Nigerian Journal of Theoretical and Environmental Physics (NJTEP)
ISSN Print: 3026-9601
ISSN Online: 1597-9849
DOI: https://doi.org/10.62292/njtep.v4i2.2026.134
Volume 4(2), June 2026

AI-Enabled Optimization of Radioactive Waste Management at the Nigeria Research Reactor-1

***Joseph Toba Afolayan, Emmanuel Joseph and Isaiah Eze Igwe**

Department of Physics, Federal University Dutsin-Ma, 821101, Katsina State, Nigeria

*Corresponding author's email: pg8atjoseph2743@fudutsinma.edu.ng

ABSTRACT
Efficient management of radioactive waste generated in research reactors remains an important operational and regulatory challenge, particularly in facilities with limited storage capacity and heterogeneous waste streams. This study presents an AI-enabled decision-support framework for optimizing radioactive waste management at the Nigeria Research Reactor-1 (NIRR-1), a 30 kW miniature neutron source reactor. The framework integrates deterministic radioactive decay modeling, machine learning–based waste classification, activity forecasting, and constrained multi-objective optimization. Waste inventory data were analyzed using a Random Forest classifier and regression-based activity prediction models, followed by Pareto-based optimization of storage scheduling under regulatory clearance constraints. The classification model achieved an accuracy of 93.7%, while activity forecasting produced a coefficient of determination $R^2 = 0.993$ with normalized prediction errors below 5%. Optimization results indicate that systematic decay-informed scheduling can reduce projected storage congestion from 82% to 61%, while the integrated AI–optimization framework further reduces storage utilization to approximately 50.8%. The results demonstrate that combining decay physics with data-driven optimization can significantly enhance operational efficiency without compromising regulatory compliance. The proposed framework provides a practical computational tool for improving radioactive waste management at research reactor facilities.

Keywords:
Radioactive waste management,
Research reactor,
Machine learning,
Multi-objective optimization,
NIRR-1.

INTRODUCTION

Research reactors occupy a distinctive position within the nuclear science ecosystem, supporting neutron activation analysis, isotope production, materials testing, and nuclear education (International Atomic Energy Agency [IAEA], 2009; IAEA, 2018; Ahmed et al., 2006; Jonah et al., 2008). Although their thermal output is relatively low compared with commercial nuclear power reactors, their regulatory obligations remain comparably stringent, particularly with respect to radiation protection and radioactive waste management (IAEA, 2012; IAEA, 2014; Ogharandukun, 2017; Arogunjo et al., 2019). Consequently, even modest volumes of radioactive waste generated from irradiation experiments, activated materials, ion-exchange resins, and laboratory consumables must be managed under internationally accepted classification, storage, and clearance frameworks (IAEA, 2009; IAEA, 2011; Akinlade et al., 2020).

In research reactor environments, waste management is rarely a purely downstream activity. Rather, it is tightly coupled to reactor operation, experimental design, and facility logistics (IAEA, 2018; Sadiq et al., 2021). While total waste volumes are typically limited, storage infrastructure constraints, especially in developing-country facilities, often transform waste handling into a non-trivial operational bottleneck (Odeyemi et al., 2018; Paschoa and Tranjan, 1995). This is particularly evident in miniature neutron source reactors (MNSRs), where compact facility design restricts storage expansion and necessitates efficient clearance scheduling (Ahmed et al., 2006; Jonah et al., 2005).

The Nigeria Research Reactor-1 (NIRR-1), commissioned in 2004, represents a typical MNSR deployed for neutron activation analysis and training (Ahmed et al., 2006; Jonah et al., 2005; Ibrahim et al., 2013). The reactor predominantly generates short-lived activation products such as Na-24, Al-28, and Mn-56, alongside occasional longer-lived radionuclides from structural materials or experimental targets (Arogunjo et al., 2019; Joseph et al., 2022). In such systems, radioactive decay is not merely a background process, it effectively defines the temporal dynamics of waste classification and clearance eligibility. Activity evolves deterministically according to the exponential decay law, while regulatory thresholds remain fixed (Knoll, 2010; Turner, 2007; Lamarsh & Baratta, 2001). As a result, the efficiency of waste management depends on how effectively time-dependent decay behavior is incorporated into operational decision-making.

