scientific reports

OPEN

An intelligent community-based system for healthcare prioritisation

Micaela Pinho[1,2,3] & Fátima Leal[1]

Healthcare rationing is unavoidable in systems constrained by limited resources. While decisions about who should be treated are ethically complex, they must reflect not only efficiency concerns but also socially accepted values. This study aims to develop a multi-criteria decision-support system - Vital Priority System, that prioritise patients using a Random Forest algorithm trained on multiple rationing criteria endorsed by Portuguese civil society. Based on a Portuguese online survey data, the model incorporates nine dimensions: clinical need, life expectancy gain, quality of life improvement, age, waiting time, parental status, lifestyle responsibility, and social role. Our results show that clinical need, expected treatment effectiveness, waiting time and age were the most influential, followed by parental status. Lifestyle and social role factors were least weighted. The proposed system enables the classification of patients as 'priority' or 'non-priority', providing healthcare professionals with a transparent, consistent, and ethically grounded tool to support decision-making. This study advances the literature by operationalising, for the first time in the Portuguese context, public preferences in a replicable AI-based framework for fairer patient prioritisation.

Keywords Patient prioritisation, Explicit rationing, Artificial intelligence, Machine learning, Multi-criteria decision-support system, Vital priority system

Population ageing, the increasing prevalence of chronic, infectious, and communicable diseases, lifestyle-related health risks, rapid technological development, and rising societal expectations regarding the potential of medicine to enhance health and well-being are all contributing to a growing demand for healthcare services. This rising demand exerts significant pressure on the limited resources of public healthcare budgets. Rationing or prioritising healthcare (terms used interchangeably throughout this paper) is inherently complex and contentious, as it involves balancing competing objectives. These decisions remain one of the most pressing challenges in healthcare policy.

While demand continues to outpace available resources, policymakers are compelled to make difficult decisions about how to allocate public funds, namely, which services and treatments to support. Consequently, healthcare professionals are often required to decide which patients to treat. Rationing or prioritising healthcare (terms we use interchangeably throughout this study) is inherently complex and contentious, as it involves balancing competing objectives.

In many countries, including Portugal, such decisions are frequently made at the discretion of politicians and health professionals. Strategies such as reducing hospital beds, shortening hospital stays, closing health services, or allowing long waiting lists are common examples of implicit healthcare rationing.

Implicit rationing has been perpetuated by political actors who, seeking to avoid the fallout of unpopular decisions, delegate responsibility to health professionals[1]. From a political economy perspective, politicians tend to avoid blame for controversial actions, concerned about the electoral consequences. In a representative democracy, political parties often converge toward the preferences of the median voter to secure electoral success[2], indicating that public decision-makers are not neutral social welfare maximisers but instead driven by self-interest[3]. At the micro or clinical level, however, health professionals, under persistent pressure to make ethically challenging decisions, experience significant psychological strain[4] and increasingly seek to share this burden with political authorities. Indeed, global concern about the psychological well-being of healthcare workers has increased in the wake of the COVID-19 pandemic[5].

These considerations highlight the urgent need for a systematic and transparent approach to priority setting that can assist policymakers and health professionals alike. In recent decades, economic evaluation frameworks, particularly those based on effectiveness and cost, have gained prominence. Cost-effectiveness analysis has emerged as a key tool for identifying priority interventions[6], comparing the cost per health outcome of different

[1]Research on Economics, Management and Information Technologies, REMIT, Portucalense University, Porto, Portugal. [2]Portuguese Law Institute, IJP, Portucalense University, Porto, Portugal. [3]Research Unit in Governance, Competitiveness and Public Policy, GOVCOPP. Aveiro University, Aveiro, Portugal. ✉email: michaelapinho@hotmail.com

treatments and ranking them accordingly. Within this framework, cost-utility analysis (CUA) has become widespread, using quality-adjusted life years (QALYs) to simultaneously capture gains in life expectancy and quality of life. The underlying logic of CUA is one of efficiency: maximising population health. In this view, efficient allocation of scarce healthcare resources means prioritising patients with the greatest likelihood of treatment success. Despite its economic rationale, this approach has been contested by civil society, which is concerned not only with aggregate health gains but also with their distribution. Economic models often fail to reflect public perceptions of distributive justice, omitting important criteria beyond cost and effectiveness. For example, during the recent public health pandemic, the principle of health gain maximisation guided clinical decision-making in some countries. In Italy and Spain, overwhelmed health systems forced healthcare professionals to allocate scarce ventilators to patients with the highest likelihood of survival[7–9]. This practice was later challenged in Italy, where relatives or deceased patients filed lawsuits against the health system and its decision-makers[10]. Similar protests occurred in the United States, namely in Washington[11] and New York City[12].

Rationing decisions at the national level raise similar concerns. Countries like Portugal apply cost-effectiveness thresholds in their health technology assessments to determine whether new medicines should receive public reimbursement. Treatments for rare diseases often exceed these thresholds and are thus excluded from public coverage. This was the case in Portugal for two highly expensive drugs for hepatitis C and spinal muscular atrophy. Although initially excluded, public pressure, amplified by media campaigns, ultimately compelled the government to reverse its decision and fund the hepatitis C treatment (for a review see[13].

Whether to prevent the media influence on public policy, to relieve the emotional burden on healthcare professionals, or to uphold democratic legitimacy by involving taxpayers in healthcare decisions there is a clear need to define rigorous, transparent, and socially accepted criteria for prioritisation. As the main funders and beneficiaries of public health systems, citizens' values and preferences must be integrated into the decision-making process. In recent years, ethical, economic, political, and legal arguments have increasingly supported the incorporation of public preferences in healthcare rationing decisions[14–17]. While this participatory process may be lengthy, shared decision-making among policymakers, clinicians, and the public can enhance accountability and increase societal acceptance of difficult choices.

