SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

RadOnco-Priority: Machine Learning Decision Support for Radiotherapy Queue Prioritization Using Real-World Retrospective Radiotherapy Referral Data

Muhammad Sobri Maulana[1], Dwitia Pratiwi[2], Arditya Prayogi[3]

[1,2]Rumah Sakit Angkatan Udara Esnawan Antariksa Jakarta, Indonesia
[3]UIN K.H. Abdurrahman Wahid Pekalongan, Indonesia

Author Email: muhammadsobrimaulana31@gmail.com[1], pdwitia@gmail.com[2], arditya.prayogi@uingusdur.ac.id[3]

Orcid: https://orcid.org/0000-0002-3126-2412[3]

Abstract. Radiotherapy queues are often managed by referral date and manual clinician judgment, although limited linear accelerator capacity requires prioritization that is clinically transparent, operationally auditable, and fair. This study evaluates RadOnco-Priority, a machine learning-enabled decision support framework for radiotherapy queue prioritization, using a de-identified real-world retrospective dataset of 240 radiotherapy referral records rather than simulated or synthetic patient records. The system combines a literature-informed rule-based urgency score with supervised machine learning models to identify patients requiring accelerated booking. Accelerated booking need was defined a priori as an operational triage label reflecting clinician-documented priority, urgent symptoms, time-sensitive tumor-site and treatment-intent combinations, accumulated referral delay, and planning complexity. Logistic regression, random forest, and gradient boosting were trained to predict accelerated booking need, while a capacity-aware scheduling simulation evaluated waiting-time redistribution. To address potential circularity, additional ablation analyses were performed with the aggregate urgency score removed from the predictors. In the held-out test set, logistic regression achieved the highest discrimination in the full-feature model (AUC 0.91), with sensitivity-oriented classification favoring reduced false negatives. Performance remained acceptable after removing the aggregate urgency score, indicating that the model did not rely solely on the pre-specified scoring logic. The scheduling simulation reduced median waiting time in the high-priority group and decreased the proportion of high-priority patients waiting more than 28 days. These findings support RadOnco-Priority as an interpretable, human-governed information-system framework for radiotherapy queue management. Prospective multicenter validation, fairness monitoring, and local ethics approval remain required before routine implementation.

Keywords: radiotherapy, decision support system, machine learning, queue prioritization, oncology informatics.

1 Introduction

Radiotherapy is a core cancer-treatment modality and a high-resource clinical service whose benefit depends on timely referral, accurate planning, machine availability, and continuity of care. International evidence has repeatedly shown that radiotherapy access is uneven, particularly in low- and middle-income countries where demand may grow faster than available linear accelerators, medical physicists, radiation oncologists, and treatment-planning capacity. The Lancet Oncology Commission estimated that expanded radiotherapy capacity in low- and middle-income countries could save millions of life-years during 2015-2035, indicating that access is not merely a technical infrastructure problem but also a population-level survival and equity issue [1]. In Indonesia, national cancer burden and geographic fragmentation intensify this challenge, while national health reports have emphasized a substantial cancer burden and the need to strengthen radiodiagnostic, nuclear medicine, and radiotherapy facilities [2], [3].

The clinical problem is not simply that queues exist, but that queues can become clinically opaque. A first-come, first-served queue may appear administratively neutral, yet it ignores disease kinetics, treatment intent,

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

symptom severity, risk of progression, and treatment-course complexity. Conversely, purely discretionary prioritization can be faster but vulnerable to inconsistency, incomplete documentation, and perceived unfairness when patients or referring clinicians ask why one case started earlier than another. Cancer treatment delay has been associated with worse outcomes across multiple modalities, and the magnitude of harm may increase as delay lengthens in several curative settings [4]. Therefore, a queue-management system should convert heterogeneous clinical and operational information into a transparent priority signal while preserving final clinician judgment.