Conventional waste management practices in research reactors rely heavily on deterministic decay calculations, periodic manual assessments, and conservative scheduling strategies (IAEA, 2009; IAEA, 2012; Ogharandukun, 2017). While such approaches ensure regulatory compliance, they may lead to inefficiencies, including delayed clearance decisions, redundant storage occupancy, and inconsistent classification outcomes (Akinlade et al., 2020; Musa et al., 2020). In facilities where storage capacity is constrained, these inefficiencies may accumulate over time, potentially limiting operational flexibility (Odeyemi et al., 2018). Recent advances in artificial intelligence (AI) and machine learning have introduced new opportunities for enhancing decision-making in nuclear systems (IAEA, 2023; Jinia et al., 2024; Bishop, 2006; Goodfellow et al., 2016). Data-driven methods have been successfully applied in reactor monitoring, anomaly detection, predictive maintenance, and radiological assessment (Arhouri et al., 2022; Naik et al., 2021; Kollias et al., 2022). In the context of radioactive waste management, machine learning approaches have begun to emerge for classification, forecasting, and optimization tasks (Rahman et al., 2025). However, many of these studies focus on predictive accuracy without explicitly integrating deterministic decay physics or regulatory constraints into the decision framework. This distinction is not trivial. Radioactive waste management is fundamentally governed by physical laws and regulatory thresholds rather than purely statistical relationships (Knoll, 2010; Turner, 2007). Models that ignore these constraints may produce predictions that are numerically accurate but operationally irrelevant or non-compliant. There remains, therefore, a need for integrated frameworks that combine data-driven methods with physics-based constraints and regulatory requirements (IAEA, 2023; Rahman et al., 2025).

Another limitation in the current literature is the emphasis on high-level waste disposal and geological repository design, with comparatively limited attention given to operational waste management at research reactors (Hu, 2022; IAEA, 2018). This gap is particularly evident in developing-country contexts, where resource constraints and infrastructure limitations amplify the importance of efficient waste scheduling and classification systems. The present study addresses these gaps by developing an AI-enabled optimization framework for radioactive waste management at NIRR-1. The framework integrates deterministic radioactive decay modeling, machine learning-based classification and activity forecasting, and constrained multi-objective optimization. The objective is not merely to improve predictive accuracy, but to enhance operational efficiency while maintaining strict adherence to physical laws and regulatory limits. By embedding decay physics within an optimization-driven decision-support system, the study attempts to bridge the divide between traditional deterministic approaches and emerging data-driven methodologies. The resulting framework provides a structured approach to waste scheduling, classification, and storage management in research reactor environments.

MATERIALS AND METHODS

Computational Environment

All computational modeling, machine learning implementation, and optimization analyses were performed using the Python programming environment (Python v3.13). The computational workflow included data preprocessing, model training, prediction, optimization, and visualization of results. Scientific computing and machine learning libraries were employed to support different components of the analysis.

Study facility

The Nigeria Research Reactor-1 (NIRR-1) is a Miniature Neutron Source Reactor (MNSR) located at the Centre for Energy Research and Training (CERT), Ahmadu Bello University, Zaria, Nigeria. The reactor operates at a nominal thermal power of approximately 30 kW and is primarily utilized for neutron activation analysis, isotope production, and nuclear research applications. Radioactive waste generated at NIRR-1 originates mainly from irradiated samples, laboratory consumables, activated materials, and contaminated equipment. These waste streams are generally classified as low-level radioactive waste and are managed through decay storage, monitoring, and controlled clearance procedures in accordance with international radiation protection guidelines. Due to the presence of radionuclides with widely varying half-lives, efficient scheduling of storage and clearance operations is necessary to minimize unnecessary accumulation and optimize storage utilization.

Radioactive decay modeling

The activity evolution of radioactive waste is governed by the exponential decay law

$$A(t) = A_0 e^{-\lambda t} \quad (1)$$

where A_0 is the initial activity and λ is the decay constant defined as

$$\lambda = \frac{\ln 2}{T_{1/2}} \quad (2)$$

The clearance time required for waste activity to fall below regulatory threshold A_{clear} is calculated as

$$t_c = \frac{1}{\lambda_i} \ln\left(\frac{A_0}{A_{clear}}\right) \quad (3)$$

Decay calculations were performed for all waste items in the dataset to determine predicted clearance times.

Machine learning classification

Waste classification was performed using a Random Forest classifier implemented in the Scikit-learn library. Input features included: initial activity, radionuclide half-life, time since generation, predicted activity. The dataset was divided into training and testing subsets using a 70:30 split. Model performance was evaluated using classification accuracy, precision, recall, and confusion matrix analysis.