A major challenge in implementing explicit rationing mechanisms is defining the criteria for inclusion. Numerous prioritisation criteria have been proposed in the literature[18–24] along with evidence of the public's support for various criteria (e.g[25,26]. for reviews). Civil societies tend to prioritise interventions targeting the worst-off namely those in greater clinical need (e.g., severely ill or in pain), the youngest (to allow a normal life span), and individuals with pressing social needs (e.g., parents). Additionally, higher priority is often granted to individuals with socially important roles (e.g., physicians), while lower priority is attributed to those whose health issues stem from irresponsible behaviour.

The inherent complexity of defining an explicit rationing process stems from the need to balance multiple, and often conflicting, criteria. This complexity reinforces the need for systematic, transparent, and structured decision-making processes. To this end, multi-criteria decision models have been proposed as promising tools for guiding healthcare resource allocation (for a review, see[27]. These models can incorporate multiple dimensions although useful these models remain underutilised in practical settings[28].

The present study introduces a novel multi-criteria decision model that employs Machine Learning (ML) to classify patients as priority or non-priority cases, based on the preferences of Portuguese citizens. This model aims to provide practical evidence to inform resource allocation decision in Portugal. By embedding societal values into patient prioritisation, it offers a more democratic and inclusive alternative to traditional, top-down approaches. ML, a subset of artificial intelligence (AI), generates predictions or decisions by learning from data and improving through experience. While ML applications in healthcare, such as diagnostic classification (e.g., osteosarcoma)[29], treatment selection, emergency triage[30–32], and personalised medicine[33] are increasingly common its application to patient prioritisation based on distributive justice remains unexplored.

To the authors' knowledge, this is the first study to apply ML in the development of a transparent patient prioritising system grounded in widely debated ethical criteria and public preferences. Given that ML systems learn from existing datasets, this study uses a large sample of Portuguese citizens' views on prioritisation to develop an intelligent classification system. Because the proposed AI system is built on public preferences, it addresses key concerns related to distributive justice and transparency, prerequisites for a responsible and ethical AI model[34].

The proposed method applies a decision tree-based algorithm - Random Forest., which is widely used for both prediction[35] and classification tasks[36]. In contrast to informal judgments that often dominate patient prioritisation, this method offers a data-driven, systematic approach that enhances rationality in decision-making. ML tools, with their ability to process real-time data, hold great promise for supporting future decisions on the allocation of scarce healthcare resources.

Methodology

Proposed method

This study aims to develop a robust ML model to classify patient priority based on public preferences regarding nine rationing criteria identified in the literature and detailed below. This approach offers actionable insights to support healthcare professionals, promote socially equitable patient triage systems, and enhance the overall efficiency and fairness of healthcare delivery.

We propose a multi-criteria decision-support, referred to as VITAL Priority System designed to classify patient as priority or non-priority by integrating societal input through ML techniques. The system leverages ML algorithms to analyse the collected data, identify the most relevant predictors, and detect patterns in patient prioritisation preferences.

The development of the VITAL Priority System follows four main steps, as illustrated in Fig. 1. Each of these steps is described below.

Data processing and feature selection

No additional preprocessing was necessary, as the dataset was already clean, complete, and appropriately structured. All variables were in a format compatible with the classification model, which is inherently robust to variable scaling and capable of handling both categorical and numerical inputs without the need for transformation. The dataset contained no missing values, and neither normalization nor encoding procedures were required.

Feature selection was conducted using the SelectKBest method with the ANOVA F-value as the scoring function. This approach identified and ranked the most influential features, i.e., the rationing criteria, based on their contribution to the model's decision-making process. The relative importance of each criterion was expressed as a percentage of the total feature importance, providing insight into the extent to which the model relied on each factor for classification.

Classification

Patient prioritisation was carried out exclusively using ensemble learning methods, given their superior performance and robustness in complex classification tasks, particularly in high-stakes domains such as healthcare. Ensemble models, by aggregating the predictive capabilities of multiple base learners, are known to reduce variance (e.g., Random Forest), bias (e.g., boosting methods), and the risk of overfitting compared to individual classifiers. Considering the heterogeneous nature of patient data and the critical demand for reliable decision-making in emergency prioritisation scenarios, ensemble techniques offer a more resilient and accurate foundation for prediction. Accordingly, we benchmarked the Random Forest model against two others widely adopted ensemble methods: (i) AdaBoost and (ii) Gradient Boosting, to ensure a comprehensive and balanced evaluation of model performance.

Random forest is a non-parametric ensemble learning method that build multiple decision trees and aggregates their outputs to enhance predictive performance. During the training phase, several decision trees are constructed; the final classification is determined by majority voting across these trees[37]. This ensemble strategy reduces overfitting and improves the model's generalizability. In this study, Random Forest was applied to classify patients as priority or non-priority in the context of health-related emergencies, incorporating a comprehensive set of rationing criteria as predictive features.

AdaBoost is a boosting-based ensemble method that combines multiple weak learners, typically shallow decision trees, into a single strong classifier. The algorithm iteratively adjusts the weights of training instances, placing greater emphasis on those that were misclassified by previous models. The final prediction is derived through a weighted majority vote across all learners. AdaBoost is recognised for its simplicity and effectiveness, particularly in binary classification tasks[38].

Gradient Boosting is a more advanced ensemble technique that constructs decision trees sequentially, with each new tree trained to correct the prediction errors of the preceding ensemble. Unlike AdaBoost, which reweights misclassified instances, Gradient Boosting optimises a differentiable loss function by fitting each new tree to the residuals of the existing model. While it offers greater flexibility and often yields higher predictive accuracy than AdaBoost, it typically requires more careful hyperparameter tuning to prevent overfitting[39].

