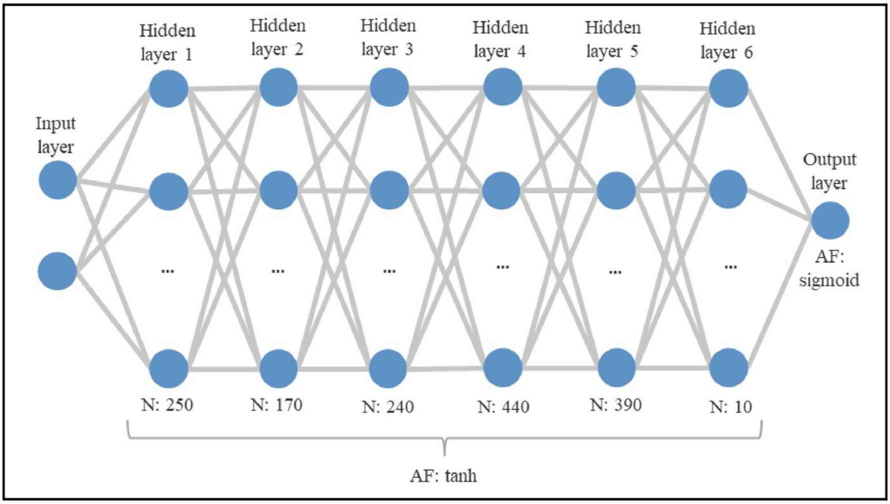

Fig. 2 Overview of the DNN applied to the data set (AF: activation function, N: neurons)

Table 4 Description of the 14 features of the data set (CCI: Charlson Comorbity Index)

No	Feature	No	Feature	No	Feature
1	Medical specialty	6	Sex	11	Number of comorbidities
2	Estimated cut suture duration	7	Weight	12	CCI
3	Estimated anes-thesia duration	8	Height	13	Planned type of anesthesia
4	Estimated surgery duration	9	Body mass index	14	Estimated ASA-score
5	Age	10	Main diagnosis		

3.5 Simulation methodology

The major goal of the simulation study is to evaluate varying weights in the performance-flexible AI-based planning approach for different scenarios. We define nine scenarios based on varying (planned) OR and ICU capacity patterns (see Table 5), estimated from the data set introduced above of the central ORs at the University Hospital of Augsburg, Germany.

OR capacity is categorized as scarce, normal, or high, defined by 20, 30, or 40 elective surgeries per day. ICU capacity is defined by the ratio of 5%, 10%, or 20% ICU patients per day (i.e., patients with ICU stay needed). In addition, we define three different patient cohorts: all patients, patients with ASA-scores[2] smaller than 3 or greater than 2. For every weight α, scarcity pattern and patient group, Monte-Carlo simulations with $1{,}000$ runs are applied. Each run involves simulating patients predicted to be transferred to the ICU after elective surgery (i.e., ICU) and those predicted not to require ICU care after elective

[2] ASA-score is a extensively employed scoring system utilized for categorizing patients based on their physical condition (Saklad (1941)).

surgery (i.e., non-ICU). The key performance indicator (KPI) is the ratio of realized ICU patients based on the anesthetist's assessment (D^1) and planned ICU patients by the surgeon with ML support (D^2), with a target and benchmark ratio of 1 indicating perfect prediction. A KPI greater than 1 indicates underestimation of ICU needs, while a KPI smaller than 1 indicates overestimation. Please find a flowchart of the simulation study in Fig. 3.

For the performance-flexible AI-based approach, we use Python for the ML models and R for the simulation study.

4 Results

In the following, we present our results by applying the methods to the ICU and elective surgery data. First, the results of the two ML models, LR and DNN are discussed. Afterwards, we present the simulation results with a focus on ICU capacity management. Below, absolute values are expressed as percentages ($\%$), and differences are presented as percentage points (PP).

4.1 Comparing LR with DNN

We evaluate the models using the two loss functions, *PFBCE* and *NPFBCE*, and compare them with the methods *BCET*, *WBCE* and *RWWCE* presented in the literature section. The application of the *NPFBCE* loss function to the data set confirms the previous assumptions; increasing the weight α enhances sensitivity while reducing specificity, with no significant change in accuracy. The results of the LR and DNN are shown in Fig. 4.

A higher weighting of *NPFBCE* generally allows a focus on one performance measure, namely sensitivity or specificity, which can assist healthcare decision-makers in capacity planning. For example, for situations where ICU capacity is scarce, it is crucial to prioritize the accuracy of the positive class, i.e., sensitivity. By adjusting the weights in the *NPFBCE* loss function, higher sensitivity can be achieved. With a positive class weight of $\alpha = 0.8$, *NPFBCE* in the DNN (in LR) achieves a sensitivity of 98.29% (97.73%), a specificity of 76.78% (78.31%), and an accuracy of 87.53% (88.02%).

The example shows that both models can achieve high accuracy while focusing on a specific performance measure, e.g., sensitivity, when ICU capacity is limited. Conversely, if ICU capacity is high, there is no need to focus on sensitivity.

Using the base case with a positive class weight of $\alpha = 0.5$, corresponding to *BCE* and state-of-the-art ML approaches, results in a sensitivity of 93.08% (91.33%), a specificity of 85.46% (86.65%) and an accuracy of 89.27% (88.99%) for *NPFBCE* for the DNN (LR). The use of the extreme values, i.e., $\alpha = 0.0$ or $\alpha = 1.0$, is not recommended because either the sensitivity or the specificity is zero. As the results

Table 5 Scenarios of the simulation study

Scenario	OR capacity [surgeries]		
Ratio of ICU patients	20 (scarce)	30 (normal)	40 (high)
0.05 (scarce)	1	4	7
0.10 (normal)	2	5	8
0.20 (high)	3	6	9