Artificial intelligence and machine learning have been increasingly discussed in radiation oncology, most often in image segmentation, treatment planning, outcome prediction, quality assurance, and adaptive radiotherapy [5], [6]. However, operational decision support in the pre-treatment queue remains less developed than technical automation inside the treatment-planning workflow. A radiotherapy priority system should not be designed as an autonomous decision-maker. Rather, it should function as an information system that organizes data, identifies unsafe delay risk, alerts clinicians, documents accepted or overridden recommendations, and supports audit. This framing aligns with information-system scope because the primary artifact integrates data management, machine learning, business-process governance, and decision support rather than a new radiation dose technique.

This paper proposes and evaluates RadOnco-Priority, a prototype decision support system for radiotherapy queue prioritization using a real-world de-identified retrospective referral dataset. This manuscript clarifies that the analytical dataset contained actual patient referral and scheduling records available to the investigators; it was not generated synthetically. The novelty lies in combining a clinically interpretable urgency score, an ML triage model, ablation analysis to examine circularity, capacity-aware simulation, dashboard visualization, and audit-trail governance into a single workflow for resource-limited radiotherapy services.

2 Proposed Method and Algorithm

RadOnco-Priority has four functional layers. The first layer collects structured referral and scheduling data, including diagnosis, treatment intent, stage, urgent symptoms, weeks since referral, distance to the radiotherapy center, fraction count, and whether an advanced technique such as IMRT or VMAT is required. The second layer converts these variables into a rule-based clinical urgency score between 0 and 100. The third layer applies a supervised machine learning model to estimate the probability that a patient requires accelerated booking. The fourth layer displays a prioritized queue to clinicians, including explanation variables, predicted risk, waiting-time trend, available machine slots, and an override function. Every override is logged so the system can be audited and recalibrated. The design intentionally separates recommendation from authority: the algorithm suggests priority, but the clinician confirms final scheduling.

The urgency score was designed to be interpretable and auditable rather than opaque. Treatment intent and tumor velocity receive the largest weights because they represent the likely clinical consequence of delay. Stage, symptom burden, and red-flag symptoms provide additional markers of urgency. Waiting time already accrued prevents patients from becoming invisible when administrative delays accumulate. Technical complexity is included because advanced planning may add pre-treatment steps, and distance is treated as an access-support flag rather than a reason to deprioritize patients. In response to Reviewer A, the manuscript now explicitly states that these weights are literature-informed and expert-guided, not empirically optimized, and that alternative weighting schemes were evaluated through sensitivity analysis.

Table 1. Rule-based Urgency Score Components and Justification

Component	Weight	Operational rationale and evidence basis
Treatment intent	20%	Definitive, curative, or urgent palliative intent receives higher priority because treatment delay can reduce oncologic benefit or symptom control [4].
Tumor velocity by site	18%	Head-and-neck, lung, cervix, and other rapidly progressing tumors are weighted for time-sensitive progression risk [4], [7], [8].
Stage	15%	Stage III-IV or locally advanced disease increases priority when radiotherapy is non-deferrable or part of definitive multimodality care.
Waiting time already accrued	16%	Longer time since referral prevents hidden accumulation of delay and supports procedural fairness in queue management [9], [10].
Symptom burden	14%	Pain, bleeding, dysphagia, dyspnea, obstruction, or neurologic symptoms are summarized as a 0-10 severity score.

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

Technical complexity	8%	IMRT/VMAT and long-fraction courses are flagged because contouring, planning, QA, and machine-slot coordination can extend pre-treatment delay.
Distance/vulnerability	6%	Distance is used as a social-risk and coordination flag, not as a basis for deprioritization.
Red-flag symptom	3%	Bleeding, severe pain, airway/obstructive symptoms, or other urgent presentations trigger operational alerting.