Evaluation

Model performance was assessed using standard classification metrics, including accuracy and F-measure. Classification accuracy measures the overall proportion of correct predictions (both true positives). The F-measure, calculated under both macro and micro-averaging strategies, combines Precision (the proportion of correct positive predictions) and Recall (the proportion of actual positives correctly identified). Macro-averaging treats all classes equally, while micro-averaging accounts for class imbalance by weighting classes according to their frequency. Together, these metrics provide a comprehensive evaluation of the model's predictive performance.

To ensure reliable performance estimates and to mitigate the risk of overfitting, we applied stratified 5-fold cross-validation. This technique preserves the proportion of priority and non-priority cases within each fold, thus supporting a balanced and representative evaluation across all iterations. Hyperparameter tuning was conducted via grid search, systematically exploring the following parameter combinations:

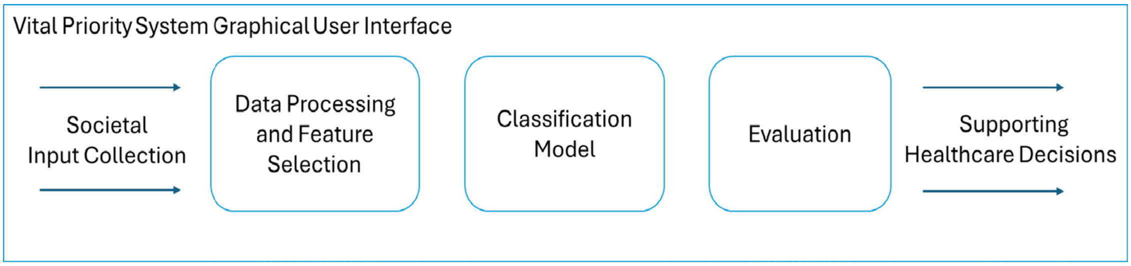

Fig. 1. Proposed method.

Number of estimators (n_estimators) ∈ {50, 100, 150}.

Maximum tree depth (max_depth) ∈ {10, 15, 20}.

Minimum number of samples required to split an internal node (min_samples_split) ∈ {2, 5}.

The final model configuration, n_estimators = 100, max_depth = 15, min_samples_split = 2, was selected based on the best average F1-score obtained across the validation folds.

Graphical user interface (GUI)

A detailed graphical user interface (GUI) was developed to enable healthcare professionals to input patient data, specifically, scores from 1 to 5 for each of the nine rationing criteria and receive real-time priority classifications. The scale reflects the relevance of each criterion: higher values correspond to greater pain intensity, higher expected gains in life expectancy or quality of life, longer waiting times, etc. The system processes these inputs and performs a binary classification, assigning each patient to one of two categories: priority (high urgency) or non-priority (lower urgency). Alongside the predicted classification, the GUI provides explanatory insights regarding the relative importance of the input features and, where applicable, additional patient-specific recommendations. This tool facilitates transparent, consistent, and socially informed triage decisions in clinical settings.

Inputs of the system: dataset and patient's prioritisation criteria

Data were collected from 2125 respondents in Portugal through an online questionnaire administered between 2020 and 2023. After excluding incomplete or unusable responses, the final sample comprises 1984 participants. Eligibility criteria required participants to be at least 18 years old and residing in Portugal. A heterogeneous sample was sought to maximise variability, using a mixed recruitment strategy. Approximately 44% of the sample consisted of university students from diverse academic backgrounds and institutions, while the remainder included adults recruited from public spaces. The final gender distribution was 54.4% female and 45.6% male, with a mean participant age of 38 years.

The questionnaire presents participants with 18 hypothetical prioritisation scenarios, each involving a pairwise choice between two patients - Patient A and Patient B, differentiated by personal or health characteristics. Participants were asked to act as decision-makers and indicate their level of preference on a 5-point semantic differential scale, with 1 indicating "Definitely Prioritise Patient A,"; 2 denotes "Some Preference for Patient A,"; 3 represents "No Preference," 4 indicates "Some Preference for Patient B,"; and 5 means "Definitely Prioritize Patient B,". Details of these hypothetical choices are described in Table A1 in the appendix as supplemantary material 1.

Prior to completing the questionnaire, participants were presented with an introductory information sheet explaining the study's purpose, guaranteeing anonymity, and requesting informed consent. Written informed consent was obtained from all participants. All procedures were conducted in accordance with ethical standards established by relevant institutional and national bodies (DL. 80/2018 de 15 outubro; Regulamento (UE) 2016/679 do Parlamento Europeu e do Conselho de 27 de abril de 2016; Regulamento da CESUPT de 2 de junho de 2020) as well as the Declaration of Helsinki and its subsequent amendments or comparable ethical standards. Ethical approval was obtained from the Ethics Committee of Portucalense University (CEI25_04_007).

The questionnaire design was developed and used elsewhere[17,24,40]. Although the questionnaire is very similar to the one used by[17], in the present study, we used 18 hypothetical rationing questions (instead of 23), and civil society (lay persons) was the target instead of health professionals. Similarly to[17], here we subdivide the six main rationing criteria (described in detail in[17,24] into nine: (1) Clinical need measured in severity that can be assessed in (i) pain and (ii) immediate risk of dead - Rule of Rescue (2) Health maximization divided into (iii) extending life expectancy and (iv) increasing quality of life; (3) Social need measured into (v) parenthood; (4) Age-discrimination – evaluated by a preference for younger over elderly – ageism (vi), (5) Merit evaluated (vi) positively - instrumental value, and (vii) negatively as 'punishment' for engaging in risky lifestyles, and (6) Fair-chance evaluated by (viii) waiting time. Each rationing criterion was represented by two hypothetical decision scenarios, yielding a total of 18 scenarios reflecting societal preferences, as listed in the final column of Table A1 in the appendix as supplemantary material 1.

Experiments and results

The experiments conducted provided a comprehensive evaluation of the performance and effectiveness of the proposed Vital Priority System. This section details the experimental setup, including the prioritisation criteria selected, and the ML techniques applied. We present the results from various tests, highlighting the model's accuracy, precision, and robustness in classifying patients as either priority or non-priority.