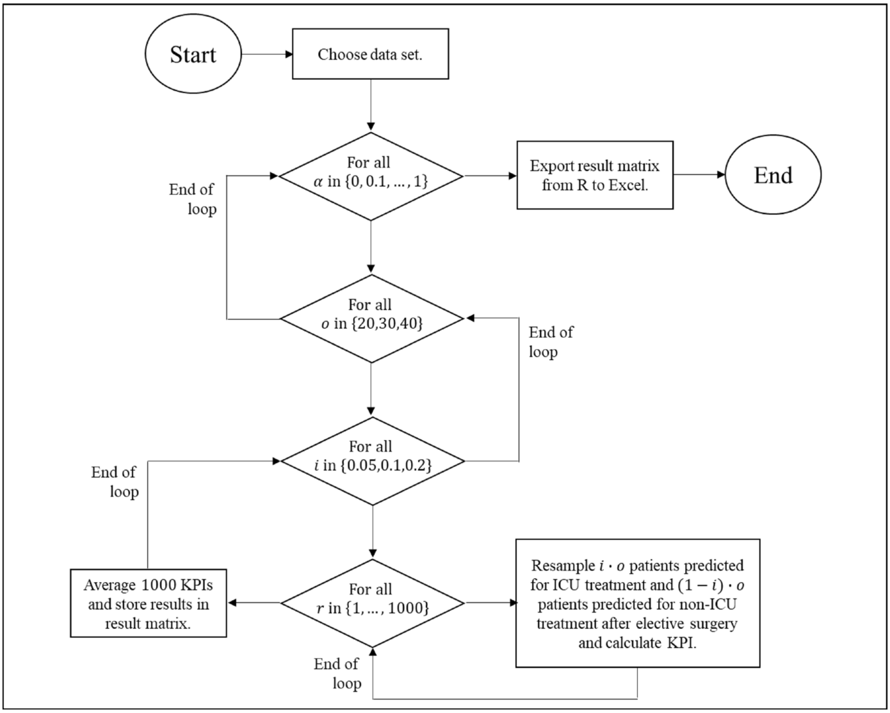

Fig. 3 Flowchart of the simulation study; α: weight of positive class, o: OR capacity, *i*: ICU capacity, r : index for runs

of the LR are similar to those of the DNN, but especially the sensitivity of the latter is higher, we focus on the DNN for further analyses.

Until now, we have solely examined the results of the *NPFBCE* loss function. However, these findings can also be applied to the non-normalized *PFBCE* loss function. When comparing *PFBCE* loss function with the normalized version, there is no distinct difference. For example, for $\alpha = 0.7$ and the DNN, we obtain a sensitivity of 97.37% (95.92%), a specificity of 79.24% (81.68%) and an accuracy of 88.31% (88.80%) for *NPFBCE* (*PFBCE*). In the remainder, we use the results of *NPFBCE* for the simulation performed in the following section.

A comparison of the results of the *(N)PFBCE* loss functions with the loss functions in literature confirms our previous findings. For example, with a high sensitivity setting $(\alpha = 0.8)$ in DNN, *NPFBCE* (*PFBCE*) achieves an improved sensitivity of $+7.05PP$ $(+7.04PP)$, a reduced specificity of $-10.32PP$ $(-9.70PP)$, and a reduced accuracy of $-1.65PP(-1.33)$ compared to *WBCE*. This trend is also evident when comparing *(N)PFBCE* with *RWWCE*. These results show that our approach places higher emphasis on sensitivity than *WBCE* and *RWWCE* while maintaining comparable accuracy. Only *BCET* achieves similar results to *(N)PFBCE* function in terms of sensitivity. With $\alpha = 0.8$, *NPFBCE* (*PFBCE*) shows a decreased sensitivity of $-2.56PP$ $(-3.04PP)$, an increased specificity of $+5.91PP$ $(+6.80PP)$, and an increased accuracy of $+1.68PP$ $(+1.88PP)$ compared to *BCET*. Thus, while *BCET* slightly outperforms in sensitivity, it has methodological limitations compared