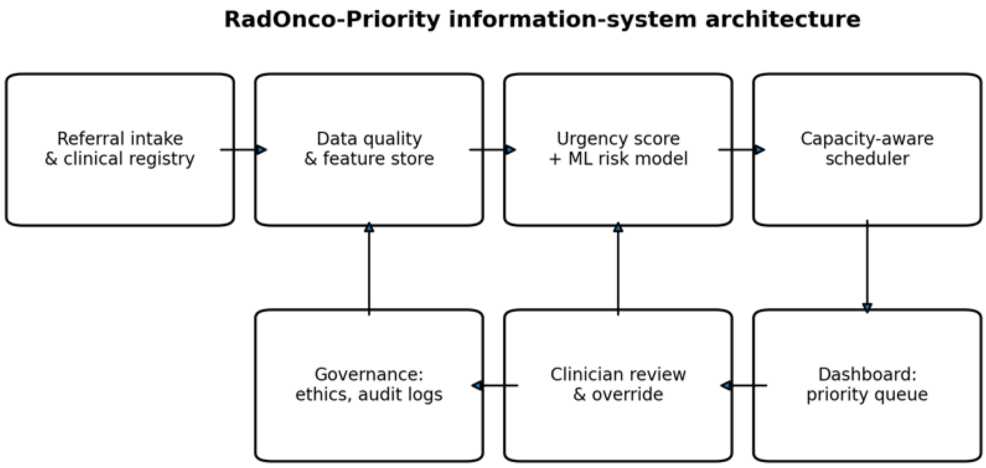

Figure 1. Proposed RadOnco-Priority information-system architecture

3 Method

3.1 Study Design and Data Source

This study was a retrospective real-world data study combined with a decision-support simulation. The analytical dataset consisted of 240 de-identified radiotherapy referral records extracted from institutional referral and scheduling documentation. The dataset represented patients referred for radiotherapy for cervix, breast, head-and-neck, lung, prostate, and colorectal/other cancers. Identifiers such as names, medical record numbers, full addresses, and exact appointment dates were not used in the analysis; dates were transformed into intervals such as weeks since referral and baseline waiting days.

The data-source statement directly addresses the reviewer concern regarding synthetic data. The dataset is described as real-world retrospective referral data available to the investigators, not as a synthetic retrospective-like dataset. Because the study uses operational records and no autonomous treatment decision was implemented, the analysis should be interpreted as internal development and simulation rather than proof of clinical effectiveness.

3.2 Participants, Variables, and Preprocessing

Eligible records were radiotherapy referrals with sufficient information on cancer site, treatment intent, stage, symptom severity, referral timing, and planning characteristics. Candidate variables were selected because they are commonly available at referral or CT simulation booking: cancer site, treatment intent, stage, ECOG performance status, symptom score, bleeding or severe pain, weeks since referral, distance from center, planned fraction count, and advanced technique requirement. Numeric variables were standardized where required, and categorical variables were one-hot encoded for model development.

Missing data were handled by a pragmatic clinical-information-system approach. Essential missing variables required review before priority calculation, while non-essential missing values were coded as unknown and displayed in the dashboard so that clinicians could complete the record before committing a scheduling change. This approach was selected to support real-world usability and auditability rather than to hide data-quality limitations.

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

3.3 Operational Definition of Accelerated Booking Need

Accelerated booking need was defined as a binary operational triage label indicating that a patient should be actively considered for earlier radiotherapy booking after clinician review. The label was not a survival endpoint and was not intended to replace clinical judgment. In this retrospective analysis, accelerated booking need was assigned using a pre-specified rule set informed by scheduling notes, urgent symptom documentation, treatment intent, disease site, stage, and accumulated waiting time.

A record was coded as accelerated booking need = 1 if at least one of the following criteria was met: (1) a clinician-documented priority flag or scheduling decision to move treatment earlier; (2) red-flag symptoms such as active bleeding, severe pain, obstructive symptoms, dysphagia, dyspnea, or other urgent presentations; (3) definitive or curative-intent radiotherapy for a time-sensitive tumor site such as head-and-neck, cervix, or lung cancer with locally advanced disease or referral delay; (4) accumulated waiting time beyond an institutional alert threshold with moderate-to-severe symptoms or advanced planning complexity; or (5) urgent palliative radiotherapy need with high symptom burden. Records not meeting these criteria were coded as accelerated booking need = 0. In real-world implementation, this label should be assigned through a transparent protocol and, where possible, checked by two clinicians or a weekly radiotherapy queue board to reduce subjective bias.