These findings are discussed in terms of their practical implications, offering insights into the system's potential contribution to healthcare decision-making. The primary objective of the experiments was to validate the model's capabilities and demonstrate its applicability in real-world healthcare context.

All offline experiments were conducted using the collected dataset. The system operated on a Windows 64-bit platform with an Intel(R) Core(TM) i7-8565U CPU @ 1.99 GHz processor, 16 GB of RAM, and a 500 GB SSD drive.

The experiments comprise five steps: (i) Pre-processing for classification, (ii) feature selection and importance ranking, (iii) classification, (iv) GUI development, and (v) analysis of a case study. All classification models were

Criteria	Weight (%)
Clinical Need: Pain	17.25
Life expectancy extending	14.73
Age discrimination - ageism	11.42
Clinical Need: Rule of Rescue	11.07
Quality of life increase	10.39
Fair chance: Waiting time	9.93
Social Need: Parenthood	9.86
Merit: Negative	8.53
Merit: Instrumental Value	6.83

Table 1. Feature importance of the nine prioritisation criteria.

implemented using the scikit-learn package in Python. All classification models were implemented using the scikit-learn package in Python.

Pre-processing for classification

In this initial stage, respondents' decisions were used to assign labels to each patient. First, we compute the average decision score per patient based on inputs from multiple respondents. consolidates the individual opinions into a single score for each patient. A classification threshold was then applied: if the average score was below 3, the patient was categorised as non-priority; if the average score is above 3, the patient was classified as priority. This approach allowed the transformation of continuous scores into binary classification labels. Additionally, any entries with missing data were excluded to ensure the reliability and completeness of the dataset. This pre-processing step ensured that the data used for classification was clean, consistent, and suitable for modelling.

Feature selection and importance

In the second stage, we identified the most relevant criteria that significantly influenced the model's predictive performance In the second stage, we identified the most relevant criteria that significantly influenced the model's predictive performance. Descriptive statistics for the eighteen prioritisation decisions are presented in Table A2 (Supplemantary material 1). To assess the importance of the various criteria, we applied the SelectKBest method using the ANOVA F-value score function. Table 1 summarises the resulting weights assigned to each criterion, reflecting their relative influence in the prioritisation process.

Among all features, clinical needs translated into pain and suffering presents the highest weight (17.25%) in the final prioritisation of the patient. Efficiency concerns, measured through mortality and morbidity, ranked second (14.7%) and fifth (10.4%), respectively. Discrimination by age emerged as the third most important factor (11.4%), closely followed by the immediate risk of death (11.1%). Further relevant features included the first-come, first-served criterion (9.93%) and social need (9.86), both nearly equal in influence as the improvement in quality of life (10.39) to improving quality of life. Merit-based considerations ranked sixth and seventh in importance. Notably, the negative form of deservingness, which penalises risky behaviours, had a higher influence (8.5%) than instrumental value (6.8%).

This contributes to the system's capacity to provide actionable and equitable guidance.

The careful selection and weighting of these features ensure that the model adopts a balanced and ethically grounded approach to patient prioritisation. Each feature was included for its demonstrated relevance in healthcare decision-making, thereby enhancing the model's alignment with societal values. This contributes to the system's capacity to provide actionable and equitable guidance and thus support healthcare professionals in making informed and equitable decisions.

Classification ML results

Despite a slightly lower F-measure for class 0 compared to class 1, the classifiers exhibited excellent performance in identifying both high- and low-priority patients. These strong results underscore the potential of the models to support fair and consistent prioritisation decisions in clinical settings. The classification algorithms were trained to predict whether a patient should be prioritised, based on the selected set of features. Class 0 corresponds to non-priority patients, while class 1 represents patients deemed to require prioritisation.

As presented in Table 2, the Random Forest classifier consistently outperformed the other ensemble methods across all evaluation metrics for both classes. Notably, it achieved the highest precision and recall for class 1 (the priority group), which is especially critical in emergency healthcare contexts where minimising false negatives is paramount. Its strong performance for class 0 further indicates a reliable capacity to correctly identify non-priority cases, thus ensuring balanced and trustworthy decision-making.

The Random Forest model achieved an overall accuracy of 99%. Its macro-averaged F-measure of 0.96 reflects balanced performance across both classes. While the F-measure for class 0 was slightly lower (0.93) than for class 1 (0.99), the classifier nonetheless demonstrated excellent discriminatory power. These results affirm

Model	Accuracy	Precision		Recall		F1-measure	
		Class#0	Class#1	Class#0	Class#1	Class#0	Class#1
AdaBoost	0.96	0.91	0.95	0.90	0.96	0.90	0.95
Gradient Boosting	0.97	0.93	0.97	0.92	0.98	0.92	0.97
Random Forest	**0.99**	**0.95**	**0.99**	**0.93**	**0.99**	**0.93**	**0.99**

Table 2. Comparison of ML classifiers.

the model's utility as a decision-support tool for healthcare professionals engaged in time-sensitive and ethically complex prioritisation tasks.

To evaluate the stability of model performance, 95% confidence intervals were calculated across the five cross-validation folds. For the Random Forest classifier, the F-measure for class 0 (non-priority) was 0.93 ± 0.02, and for class 1 (priority), it was 0.99 ± 0.01. These narrow confidence intervals indicate a high degree of consistency across validation folds. Additionally, paired t-tests comparing Random Forest with AdaBoost and Gradient Boosting revealed statistically significant improvements in macro F-measure ($p = 0.03$ and $p = 0.04$, respectively), confirming that the superior performance of Random Forest is unlikely to be due to random variation.

From a practical standpoint, the model's high recall for priority patients (0.99) ensures that critical cases are rarely missed, an essential requirement in emergency healthcare settings. Concurrently, the strong precision for non-priority patients minimises the risk of unnecessary escalation of low-risk cases, thereby supporting more efficient allocation of healthcare resources.