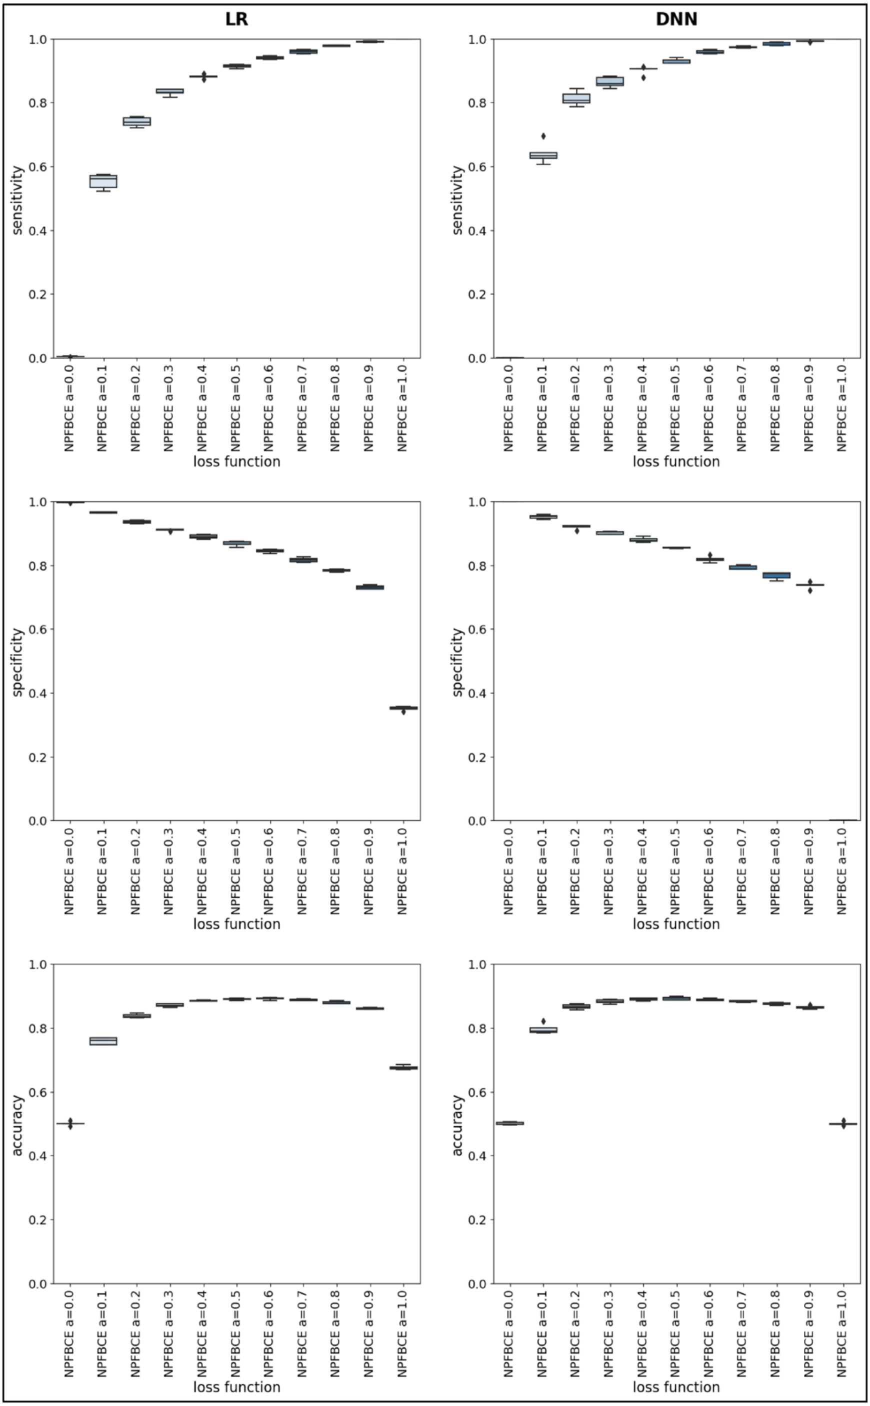

Fig. 4 Performance measures sensitivity, specificity, and accuracy of LR and DNN for *NPFBCE*

to our performance-flexible AI-based planning approach. Specifically, the sensitivity of *BCET* is determined after training, preventing influence on the model's learning process. This aspect is important not only for learning transferable patterns but also for applying techniques such as feature importance analysis during training. The performance of the different loss functions with $\alpha = 0.6$ and $\alpha = 0.8$ is shown in Fig. 5.

While a reduction in specificity can be acceptable or even desirable when prioritizing sensitivity, such as in ICU admission planning where missing true ICU needs could have critical consequences, lower specificity also implies an increased number of false positives. This may lead to overestimation of ICU demand, potentially resulting in inefficient allocation of scarce resources or unnecessary strain on ICU capacity. For healthcare practitioners, such false positives could mean unnecessary preparation or reservation of ICU beds, delaying other patients or impacting workflow. Therefore, in practice, it is crucial to carefully consider the clinical context and capacity constraints when choosing the weighting parameters in the loss function, balancing the trade-off between sensitivity and specificity to optimize both patient safety and resource utilization.

The results show that the *(N)PFBCE* loss function should be used in management decisions involving capacity constraints. Besides the focus on sensitivity, the function also outperforms the loss functions *WBCE* and *RWWCE*. Although *BCET* is comparable in performance, it has methodological weaknesses compared to *(N)PFBCE*. These findings can be transferred to the other weights with a focus on specificity. Additionally, with a high focus on sensitivity, *NPFBCE* performs better than *PFBCE* in our data set.

The performance measures of all loss functions for LR and DNN functions are shown in the Appendix. Due to memory constraints, only certain weights were considered for the two loss functions *WBCE* and *RWWCE*.

4.2 Evaluation of different capacity patterns using simulation

In the following, we discuss the outcomes of the simulation study across different data sets, capacity patterns, i.e., simulation scenarios defined in Table 5, and weights. As discussed before, we focus on one KPI which is the ratio of realized (anesthetist in D^1) and planned ICU patients (surgeon with ML support in D^2). A KPI of 1 (benchmark) indicates that all patients needing ICU treatment after elective surgery receive adequate care.

For the data set with all patients, normal ICU and OR capacity (scenario 5), a KPI of 1 is achieved for positive class weights α between 0.8 and 0.9. This indicates that a considerable focus on sensitivity is necessary to meet ICU demand under normal conditions. For state-of-the-art ML with a positive class weight of $\alpha = 0.5$, the KPI is 1.29, indicating an overbooked ICU (1.29 patients per unit of ICU capacity). This overbooking could lead to postponed or cancelled surgeries, or premature ICU discharges, compromising patient health and increasing managerial effort. Our approach, with α between 0.8 and 0.9, avoids these issues. These findings are consistent across scenarios with varying ICU and OR capacities. Performance-flexible AI-based planning of elective surgeries improves ICU capacity management, outperforming state-of-the-art ML predictions in all scenarios. Depending on OR and ICU capacity patterns, our approach with α between 0.7 and 1 is recommended. This approach is particularly important for scarce ICU capacities, while scarcity in the OR plays a subordinate role for performance-flexible ICU planning. For both scarce ICU and OR capacities, the gap between state-of-the-art planning and our

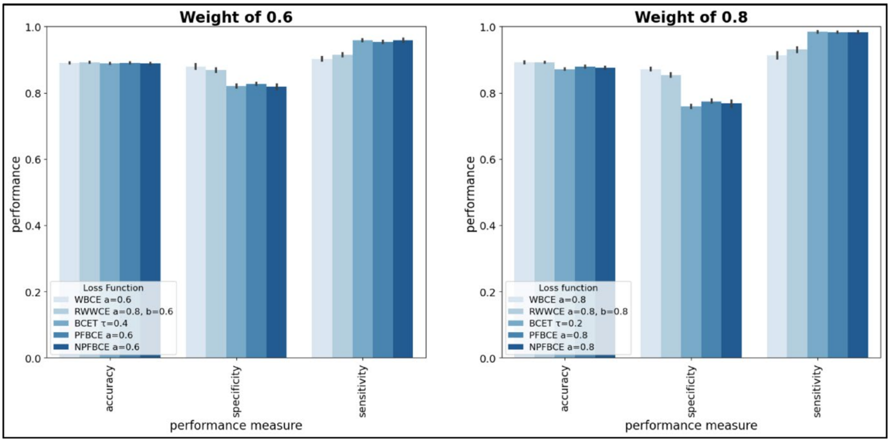

Fig. 5 Comparison of loss functions for $\alpha = 0.6$ and $\alpha = 0.8$ of DNN

performance-flexible AI-based planning approach becomes bigger but superior positive class weights of α remain unchanged compared to normal OR capacity.

Comparing these results for the data set with all patients with the data sets with high and low ASA-score patients highlights the importance of performance-flexible AI-based planning for patient cohorts with average (see previous paragraph) and high ASA-scores. For the data set with high ASA-score patients, positive class weights of α between 0.9 and 1 lead to a KPI near to 1. In contrast, using of state-of-the-art planning for this cohort, we expect up to 2.80 patients per unit of ICU capacity, reflecting significant overbooking. This scenario mimics peak COVID-19 conditions, where only severely ill patients (i.e., high ASA score) had access to elective surgeries. Under such circumstances, state-of-the-art planning exacerbates ICU overbooking, necessitating our performance-flexible AI-based planning approach (see Fig. 6).