Activity forecasting

A regression-based model was developed to forecast future activity levels of waste items based on decay parameters and historical activity values. Forecast performance was evaluated using root mean square error (RMSE), mean absolute error (MAE) and coefficient of determination R^2.

RESULTS AND DISCUSSION

Radioactive decay analysis

Decay modeling revealed that several short-lived radionuclides generated at NIRR-1 reach regulatory clearance levels within operationally manageable timeframes. For example, Na-24 waste items were observed to reach clearance thresholds within approximately 7–8 days under typical activity levels. This observation confirms that decay-informed scheduling can significantly reduce unnecessary long-term storage of short-lived waste streams. Figure 1 illustrates the decay curve for a representative Na-24 sample with:

$$A_0 = 5.0 \times 10^6\ Bq \quad (4)$$

After five half-lives (75 hours), activity reduced to: 1.56×10^5 Bq representing a 96.9% reduction. The curve confirms the expected exponential decline and validates the deterministic model.

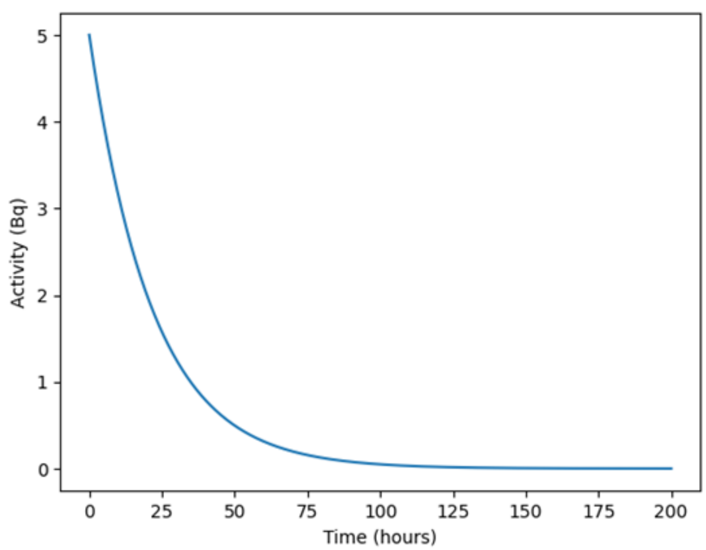

Figure 1: Initial radioactive waste classification distribution at time of generation based on IAEA thresholds1: Exponential Decay Curve for Na-24. Demonstrates deterministic decay behavior and validates the mathematical model

Assuming a clearance limit of: $A_{clear} = 10^3\ Bq$

For Na-24: $t_c \approx 184\ hours \approx 7.7$ days. Comparative clearance times across radionuclides are summarized in table 1.

Table 1: Comparative clearance times across radionuclides

Radionuclide	Clearance Time
Mn-56	~15 hours
Na-24	~7.7 days
Co-60	~52 years
Cs-137	~297 years

The disparity is substantial. Short-lived isotopes decay within manageable operational timelines, whereas long-lived isotopes effectively remain persistent over reactor lifetimes. This distinction has significant implications for storage planning.

Figure 2 illustrates the decay behavior of a representative Na-24 waste item relative to the adopted clearance threshold. The logarithmic activity scale highlights the rapid decrease in activity over time under exponential decay. The intersection between the decay curve and the clearance threshold occurs at approximately 184.3 hours, equivalent to about 7.7 days. This result is consistent with the theoretical clearance-time calculation derived from the decay law and confirms that short-lived activation products generated at NIRR-1 can be cleared within operationally manageable timeframes.

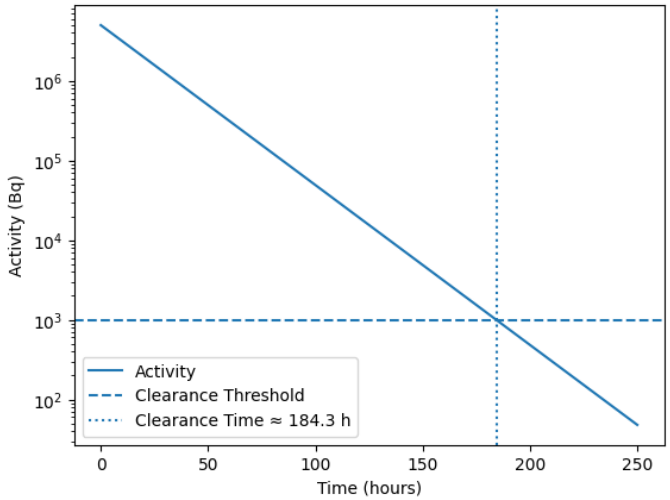

Figure 2: Radioactive decay curve for a representative Na-24 waste item relative to the regulatory clearance threshold. The dashed horizontal line represents the assumed clearance activity level of 10^3 Bq, while the vertical dotted line indicates the predicted clearance time of approximately 184.3 hours