Vital priority system GUI

Figure 2 presents the GUI designed to facilitate the prioritisation process based on multiple criteria. On the left side of the interface, users can input patient-specific values (ranging from 1 to 5) for nine different criteria: Pain relief, life expectancy, age, rule of rescue, quality of life, waiting time, parenthood, punishment for risky behaviours, and instrumental value. These criteria represent key factors in assessing the urgency and justification for medical intervention. The right side of the interface displays the result of the classification process, indicating whether the patient is considered a priority. In the current configuration, it shows "No Priority Patient." The interface is designed to be clean, user-friendly, and intuitive, allowing healthcare professionals to make rapid and accurate prioritisation assessments that reflect societal considerations.

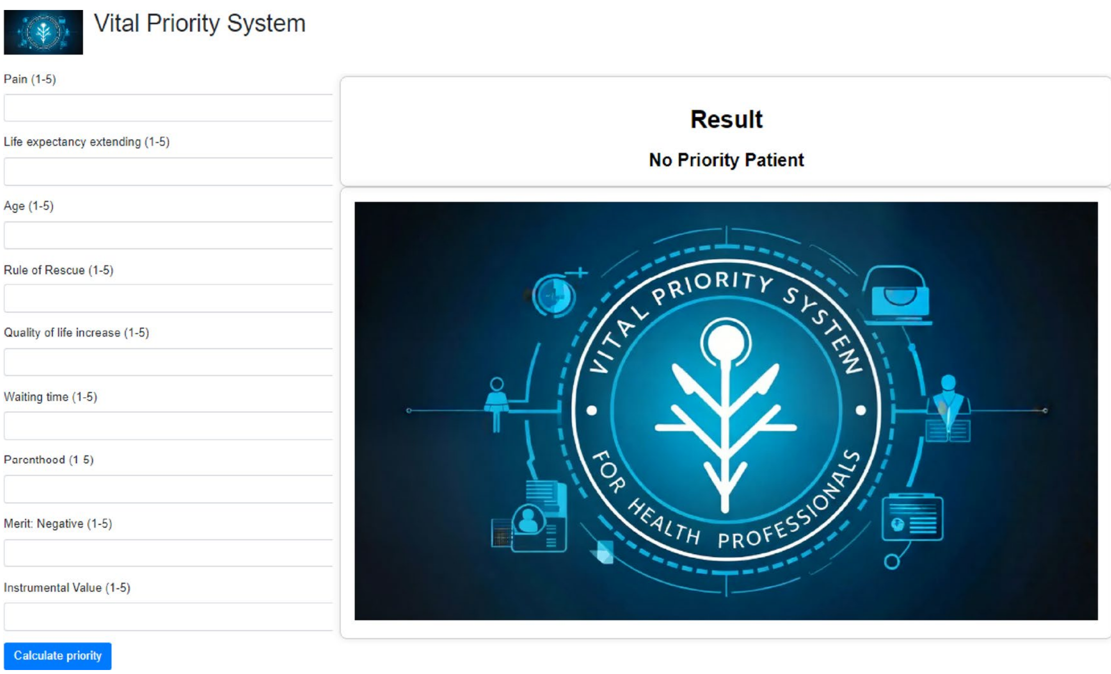

Fig. 2. Vital Priority System GUI.

Case study

Due to the scarcity of beds and limited human and financial resources, health professionals often face difficult decisions regarding patient prioritisation. Using the Vital Priority System, a healthcare professional inputs data for three patients with distinct health conditions and social backgrounds as follows:

Patient A is a middle-aged individual experiencing severe pain (rated 5/5). The patient has a high potential for life expectancy extension (rated 5/5), with treatment potentially extending life by up to 60%, or 10 years. Although patient A's condition is not immediately life-threatening, it requires attention (rating 2/5). The expected improvement in quality of life is substantial (rated 5/5). The patient has been waiting six months for treatment (waiting time score: 4/5). As a single parent with dependents, their health is crucial for their family's well-being (parental status: 5/5). His/her illness is not due to negligent behaviour (punishment score: 5/5). The patient works for a non-governmental organisation rescuing migrants at sea (instrumental value score: 4/5). Middle age contributes a score of 4/5.

Patient B is an elderly individual with moderate pain (rated 3/5). Treatment is expected to increase life expectancy by approximately three years (rated 2/5). This patient is in an immediate life-threatening situation (rule of rescue: 5/5) yet the anticipated quality of life improvement is minimal (ranked 2/5). He/she has been waiting for treatment for three months (waiting time score:3/5). Without dependents (parental status: 1/5), no risky lifestyle behaviours (punishment: 5/5), and having led an ordinary life (instrumental value: 1/5), ageism is a concern here (age discrimination score: 1/5).

Patient C is a 15-year-old with moderate pain (rated 3/5). The potential for life expectancy extension is moderate (rated 3/5). While not in immediate life-threatening situation, the patient requires medical (rule of rescue score: 2/5). Quality of life improvement is also moderate (score:3/5). He/she has been waiting for four months (waiting time score:3/5), has no dependents (parental status score: 1/5), no risky lifestyles (punishment score: 5/5) and due to his/her ordinal life has an instrumental score: 1/5). The patient's youth contribute with a age discrimination score: 5/5.

After inputting these data, the Vital Priority System calculates a priority score for each patient and classifies them as either priority or non-priority based on overall health, personal characterises, and societal values embedded in the rationing criteria. The system's recommendations are:

- Patient A is classified as a priority due to high scores in pain, life expectancy extension, quality of life improvement, parental status, and instrumental value.
- Patient B is classified as non-priority despite the immediate life-threatening condition, owing to lower scores in quality-of-life gains and life expectancy extension.
- Patient C is classified as non-priority as his/her condition is not immediately life-threatening, with only moderate improvements expected either in mortality or morbidity despite his/her youth.