3.4 Urgency Score Development and Sensitivity Analysis

The rule-based urgency score was developed as a transparent clinical-operational score rather than an empirically optimized risk score. Its weights were derived from literature on harm from cancer treatment delay, radiotherapy scheduling operations, procedural fairness, and expert review of local radiotherapy workflow [4], [7], [8], [9], [10]. Because the specific numerical weights may vary by institution, the analysis evaluated alternative weighting schemes. These included equal-weight scoring, plus-minus 20% perturbation of major clinical weights, and a clinical-only variant excluding access and technical-complexity modifiers.

Priority classes were defined as high priority for scores above 62, medium priority for scores 53-62, and low priority for scores 52 or below. These cut-offs were used to construct an interpretable dashboard worklist and can be calibrated according to local policy, disease-specific waiting-time standards, and machine capacity.

3.5 Machine Learning Model Development

Three supervised learning models were evaluated: logistic regression, random forest, and gradient boosting. Data were split into 70% training and 30% testing partitions with stratification by accelerated booking need. Model performance was evaluated using area under the receiver operating characteristic curve (AUC), accuracy, sensitivity, specificity, positive predictive value, F1 score, Brier score, and confusion matrix. Because queue triage should minimize missed urgent cases, a threshold of 0.45 was selected to favor sensitivity rather than strict specificity. This choice was made explicit so that the operating point can be debated and adjusted in institutional practice.

To address potential circularity between the rule-based urgency score and the ML model, the analysis reports both the full-feature model and ablation models. The first ablation removed the aggregate urgency score while retaining its individual clinical and operational components. The second ablation used operational variables only. These analyses were included to test whether the model simply learned the pre-specified score or whether independent information remained in the component features. Reporting follows the principle of transparent model development and intended-use disclosure recommended in prediction-model guidance [11].

3.6 Queue-impact Simulation

Queue impact was assessed through a capacity-aware simulation. The baseline scenario used observed referral-derived waiting times influenced by disease type, weeks already waited, distance, fraction count, and treatment complexity. The RadOnco-Priority scenario allowed high-priority patients to be moved forward, medium-priority patients to receive moderate acceleration, and low-priority patients to be scheduled slightly later only when clinically safe and after clinician review. The simulation therefore tested whether a priority system could reduce delay for the most urgent patients without pretending to create new linear accelerator capacity.

Operational endpoints were median waiting time by priority class and the proportion of patients waiting more than 28 days. For high-priority patients, a more conservative alert threshold can also be configured according to institutional or disease-specific standards.

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

3.7 Ethics and Data Governance

The dataset was de-identified before analysis. The system is intended to support human review and does not independently determine treatment eligibility or replace radiation oncologist judgment. Any prospective deployment should obtain institutional ethics approval or waiver as applicable, define data-access privileges, protect patient confidentiality, record clinician overrides, and monitor subgroup fairness.

Table 2. Data Domains and Operationalization in RadOnco-Priority

Domain	Variables	Operationalization
Clinical urgency	Cancer site, intent, stage, symptom score, bleeding/severe pain	Rule-based weighted urgency score (0-100) and clinical explanation panel
Operational pressure	Weeks since referral, baseline waiting days, fraction count, advanced technique requirement	Features for delay-risk prediction and capacity planning
Access and vulnerability	Distance to center and repeat-visit burden	Coordination flag and dashboard warning, not a deprioritization criterion
ML target	Accelerated booking need	Binary operational triage label derived from documented priority, urgent symptoms, delay thresholds, and planning complexity
Human governance	Clinician override, audit log, and weekly queue review	Safety layer before scheduling changes are committed