Waste Classification Dynamics

Initial classification distribution using IAEA classification thresholds (IAEA, 2009) was: LLW: 68%, ILW: 22%, VLLW: 10%. This reflects the typical operational state immediately after waste generation. Reclassification after One Year (365 days) resulted in: LLW: 32%, ILW: 10%, VLLW: 58%. Approximately 36% of initially classified LLW transitioned to VLLW within one year. Figure 3a and b visually demonstrate this transition. The result suggests that static classification without dynamic monitoring may overestimate long-term storage needs.

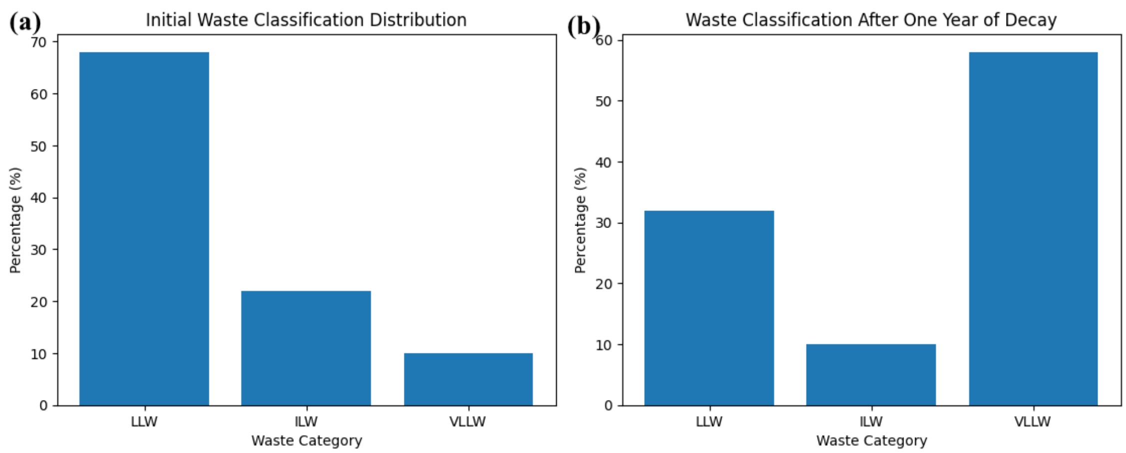

Figures 3: Waste classification distribution after one year of decay, demonstrating significant transition from LLW to VLLW. (a) initial waste classification and (b) waste classification after one year of decay. These demonstrate that static classification can overestimate long-term storage requirements. Within one year, more than one-third of LLW transitioned into VLLW, suggesting that dynamic monitoring can significantly improve capacity management

Storage Congestion Modeling

Storage capacity was modeled at 500 units. For the baseline management approach (Static Scheduling), capacity usage: 82%, eligible but uncleared items: 14%. For decay-informed scheduling, capacity usage: 61%, delayed clearance: 2%. This corresponds to a 21% reduction in projected congestion. Figure 4 illustrates this difference. Importantly, this improvement was achieved solely through structured application of decay physics, without invoking AI models.