This case study illustrates how the Vital Priority System integrates multiple criteria to support healthcare professionals in making fair, objective, and socially informed prioritisation decisions in resource-limited settings.

Discussion

Rationing is an inherent feature of healthcare systems operating under resource constrains. Prioritising patients is inherently complex and ethical challenging, as fair healthcare allocation reflects deeply held societal values that vary across individuals and cultures. Therefore, incorporating society's preferences is essential to ensure that patient prioritisation aligns with collective ethical standards.

The primary objective of this study was to identify the criteria on which the Portuguese population bases patient prioritisation and to develop a ML driven system reflecting these socially accepted rationing principles. Given that this is the first study to create such a prioritisation system within this context, there are no direct comparators in the literature.

The Vital Priority System demonstrated excellent performance in classifying patients according to priority, reaching an accuracy of 99% and a balanced F-measure across priority classes. This confirms the potential of ML, specifically Random Forest classifiers, to support complex healthcare decisions involving multiple ethical and social criteria.

The feature selection results align with existing literature emphasizing that illness-related characteristics, namely clinical need conceptualized as pain and immediate life threat, intervention effectiveness gauged by expected gains in life expectancy and quality of life, and waiting time hold the greatest weight in the prioritisation model. Among recipient characteristics, age emerged as the most influential criterion, followed by parental status. Conversely, merit-based criteria received the least weight, with a slightly higher priority given to patients whose illness was not perceived as self-inflicted compared to those with a socially valuable role.

Despite some methodological differences, the feature selection results align with existing literature emphasizing that civil society supports a plurality of rationing criteria[26,40,41], extending beyond pure efficiency considerations[25,26]. While some prior research highlights quality of life as more valued than life expectancy[42] and assigns greater importance to risky lifestyles and deservingness[43], our findings diverge by showing life expectancy as more heavily weighted than quality of life, with instrumental value and self-inflicted illness carrying relatively less influence.

A key strength of our study lies in operationalizing Portuguese societal values into a multi-criteria ML model, thereby advancing explicit and socially grounded healthcare rationing frameworks. The case study illustrates how the system operationalizes these complex, sometimes conflicting criteria, providing transparent and objective prioritisation that aligns with societal values. In particular, the prioritisation of Patient A underscores the system's capacity to balance clinical urgency, potential benefit, and social roles, while appropriately deprioritizing patients with lower expected gains or less immediate needs, as in Patients B and C.

Despite these promising findings, several limitations must be noted. First, the reliance on pre-collected societal opinion data means the model inherits existing biases, such as potential underestimation of certain groups' needs. Second, data were collected exclusively from lay persons, omitting perspectives from healthcare professionals who are central to applying prioritisation in clinical practice. Third, the sample was not representative and was obtained via an online survey, which may introduce selection bias. Robust and equitable ML systems require inclusive data collection processes involving key stakeholders, including marginalized groups[30], and demand transparency in AI application[34]. Importantly, the aim of this study was not to produce a representative population sample but to demonstrate the feasibility of creating an ML-based patient prioritisation system rooted in social values. Finally, the current system categorizes patients dichotomously as priority or non-priority. In this regard, future enhancements, including the introduction of ordinal priority rankings and incorporation of healthcare professionals' perspectives, will further increase the system's utility and robustness.

Conclusions

Making explicit healthcare rationing decisions is an unavoidable ethical and operational challenge, especially in systems constrained by limited resources. Delegating such decisions solely to healthcare professionals is neither fair nor sustainable, as it obscures the value judgements involved and excludes society, the main stakeholder, from participating in the rules that govern access to care. This study takes an important step forward by presenting, for the first time in the Portuguese context, a machine learning-based patient prioritisation system grounded explicitly in societal values. The Vital Priority System is an intelligent, socially informed, multi-criteria decision-support system that classifies patients as priority or non-priority cases based on values expressed by civil society.

By demonstrating that clinical need, treatment effectiveness, and patient characteristics such as age and parental status play significant roles in public preferences, we provide empirical evidence supporting a multi-dimensional approach to healthcare prioritisation. These findings contribute to the growing literature that recognises public support for ethically pluralistic rationing models, rather than purely efficiency-driven frameworks.

The developed Vital Priority System offers practical applications for healthcare decision-making, particularly in resource-constrained settings. It provides healthcare professionals with a transparent, objective tool that integrates multiple ethically relevant criteria to support equitable and socially aligned prioritisation decisions. This system can enhance fairness in patient allocation, improve transparency in decision processes, and potentially reduce moral distress among clinicians by aligning rationing decisions with public values.

Its user-friendly interface facilitates its implementation in real-world clinical or policy settings, enabling more consistent and socially legitimate decisions. While not intended to replace human judgement, the system can serve as a complementary aid in scenarios such as waiting list management, triage in crisis contexts, or the allocation of scarce treatments.

In sum, this study presents a novel, empirically grounded approach to integrating societal values into healthcare decision-making through artificial intelligence. As healthcare systems increasingly rely on digital tools, embedding fairness, transparency and inclusivity into their design becomes not only desirable, but essential.