For the data set with low ASA-score patients, using state-of-the-art planning, we can expect only 1.10 patients per unit of ICU capacity. For the data set with low ASA-score patients and scenario 6 (4), i.e., high (scarce) ICU and normal OR capacity, the KPI is 1 for positive class weights of α in between 0.6 and 0.7 (0.8 and 0.9). Thus, state-of-the-art planning might be applied in the unlikely situation of rather healthy patients. However, our performance-flexible approach can adapt to this scenario by allowing the user to flexibly adjust weights depending on the circumstances, supporting the application of our integrated approach for ex-ante ICU planning of elective surgeries. A summary of the results of the simulation study is shown in Fig. 7.

	Recommendation
All patients	Our approach
High ASA-score	Our approach
Low ASA-score	Standard ML model

Fig. 6 Model recommendations based on different patient cohorts

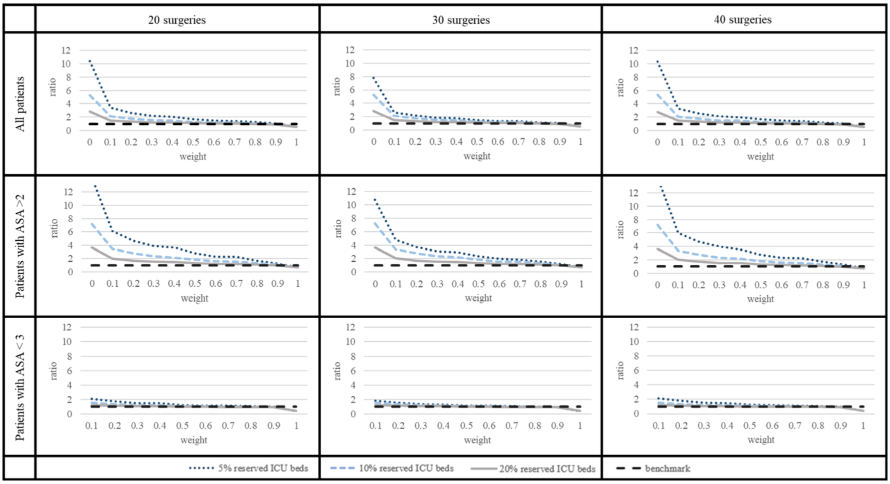

Fig. 7 Ratios of realized and planned ICU patients (KPI) for the different simulation scenarios (see Table 5) and data sets

5 Conclusion

This study introduces an innovative performance-flexible AI-based planning approach designed to predict ICU treatment following elective surgeries. Our approach emphasizes performance flexibility, achieved through a weighted loss function during training. By prioritizing a specific label while maintaining high accuracy, the model allows decision-makers, for example, to focus on sensitivity when ICU capacity is of crucial importance, leading to more accurate positive class predictions. The method considers both OR and ICU capacity for elective patients requiring post-surgery ICU care, resulting in stable ICU capacity ratios close to 1 across various simulation scenarios and patient cohorts with different resource availability.

In practice, this approach supports OR planners by enabling earlier and more reliable identification of patients likely to require ICU care, allowing adjustments to the surgical schedule before admission decisions are finalized. This can reduce last-minute cancellations and minimize discrepancies between initial surgical assessments and later anesthesia evaluations, particularly when ICU capacity is limited. Before implementation, local validation of the model with hospital-specific data, careful selection of appropriate weighting configurations aligned with capacity constraints, and integration into existing planning processes are essential to ensure seamless adoption and to align predictions with resource management goals.

There are some limitations to our methodology. First, data preparation steps applied to the data set can influence the model's results, and the impact of different preparation approaches remains unclear. Second, the loss functions for comparison include other objectives, leading to different requirements in implementation. Third, while the performance-flexible AI-based planning approach excels in the integrated ICU and OR case, its performance in other applications warrants further investigation.

Future research could extend our approach to address short-term scheduling. We focus on the initial assessment before the elective surgery (D^2), which occurs during the patient's visit to the surgical outpatient clinic. However, short-term adjustments may still be required. Additionally, applying the performance-flexible AI-based planning approach to other decision problems offers promising avenues for further exploration and application in diverse domains. Furthermore, the combination of the *(N) PFBCE* and thresholding could provide fruitful insights.

Appendix

Threshold

The difference between setting different thresholds is shown in Fig. 8. The figure visualizes 500 patients with two exemplary features, where the range of values can be neglected for understanding thresholds. The blue colors in the background show the different predicted probabilities. Darker colors indicate a higher predicted probability and vice versa.

The white dots represent samples of the positive class, and the black dots represent samples of the negative class. The black line indicates the threshold at which the sample is classified as positive (right side of the line, e.g., ICU) or negative (left side of the line, e.g., non-ICU) based on the predicted probabilities of the ML model. All samples on the right-hand side of the line are assigned to the positive class and all samples on the left-hand side are assigned to the negative class. By using the default state-of-the-art threshold of $\tau = 0.5$, the line is placed precisely at the midpoint. If a higher sensitivity is desired, the line is moved to the left. This results in more samples being assigned to the positive class.

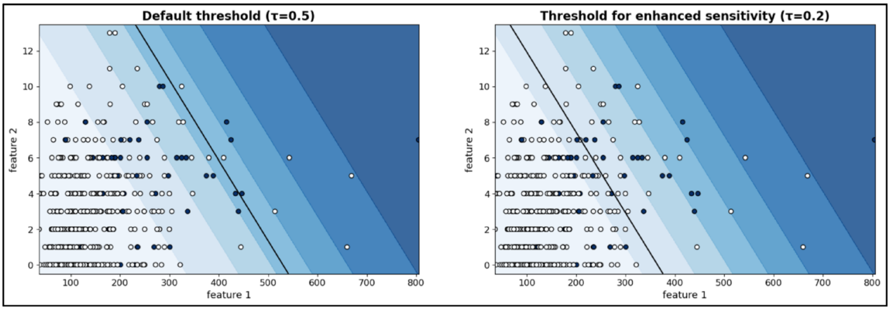

Fig. 8 Default threshold and increased sensitivity threshold for LR and two features with white dots representing samples of the positive class and black dots representing samples of the negative class. The left side of the line indicates the negative label, and the right side indicates the positive label

Results

See Figs. 9, 10, 11 and 12.