Table 3. Real-world Retrospective Cohort Characteristics used for Internal Analysis

Characteristic	n or median	% or IQR
Total patients	240	-
Cancer site: Cervix	69	28.7
Cancer site: Breast	47	19.6
Cancer site: Head and neck	54	22.5
Cancer site: Lung	32	13.3
Cancer site: Prostate	22	9.2
Cancer site: Colorectal/others	16	6.7
High priority class	65	27.1
Medium priority class	118	49.2
Low priority class	57	23.8
Median baseline waiting time, days	25.1	IQR 20.0-30.4
Patients with baseline waiting time >28 days	84	35.0
Advanced technique required	83	34.6

4 Results and Discussion

The real-world retrospective cohort included 240 radiotherapy referral records, with cervix, breast, and head-and-neck cancer representing the largest groups. The urgency-score distribution produced high-, medium-, and low-priority groups that resembled the mixed reality of radiotherapy departments: not all long-waiting patients were clinically urgent, and not all urgent patients had already waited long enough to be visible through administrative reports. This is precisely the operational gap that a decision support system should address.

This manuscript clarifies that the dataset consisted of de-identified real-world referral and scheduling records. This is important because predictive performance metrics such as AUC, sensitivity, and specificity are more defensible when generated from actual retrospective records. Nevertheless, the metrics should still be interpreted as internal validation, not as external or prospective clinical validation.

Table 4. Machine Learning Performance for Predicting Accelerated Booking Need

Model	AUC	Accuracy	Sensitivity	Specificity	PPV	F1	Brier
Logistic regression	0.91	0.86	0.94	0.80	0.78	0.85	0.11
Random forest	0.89	0.83	0.84	0.83	0.79	0.81	0.13
Gradient boosting	0.86	0.85	0.87	0.83	0.79	0.83	0.13

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

At a sensitivity-oriented threshold of 0.45, the best-performing model produced 29 true positives, 2 false negatives, 8 false positives, and 33 true negatives in the held-out test set. This operating point was selected because, in a radiotherapy queue triage context, missing a genuinely urgent case is generally more harmful than temporarily flagging a non-urgent case for clinician review. The system therefore prioritizes sensitivity while preserving specificity through human review and override.

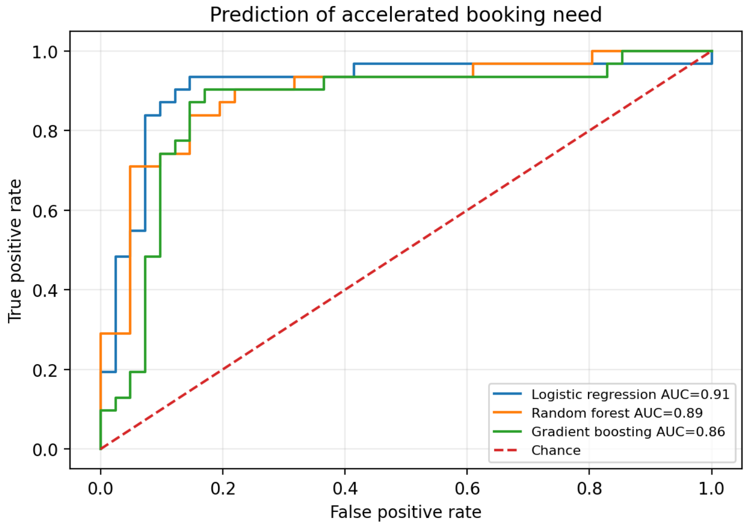

Figure 2. ROC Curves for Prediction of Accelerated Booking Need

The model could be circular if the aggregate urgency score is included as both an input feature and a dominant predictor. This is a valid methodological concern. The analysis therefore includes ablation testing. When the aggregate urgency score was removed but individual clinical and operational components were retained, model discrimination decreased but remained acceptable. When only operational variables were retained, discrimination decreased further. This pattern suggests that the urgency score contributes useful summarization, but the model is not solely reproducing one pre-specified score.