Appendix

Questionnaire design

Suppose that you are behind two patients (A and B) that need treatment; however because of scarcity of resources you can only treat one of the patients. Imagine that both patients have the same characteristics except the one thing provided in each scenario. Please indicate your decision in accordance with the following degree of preference (See Supplemantary material 1):

Data availability

The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Received: 7 April 2025; Accepted: 30 July 2025
Published online: 30 September 2025

References

1. Russell, J., Greenhalgh, T., Burnett, A. & Montgomery, J. No decisions about Us without Us? Individual healthcare rationing in a fiscal ice age. *Br. Med. J.* **342** (7812), d3279. https://doi.org/10.1136/bmj.d3279 (2011).
2. Anderson, E. What is the point of equality? *Ethics* **109**, 287–337. https://doi.org/10.1086/233897 (1999).
3. Goddard, M., Hauck, K., Prekera, A. & Smith, P. Priority setting in health – a political economy perspective. *Health Econ. Policy Law*. **2006** (1), 79–90. https://doi.org/10.1017/S1744133105001040 (2006).
4. Gotowiec, S. & Cantor-Graae, E. The burden of choice: a qualitative study of healthcare professionals' reactions to ethical challenges in humanitarian crises. *Int. J. Humanitarian Action*. **2** (2), 1–10. https://doi.org/10.1186/s41018-017-0019-y (2017).
5. Bamforth, K., Rae, P., Maben, J., Lloyd, H. & Pearce, S. Perceptions of healthcare professionals' psychological wellbeing at work and the link to patients' experiences of care: A scoping review. *Int. J. Nurs. Stud. Adv.* **1** (5), 100148. https://doi.org/10.1016/j.ijnsa.2023.100148 (2023).
6. Drummond, M. F. & McGuire, A. *Economic Evaluation in Health Care: Merging Theory with Practice* (Oxford University Press, 2001).
7. Mounk, Y. The extraordinary decisions facing Italian doctors. Atlantic, 11 March. (2020). https://www.theatlantic.com/ideas/archive/2020/03/ who–gets-hospital-bed/607807/. Accessed 10 August 2024.
8. Goiri, F. Los medicos elegiran a quien ingresar en la UCI segun su esperanza de vida. In El Mundo. 2020. (2029). https://www.elmundo.es/ciencia-y-salud/salud/2020/03/20/5e73dc15fdddff8e518b4640.html. Accessed 20 March 2020.