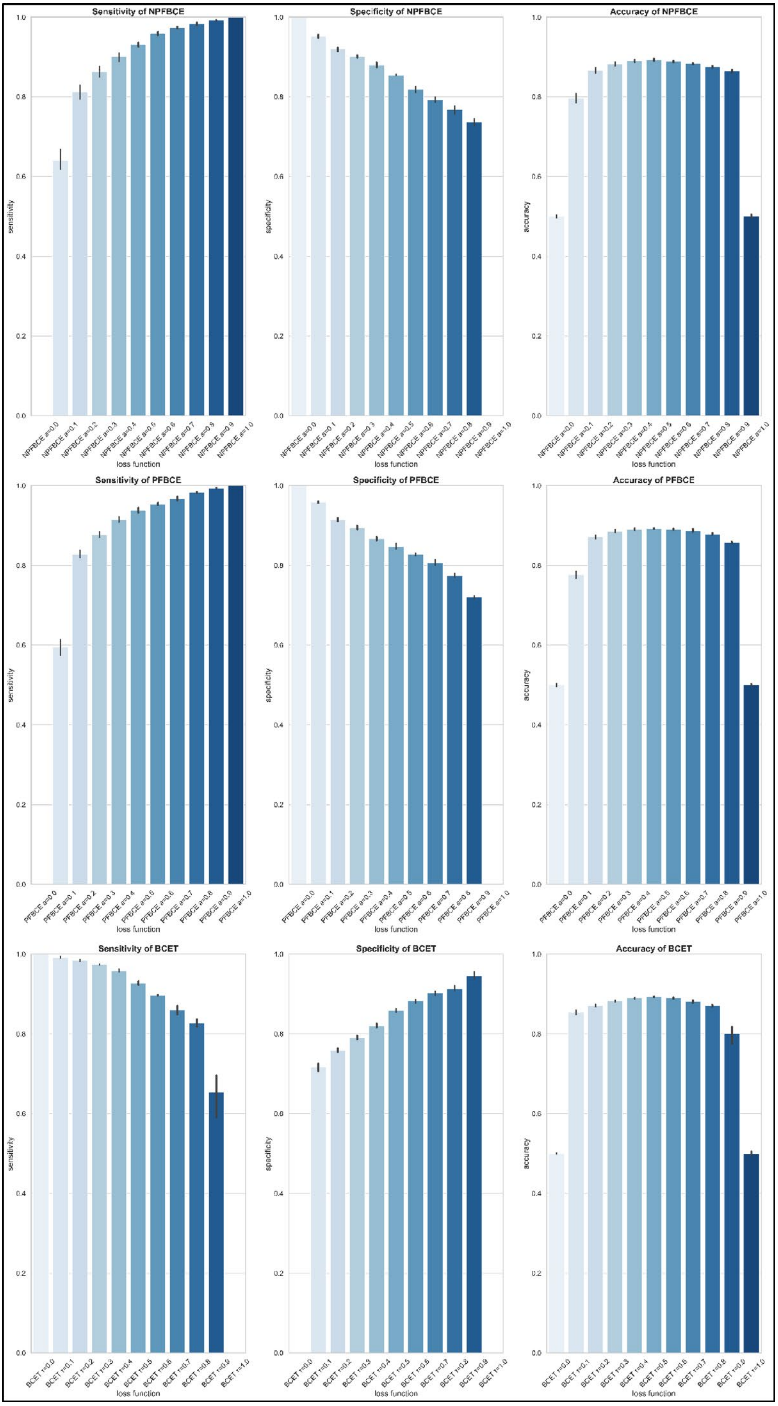

Fig. 9 Performance measures sensitivity, specificity, and accuracy of *NPFBCE*, *PFBCE* and *BCET* of DNN

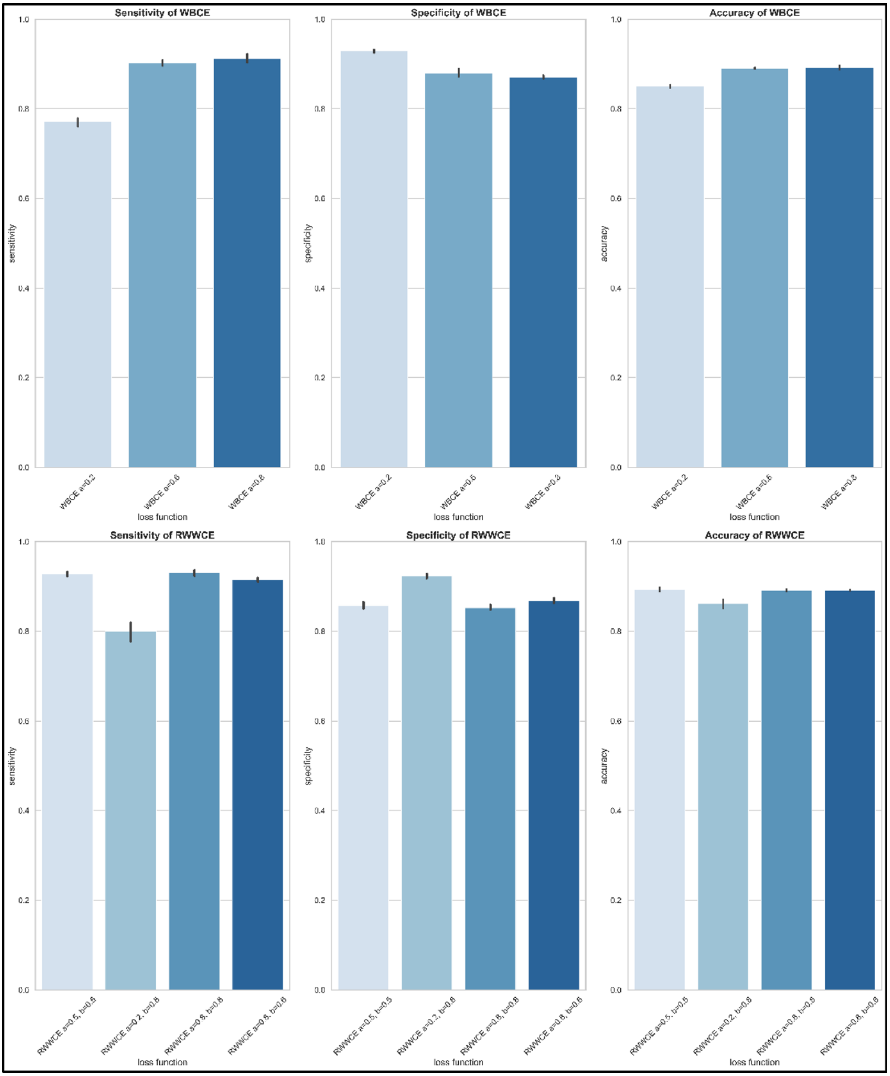

Fig. 10 Performance measures sensitivity, specificity, and accuracy of *WBCE* and RWWCE of DNN