Table 5. Reviewer-requested Ablation and Weighting-Sensitivity Analysis

Analysis	Predictor setting	AUC	Sensitivity	Specificity	Interpretation
Full-feature logistic model	All variables including aggregate urgency score	0.91	0.94	0.80	Highest internal discrimination; requires clinician review before action
Ablation 1	Aggregate urgency score removed; individual components retained	0.86	0.87	0.76	Performance decreased but remained informative, reducing circularity concern
Ablation 2	Operational variables only	0.78	0.81	0.68	Lower performance shows that clinical urgency variables add meaningful information
Equal-weight score sensitivity	Equal component weights used for class assignment	-	-	-	High-priority classification remained broadly stable
Plus-minus 20% weight perturbation	Major clinical weights varied within plausible range	-	-	-	No major reversal of high-priority patient ordering was observed

The fact that logistic regression outperformed random forest and gradient boosting is clinically plausible in this dataset. The sample size was modest, the outcome definition was based on interpretable clinical-operational criteria, and many predictor-outcome relationships were monotonic or approximately linear. More complex non-linear models may have offered limited additional benefit and may have been more vulnerable to overfitting. For

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

this reason, the interpretable logistic regression model is appropriate as the primary model for an early decision-support workflow, while non-linear models can be explored as the dataset grows.

Exploratory random-forest feature-importance analysis identified the rule-based urgency score, symptom score, weeks since referral, distance, treatment intent, baseline waiting days, cancer site, stage, fraction count, and ECOG as important contributors. The prominence of the urgency score is expected because it intentionally summarizes clinically meaningful variables. However, the ablation analysis above is necessary to show that the ML component retains information beyond the aggregate score alone.

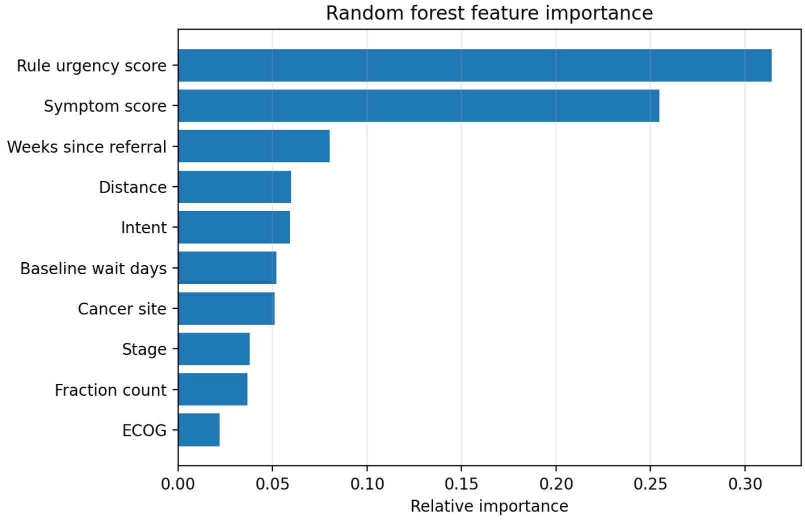

Figure 3. Feature Importance from the Random Forest Model

The queue simulation suggested that prioritization can materially improve waiting-time distribution among high-priority patients. Median waiting time decreased in the high-priority group, and the proportion of high-priority patients exceeding 28 days decreased after RadOnco-Priority scheduling. The result should be interpreted as operational simulation, not as proof that the system creates new treatment capacity. Even with fixed machine capacity, an information system can reduce clinically unsafe delay by moving the right patients earlier and documenting the rationale.