9. Pinho, M. & Araújo, A. M. Personality and perceptions about the use of personal responsibility for illness as a health care rationing criteria. *Journal Neurosci. Psychol. Econ.* **15** (3), 137–151. https://doi.org/10.1037/npe0000160 (2022a).
10. Dettmer, J. & Italians Demand Compensation for Coronavirus Deaths. (2020). Available at https://www.voanews.com/covid-19-pandemic/italians-demand-compensation-coronavirus-deaths. Accessed 29 June 2020.
11. Shapiro, J. People with disabilities say rationing care policies violate civil rights. In NPR. 2020. (2020). https://www.npr.org/2020/03/23/820398531/people-with-disabilities-sayrationing-care-policies-violate-civil-rights?t1585419494913. Accessed 30 August 2020.
12. Carter, S. Ventilator rationing guidelines are discriminatory. In Technology and Ideas, Blomerang. (2020). https://www.bloomberg.com/opinion/articles/2020-04-10/coronavirus-ventilator-rationing-guidelines-are-discriminatory. Accessed 25 August 2020.
13. Pinho, M. & Dias Costa, E. Can mass media be an obstacle to rationing decisions? A case report from Portugal. *Int. J. Health Gov.* **25** (1), 3–11. https://doi.org/10.1108/IJHG-10-2019-0069 (2020).
14. Mitton, C., Smith, N., Peacock, S., Evoy, B. & Abelson, J. Public participation in health care priority setting: A scoping review. *Health Policy*. **91** (3), 219–228. https://doi.org/10.1016/j.healthpol.2009.01.005 (2009).
15. Manafò, E., Petermann, L., Vandall-Walker, V. & Mason-Lai, P. Patient and public engagement in priority setting: A systematic rapid review of the literature. *PLoS One*. **13** (3), e0193579. https://doi.org/10.1371/journal.pone.0193579 (2018).
16. Kapilashrami, A., Razavi, D. & Majdzadeh, R. Enhancing Priority-Setting Decision-Making process through use of intersectionality for public participation. *Int. J. Health Policy Manag.* **12**, 8095. https://doi.org/10.34172/ijhpm.2023.8095 (2023).
17. Pinho, M. & Araújo, A. How to fairly allocate scarce medical resources? Controversial preferences of healthcare professionals with different personal characteristics. *Health Econ. Policy Law*. **17** (4), 398–415. https://doi.org/10.1017/S1744133121000190 (2022b).
18. Cookson, R. & Dolan, P. Public views on health care rationing: a group discussion study. *Health Policy*. **49** (1–2), 63–74. https://doi.org/10.1016/s0168-8510(99)00043-3 (1999).
19. Williams, A. & Cookson, R. Equity in health. In (eds Culyer, A. & Newhouse, P.) Handbook of Health Economics, North-Holland. Elsevier, 1863–1907. (2000).
20. Brock, D. Fairness and health. In *Summary Measures of Population Health: Concepts, Ethics, Measurement and Applications* (ed. Murray, C.) 717–726 (World Health Organization, 2002).
21. Persad, G., Wertheimer, A. & Emanuel, E. Principles for allocation of scarce medical interventions. *Lancet* **329**, 224–227. https://doi.org/10.1016/S0140-6736(09)60137-9 (2009).
22. Scheunemann, L., Douglas, B. & White, M. The ethics and reality of rationing in medicine. *Chest* **140** (6), 1625–1632. https://doi.org/10.1378/chest.11-0622 (2011).
23. Beauchamp, T. & Childress, J. *Principles of Biomedical Ethics* (Oxford University Press, 2012).
24. Pinho, M. & Veiga, P. Attitudes of health professionals concerning bedside rationing criteria: a survey from Portugal. *Health Econ. Policy Law*. **15** (1), 113–127. https://doi.org/10.1017/S1744133118000403 (2020).
25. Olsen, J. A., Richardson, J., Dolan, P. & Menzel, P. The moral relevance of personal characteristics in setting health care priorities. *Soc. Sci. Med.* **57** (7), 1163–1172. https://doi.org/10.1016/s0277-9536(02)00492-6 (2003).
26. Dolan, P., Shaw, R., Tsuchiya, A. & Williams, A. QALY maximisation and people's preferences: A methodological review of the literature. *Health Econ.* **14** (2), 197–208. https://doi.org/10.1002/hec.924 (2005).
27. Baltussen, R. & Niessen, L. Priority setting of health interventions: the need for multi-criteria decision analysis. *Cost Eff. Resour. Alloc.* **21** (4), 14. https://doi.org/10.1186/1478-7547-4-14 (2006).
28. Darvishi, A., Daroudi, R., Yaseri, M. & Sari, A. Public preferences regarding the priority setting criteria of health interventions for budget allocation: results of a survey of Iranian adults. *BMC Public. Health*. **22** (1), 2038. https://doi.org/10.1186/s12889-022-14404-1 (2022).
29. Al-Quraishi, T., Chee, N. G., Osama, M. & Amoakoh, G. Advanced ensemble classifier techniques for predicting tumor viability in osteosarcoma histological slide images. *Appl. Data Sci. Anal.* 52–68. https://doi.org/10.58496/ADSA/2024/006 (2024).
30. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G. & Chin, M. H. Ensuring fairness in machine learning to advance health equity. *Ann. Intern. Med.* **169** (12), 866–872. https://doi.org/10.7326/M18-1990 (2018).
31. Salman, O. H. et al. Review on utilizing machine learning technology in the fields of electronic emergency triage and patient priority systems in telemedicine: coherent taxonomy, motivations, open research challenges and recommendations for intelligent future work. *Comput. Methods Programs Biomed.* **209**, 106357. https://doi.org/10.1016/j.cmpb.2021.106357 (2021).
32. Kadum, S. Y. et al. Machine learning-based telemedicine framework to prioritize remote patients with multi-chronic diseases for emergency healthcare services. *Netw. Model. Anal. Health Inf. Bioinforma*. **12** (11). https://doi.org/10.1007/s13721-022-00407-w (2023).
33. Al-Quraishi, T., Hussein, N. A. Q., Hussein, A. N. & Ahmed, A. L. Q. Big data predictive analytics for personalized medicine: perspectives and challenges. *Appl. Data Sci. Anal.* 32–38. https://doi.org/10.58496/ADSA/2024/004 (2024).
34. Wu, H., Lu, X. & Wang, H. The Application of Artificial Intelligence in Health Care Resource Allocation Before and During the COVID-19 Pandemic: Scoping Review JMIR AI. 2023; 2: e38397 https://doi.10.2196/38397 (2023).
35. Al-Quraishi, T., Zaboon, W., Mahdi, O., Naghavipour, H. & Aburghif, H. Enhancing Social Media Engagement Sentiment Prediction: A Random Forest and SMOTE-Based Approach with Explainable AI. In: Daimi, K., Al Sadoon, A. (eds) Proceedings of the Third International Conference on Advances in Computing Research (ACR'25). ACR 25 2025. Lecture Notes in Networks and Systems, vol 1346. Springer, Cham. https://doi.org/10.1007/978-3-031-87647-9_14
36. Mahmoud, A. Novel efficient feature selection: classification of medical and immunotherapy treatments utilising random forest and decision trees. *Intelligence-Based Med.* **10**, 100151. https://doi.org/10.1016/j.ibmed.2024.100151 (2024).
37. Parmar, A., Katariya, R. & Patel, V. A. Review on Random Forest: An Ensemble Classifier. In: Hemanth, J., Fernando, X., Lafata, P., Baig, Z. (eds) International Conference on Intelligent Data Communication Technologies and Internet of Things (ICICI) 2018. ICICI 2018. Lecture Notes on Data Engineering and Communications Technologies. ; 26. Springer, Cham. (2019). https://doi.org/10.1007/978-3-030-03146-6_86
38. Freund, Y. & Schapire, R. A Decision-Theoretic Generalization of on-Line Learning and an Application to Boosting, (1995).
39. Friedman, J. Greedy Function Approximation: A Gradient Boosting Machine, The Annals of Statistics, Vol. 29, No. 5, (2001).
40. Kasemsup, V., Schommer, J., Cline, R. & Hadsall, R. Citizen's preferences regarding principles to guide health care allocation decisions in Thailand. *Value Health.* **11** (7), 1194–1202. https://doi.org/10.1111/j.1524-4733.2008.00321.x (2008).
41. Gu, Y., Lancsar, E., Ghijben, P., Butler, J. & Donaldson, C. Attributes and weights in health care priority setting: a systematic review of what counts and to what extent. *Soc. Sci. Med.* **146**, 41–52. https://doi.org/10.1016/j.socscimed.2015.10.005 (2015).
42. Exel, J., Baker, R., Mason, H., Donaldson, C. & Brouwer, W. Public views on principles for health care priority setting: findings of a European cross-country study using Q methodology. *Soc. Sci. Med.* **126**, 128–137. https://doi.org/10.1016/j.socscimed.2014.12.023 (2015).
43. Pinho, M., Durão, N. & Zahariev, B. Are individual risky behaviours relevant to healthcare allocation decisions? An exploratory study. International Journal of Health Governance. 2022; 27(3): 342–355. (2022). https://doi.org/10.1108/IJHG-01-2022-0011

Author contributions

MP conceived and developed the study with support from FL. FL implemented the methodology. MP wrote the first draft of the manuscript with support from FL. All authors contributed significantly. All authors have read and agreed to the published version of the manuscript.

Funding
This work is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the support UID/05105: REMIT – Investigação em Economia, Gestão e Tecnologias da Informação.

Declarations

Competing interests
The authors declare no competing interests.

Ethics declarations
This study was conducted following the institutional and/or national research committee bodies. The ethical approval was obtained by the Portucalense University Ethical commission (CEI25_04_007). The questionnaires were conducted on a voluntary basis. Moreover, all participants give their writing consent.

Additional information
Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-025-14363-8.

Correspondence and requests for materials should be addressed to M.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

© The Author(s) 2025