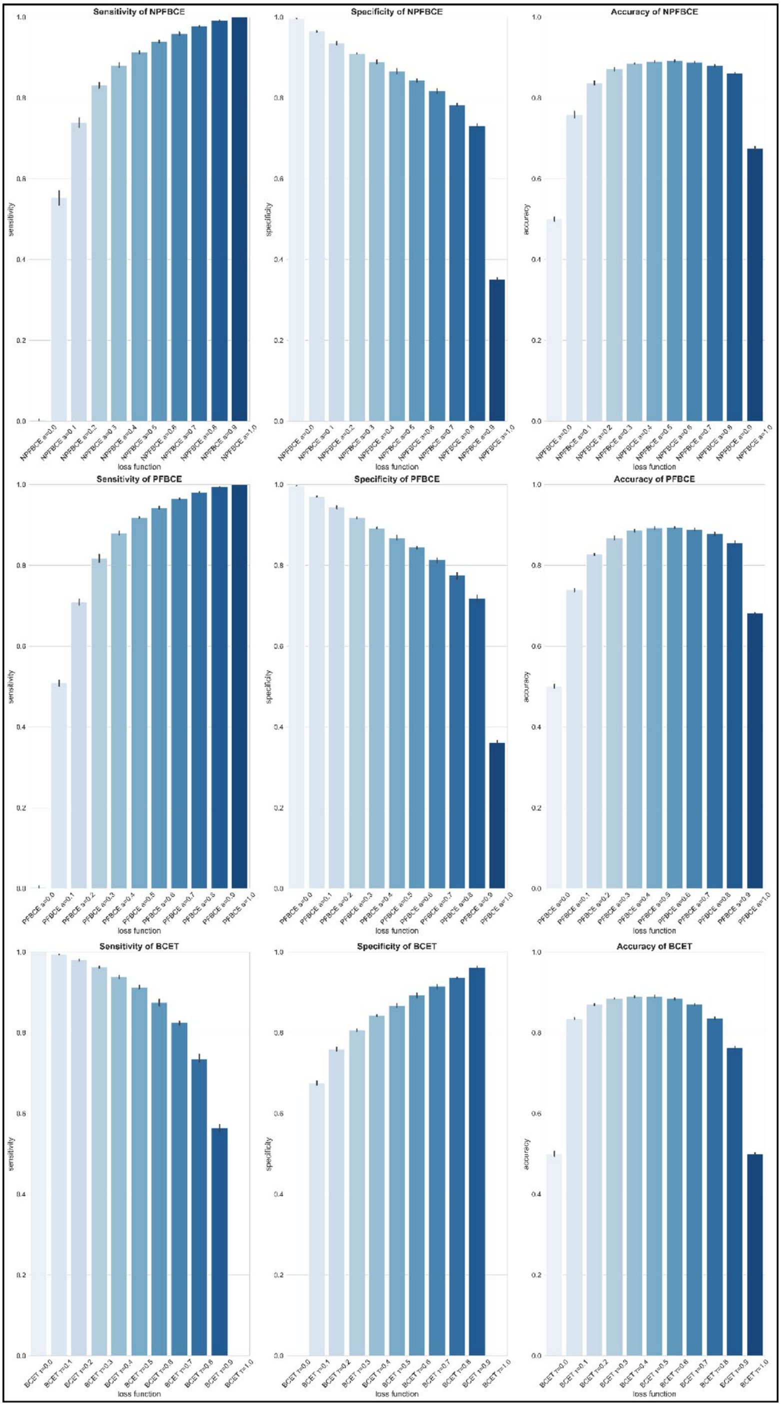

Fig. 11 Performance measures sensitivity, specificity, and accuracy of NPFBCE, PFBCE and BCET of LR