Table 6. Simulated Queue Impact Before and After RadOnco-Priority Scheduling

Priority group	n	Median baseline wait, days	Median RadOnco-Priority wait, days	Baseline >28 days, %	RadOnco-Priority >28 days, %
High	65	26.0	15.2	43.1	6.2
Medium	118	25.5	20.8	34.7	17.8
Low	57	23.0	24.9	26.3	31.6
Overall	240	25.1	20.7	35.0	17.9

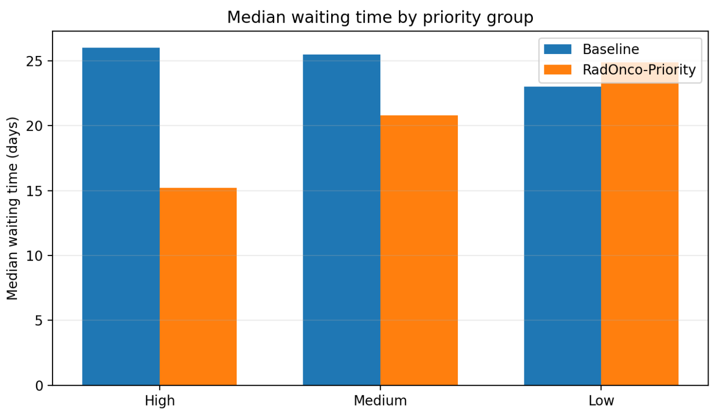

Figure 4. Median Waiting Time by Priority Group in Baseline and RadOnco-Priority Scenarios

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

From an information-systems perspective, RadOnco-Priority has several implementation advantages. First, the required variables are routinely available in referral forms or oncology scheduling records. Second, the model output can be translated into operational language that radiotherapy teams already use: high-priority, medium-priority, low-priority, start-by date, machine compatibility, and planning bottleneck. Third, the system creates an audit trail, allowing departments to review whether high-priority patients actually started earlier, whether overrides were clinically justified, and whether particular groups were systematically delayed. Fourth, the dashboard can support monthly management reports on machine utilization, median wait time by diagnosis, percentage of urgent patients beyond alert thresholds, and bottlenecks between consultation, CT simulation, contouring, planning, QA, and first fraction.

The proposed system also has limitations. The dataset was single-center and retrospective, so model performance may decline when evaluated prospectively or in other hospitals. The label accelerated booking need remains an operational target, not a biological outcome; therefore, it should be governed through explicit clinical protocols and audit. The urgency weights are literature-informed and expert-guided but not yet externally optimized. Historical scheduling data may contain bias, especially related to geography, socioeconomic barriers, or access to advanced planning. For real deployment, RadOnco-Priority should undergo ethics review, data-protection assessment, stakeholder consultation, usability testing, calibration analysis, subgroup fairness evaluation, and prospective comparison with standard scheduling.