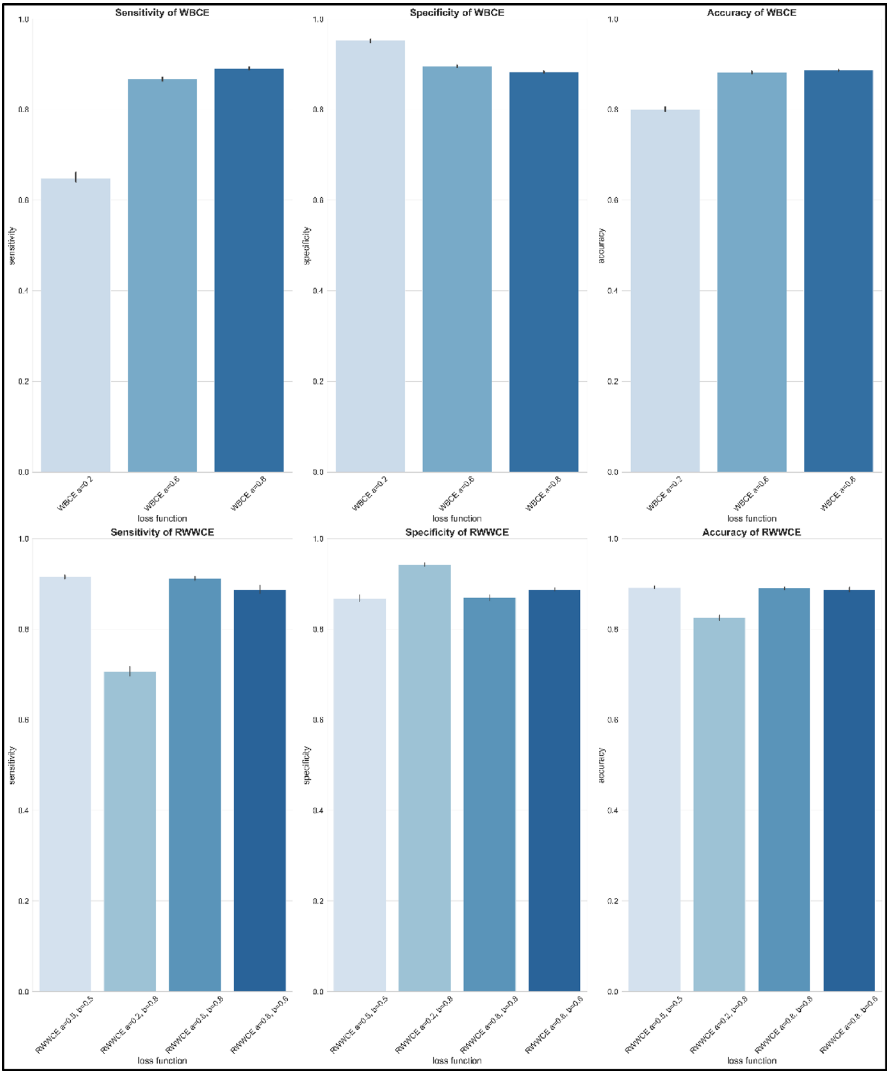

Fig. 12 Performance measures sensitivity, specificity, and accuracy of WBCE and RWWCE of LR

Acknowledgements We thank the German Federal Ministry of Education and Research for funding of the KISIK project. Grant Numbers 16SV9029 (University of Augsburg), 16SV9030 (University Hospital of Augsburg)

Funding Open access funding provided by Technical University of Denmark. This study was (partially) funded by the German Federal Ministry of Education and Research.

Data availability The data set generated during the current study is not publicly available as it contains exclusively sensitive healthcare data. Information on how to obtain it and reproduce the analysis is available from the corresponding author on request.

Declarations

Conflict of interest The authors Bartenschlager, Brunner, and Heller declare funding by the German Federal Ministry of Education and Research. The author Grieger declares no conflicts of interest.

References

Aurelio YS, de Almeida GM, de Castro CL, Braga AP (2019) Learning from Imbalanced data sets with weighted cross-entropy function. Neural Process Lett 50:1937–1949. https://doi.org/10.1007/s11063-018-09977-1

Bai J, Fügener A, Schoenfelder J, Brunner JO (2016) Operations management in intensive care unit scheduling: a literature review. SSRN Journal

Charlson ME, Pompei P, Ales KL, MacKenzie CR (1987) A new method of classifying prognostic comorbidity in longitudinal studies: development and validation. J Chronic Dis 40:373–383. https://doi.org/10.1016/0021-9681(87)90171-8

Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP (2002) SMOTE: synthetic minority over-sampling technique. JAIR 16:321–357. https://doi.org/10.1613/jair.953

Cichy RM, Kaiser D (2019) Deep neural networks as scientific models. Trends Cogn Sci 23:305–317. https://doi.org/10.1016/j.tics.2019.01.009

Cui Y, Jia M, Lin T-Y, Song Y, Belongie S (2019) Class-balanced loss based on effective number of samples. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 9268–9277. https://doi.org/10.1109/cvpr.2019.00949

Das S, Moore T, Wong W-K, Stumpf S, Oberst I, McIntosh K, Burnett M (2013) End-user feature labeling: supervised and semi-supervised approaches based on locally-weighted logistic regression. Artif Intell 204:56–74. https://doi.org/10.1016/j.artint.2013.08.003

Demeulemeester E, Beliën J, Cardoen B, Samudra M (2013) Operating room planning and scheduling. In: Denton BT (ed) Handbook of healthcare operations management: methods and applications, 1st edn. Springer, New York, pp 121–152. https://doi.org/10.1007/978-1-4614-5885-2_5

Drummond C, Holte RC (2003) Class imbalance, and cost sensitivity: why under-sampling beats over-sampling. Workshop on Learning from Imbalanced Datasets II

Elkan C (2001) The foundations of cost-sensitive learning. In: International joint conference on artificial intelligence, pp 973–978

Goic M, Bozanic-Leal MS, Badal M, Basso LJ (2021) COVID-19: Short-term forecast of ICU beds in times of crisis. PLoS ONE 16:e0245272. https://doi.org/10.1371/journal.pone.0245272

He H, Garcia EA (2009) Learning from imbalanced data. IEEE Trans Knowl Data Eng 21:1263–1284. https://doi.org/10.1109/tkde.2008.239

Heider S (2022) Leveling ICU bed occupancy using a quota system: an applied heuristic for operating theater managers. In: Trautmann N, Gnägi M (eds) Operations research proceedings 2021: Selected papers of the International Conference of the Swiss, German, and Austrian Operations Research societies (SVOR/ASRO, GOR e.V., ÖGOR), University of Bern, Switzerland, August 31-September 3, 2021. Springer, Cham, pp 216–222. https://doi.org/10.1007/978-3-031-08623-6_33

Heider S, Schoenfelder J, Koperna T, Brunner JO (2022) Balancing control and autonomy in master surgery scheduling: benefits of ICU quotas for recovery units. Health Care Manag Sci 25:311–332. https://doi.org/10.1007/s10729-021-09588-8