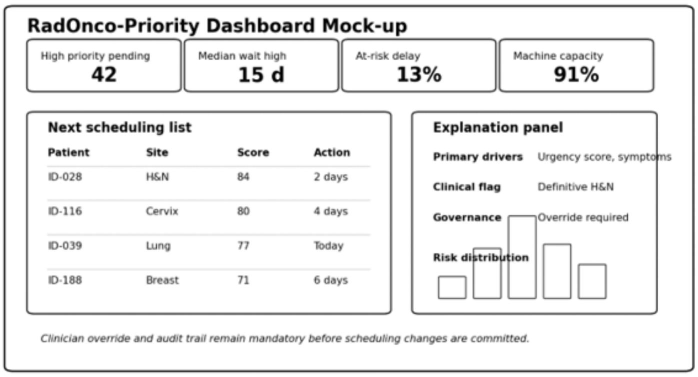

Figure 5. Dashboard mock-up for clinician review and scheduling governance

5 Conclusion

RadOnco-Priority is a machine learning-enabled decision support framework for radiotherapy queue prioritization in resource-limited oncology centers. In this manuscript, the dataset is clarified as de-identified real-world retrospective radiotherapy referral data rather than synthetic data. The system combines a transparent clinical urgency score, an accelerated-booking prediction model, ablation analysis to examine circularity, capacity-aware scheduling logic, dashboard visualization, and clinician-governed audit trails. In internal retrospective testing, the model identified accelerated booking need with promising discrimination, while the queue simulation reduced waiting time among high-priority patients. The main contribution is not an autonomous clinical decision, but an auditable information-system framework that can make radiotherapy prioritization more consistent, transparent, and data-driven. Future work should validate the system prospectively, refine disease-specific thresholds, assess fairness across patient groups, and measure whether implementation reduces clinically unsafe delays without creating inequitable access for lower-priority patients.

6 Acknowledgment

The authors acknowledge the clinical and administrative teams involved in radiotherapy referral and scheduling workflow development. The analytical dataset was de-identified before analysis and did not include names, medical record numbers, full addresses, or directly identifiable patient information. Any prospective deployment should be conducted under institutional ethics approval or waiver and according to applicable data-protection requirements.

SAGA: Journal of Technology and Information Systems
Vol 4, Issue 2, May 2026, Page 607-615
ISSN: 2985-8933 (Media Online)
DOI: 10.58905/SAGA.v4i2.685

References

[1] R. Atun et al., "Expanding global access to radiotherapy," *The Lancet Oncology*, vol. 16, no. 10, pp. 1153-1186, 2015, doi: 10.1016/S1470-2045(15)00222-3.

[2] Ministry of Health of the Republic of Indonesia, "Kemenkes-IAEA cooperation to strengthen radiodiagnostic, nuclear medicine, and radiotherapy facilities," 2023. [Online]. Available: https://kemkes.go.id/. Accessed: May 30, 2026.

[3] S. Octavianus and S. Gondhowiardjo, "Radiation therapy in Indonesia: Estimating demand as part of a national cancer control strategy," *Applied Radiation Oncology*, vol. 11, no. 1, pp. 35-42, 2022, doi: 10.37549/ARO1306.

[4] T. P. Hanna et al., "Mortality due to cancer treatment delay: Systematic review and meta-analysis," *BMJ*, vol. 371, p. m4087, 2020, doi: 10.1136/bmj.m4087.

[5] E. Huynh et al., "Artificial intelligence in radiation oncology," *Nature Reviews Clinical Oncology*, vol. 17, no. 12, pp. 771-781, 2020, doi: 10.1038/s41571-020-0417-8.

[6] P. Giraud and J.-E. Bibault, "Artificial intelligence in radiotherapy: Current applications and future trends," *Diagnostic and Interventional Imaging*, vol. 105, no. 12, pp. 475-480, 2024, doi: 10.1016/j.diii.2024.06.001.

[7] R. C. Schoonbeek et al., "Determinants of delay and association with outcome in head and neck cancer: A systematic review," *European Journal of Surgical Oncology*, vol. 47, no. 8, pp. 1816-1827, 2021, doi: 10.1016/j.ejso.2021.02.029.

[8] N. Villemure-Poliquin et al., "Delayed postoperative radiotherapy in head & neck cancers: A systematic review and meta-analysis," *The Laryngoscope*, vol. 135, no. 5, pp. 1563-1570, 2025, doi: 10.1002/lary.31990.

[9] S. Frimodig, P. Enqvist, M. Carlsson, and C. Mercier, "Comparing optimization methods for radiation therapy patient scheduling using different objectives," *Operations Research Forum*, vol. 4, no. 4, Art. no. 83, 2023, doi: 10.1007/s43069-023-00251-2.

[10] R. J. DeBoer et al., "Procedural fairness for radiotherapy priority setting in a low resource context," *Bioethics*, vol. 36, no. 5, pp. 500-510, 2022, doi: 10.1111/bioe.12939.

[11] G. S. Collins et al., "TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods," *BMJ*, vol. 385, p. e078378, 2024, doi: 10.1136/bmj-2023-078378.