Heimerl C, Kolisch R (2010) Scheduling and staffing multiple projects with a multi-skilled workforce. Or Spectrum 32:343–368. https://doi.org/10.1007/s00291-009-0169-4
Ho Y, Wookey S (2020) The real-world-weight cross-entropy loss function: modeling the costs of mislabeling. IEEE Access 8:4806–4813. https://doi.org/10.1109/access.2019.2962617
Jadon S (2020) A survey of loss functions for semantic segmentation. In: 2020 IEEE conference on computational intelligence in bioinformatics and computational biology (CIBCB). IEEE. https://doi.org/10.1109/cibcb48159.2020.9277638
Jebali A, Diabat A (2015) A stochastic model for operating room planning under capacity constraints. Int J Prod Res 53:7252–7270. https://doi.org/10.1080/00207543.2015.1033500
Jebali A, Diabat A (2017) A Chance-constrained operating room planning with elective and emergency cases under downstream capacity constraints. Comput Ind Eng 114:329–344. https://doi.org/10.1016/j.cie.2017.07.015
Lin T-Y, Goyal P, Girshick R, He K, Dollar P (2017) Focal Loss for Dense Object Detection. In: 2017 IEEE international conference on computer vision (ICCV). IEEE. https://doi.org/10.1109/iccv.2017.324
Lipton ZC, Elkan C, Narayanaswamy B (2014) Thresholding Classifiers to Maximize F1 Score. arXiv preprint arXiv 1402
Lorenzen SS, Nielsen M, Jimenez-Solem E, Petersen TS, Perner A, Thorsen-Meyer H-C, Igel C, Sillesen M (2021) Using machine learning for predicting intensive care unit resource use during the COVID-19 pandemic in Denmark. Sci Rep 11:18959. https://doi.org/10.1038/s41598-021-98617-1
Ma X, Huang H, Wang Y, Romano S, Erfan S, Bailey J (2020) Normalized loss functions for deep learning with noisy labels. In: International conference on machine learning, pp 6543–6553
Maalouf M, Siddiqi M (2014) Weighted logistic regression for large-scale imbalanced and rare events data. Knowl-Based Syst 59:142–148. https://doi.org/10.1016/j.knosys.2014.01.012
Murphy KP (2022) Probabilistic machine learning: an introduction. Adaptive computation and machine learning series. The MIT Press, Cambridge Massachusetts
Mushava J, Murray M (2024) Flexible loss functions for binary classification in gradient-boosted decision trees: an application to credit scoring. Expert Syst Appl 238:121876
Neyshabouri S, Berg BP (2017) Two-stage robust optimization approach to elective surgery and downstream capacity planning. Eur J Oper Res 260:21–40. https://doi.org/10.1016/j.ejor.2016.11.043
Pearse RM, Moreno RP, Bauer P, Pelosi P, Metnitz P, Spies C, Vallet B, Vincent J-L, Hoeft A, Rhodes A (2012) Mortality after surgery in Europe: a 7 day cohort study. Lancet 380:1059–1065. https://doi.org/10.1016/S0140-6736(12)61148-9
Qian Z, Alaa AM, van der Schaar M (2021) CPAS: the UK's national machine learning-based hospital capacity planning system for COVID-19. Mach Learn 110:15–35. https://doi.org/10.1007/s10994-020-05921-4
Rachuba S, Imhoff L, Werners B (2022) Tactical blueprints for surgical weeks—an integrated approach for operating rooms and intensive care units. Eur J Oper Res 298:243–260. https://doi.org/10.1016/j.ejor.2021.06.005
Reig B, Heacock L, Geras KJ, Moy L (2020) Machine learning in breast MRI. J Magn Reson Imaging 52:998–1018. https://doi.org/10.1002/jmri.26852
Rezeai-Dastjerehei MR, Mijani A, Fatemizadeh E (eds) (2020) Adressing imbalance in mulit-label classification using weighted cross entropy loss function. IEEE. https://doi.org/10.1109/ICBME51989.2020
Rodríguez-Espíndola O (2023) Two-stage stochastic formulation for relief operations with multiple agencies in simultaneous disasters. Or Spectr 45:477–523. https://doi.org/10.1007/s00291-023-00705-3
Saklad M (1941) Grading of patients for surgical procedurs. Anesthesiology 2:281–284. https://doi.org/10.1097/00000542-194105000-00004
Schiele J, Koperna T, Brunner JO (2021) Predicting intensive care unit bed occupancy for integrated operating room scheduling via neural networks. Nav Res Logist 68:65–88. https://doi.org/10.1002/nav.21929
Sheng J, Wu S, Zhang Q, Li Z, Huang H (2022) A binary classification study of Alzheimer's disease based on a novel subclass weighted logistic regression method. IEEE Access 10:68846–68856. https://doi.org/10.1109/access.2022.3186888
Sheng VS, Ling CX (2006) Thresholding for making classifiers cost-sensitive. AAAI 6:476–481. https://doi.org/10.1016/B978-0-12-802021-0.00010-3
Shrivastava A (2020) Adma: a flexible loss function for neural networks. arXiv preprint arXiv:2007.12499
van Houdenhoven M, van Oostrum JM, Wullink G, Hans E, Hurink JL, Bakker J, Kazemier G (2008) Fewer intensive care unit refusals and a higher capacity utilization by using a cyclic surgical case schedule. J Crit Care 23:222–226. https://doi.org/10.1016/j.jcrc.2007.07.002

van Oostrum JM, van Houdenhoven M, Hurink JL, Hans EW, Wullink G, Kazemier G (2008) A master surgical scheduling approach for cyclic scheduling in operating room departments. Or Spectrum 30:355–374. https://doi.org/10.1007/s00291-006-0068-x

Vargas LG, May JH, Spangler W, Stanciu A, Strum DP (2009) Operating room scheduling and capacity planning. In: Stonemetz J, Ruskin K (eds) Anesthesia informatics. Springer, New York, pp 361–392. https://doi.org/10.1007/978-0-387-76418-4_19

Wachtel RE, Dexter F (2008) Tactical increases in operating room block time for capacity planning should not be based on utilization. Anesth Analg 106:215–226. https://doi.org/10.1213/01.ane.0000289641.92927.b9

Zare N, Haem E, Lankarani KB, Heydari ST, Barooti E (2013) Breast cancer risk factors in a defined population: weighted logistic regression approach for rare events. J Breast Cancer 16:214–219. https://doi.org/10.4048/jbc.2013.16.2.214

Zhou Z-H, Liu X-Y (2006) Training cost-sensitive neural networks with methods addressing the class imbalance problem. IEEE Trans Knowl Data Eng 18:63–77. https://doi.org/10.1109/TKDE.2006.17

Authors and Affiliations

Milena Grieger[1] · Jens O. Brunner[1,2,5] · Axel R. Heller[3] · Christina C. Bartenschlager[3,4]

✉ Jens O. Brunner
jotbr@dtu.dk

1 Health Care Operations/Health Information Management, Faculty of Business and Economics, Faculty of Medicine, University of Augsburg, Augsburg, Germany

2 Department of Technology, Management, and Economics, Technical University of Denmark, Kongens Lyngby, Denmark

3 Anaesthesiology and Operative Intensive Care, University Hospital of Augsburg, Augsburg, Germany

4 Applied Data Science in Health Care, Nürnberg School of Health, Ohm University of Applied Sciences Nuremberg, Nuremberg, Germany

5 Center for Excellence in Healthcare Operations Planning, Next Generation Technology, Slagelse, Region Zealand, Denmark