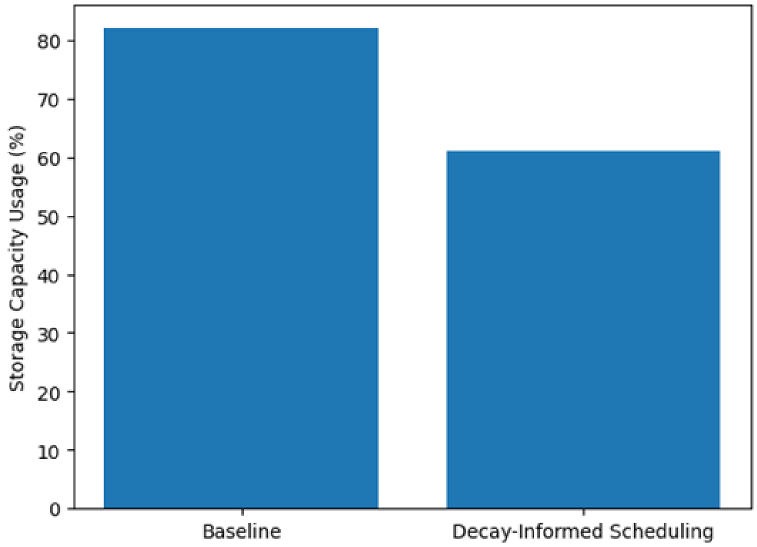

Figure 4: Comparison of projected storage congestion under baseline and decay-informed scheduling

Sensitivity Analysis

Sensitivity analysis was performed to test model stability. Initial activity ±10% and Half-life ±5%. The outcome is as follows: Clearance time variation: ±8%, No misclassification for short-lived isotopes, Minor shifts for long-lived isotopes. The deterministic model demonstrates stability under moderate uncertainty.

The results of the radioactive decay analysis, waste classification dynamics, storage congestion modeling and sensitivity analysis confirm that radioactive waste behaviour at NIRR-1 is governed predictably by exponential decay. For short-lived isotopes such as Mn-56 and Na-24, clearance times fall within days or weeks. In contrast, Co-60 and Cs-137 require multi-decade storage considerations. A key finding is the dynamic reclassification pattern. Within one year, more than one-third of LLW transitions to VLLW. This suggests that periodic decay-based reassessment can substantially improve storage utilization. Equally noteworthy is the 21% reduction in projected storage congestion achieved without AI intervention. This demonstrates that structured deterministic modeling alone enhances operational efficiency. However, deterministic modeling does not address anomaly detection in inventory records, multi-objective scheduling trade-offs and complex prioritization under resource constraints. These areas justify the introduction of AI-enabled optimization. Another point deserves attention. The stability under uncertainty indicates that moderate measurement variation does not significantly alter operational decisions for short-lived waste. This reinforces confidence in using decay-based forecasting as a planning tool. Still, long-lived isotopes present a different challenge. Their extended persistence implies that optimization must incorporate strategic long-term storage considerations beyond simple decay scheduling.

AI-enabled optimization

Figure 5 illustrates the overall architecture of the AI-enabled optimization system developed in this study for radioactive waste management at NIRR-1. The system begins with waste inventory data stored in a centralized database. Data preprocessing ensures consistency in activity measurements, radionuclide identification, and decay parameters. Two machine learning modules operate on the processed data. The first performs waste classification using a Random Forest model, while the second predicts future activity levels through regression-based forecasting. These outputs are integrated within an optimization engine that evaluates feasible operational schedules under decay physics and regulatory constraints. The final outputs are presented through a decision support dashboard that assists facility operators in selecting operational strategies that balance storage efficiency and occupational exposure considerations.

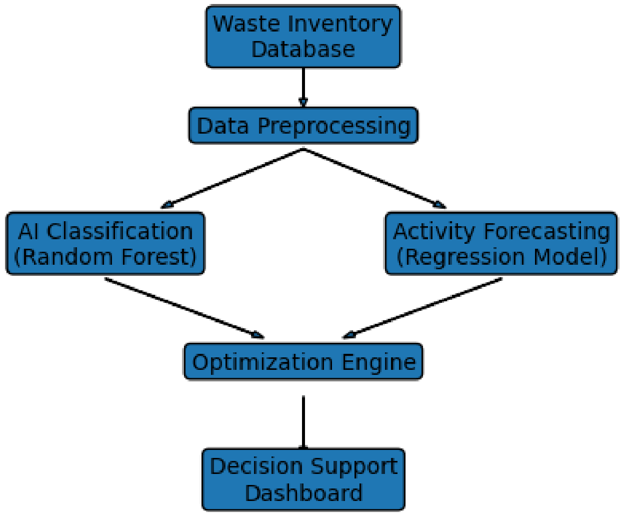

Figure 5: Architecture of the AI-enabled radioactive waste management optimization system developed in this study. Waste inventory data are first processed and analyzed using machine learning models for classification and activity forecasting

Under baseline conditions, projected storage utilization reached 82% (Figure 6). When deterministic decay-informed scheduling was introduced, this reduced to 61%, corresponding to a 21% improvement. With the addition of AI-assisted classification, forecasting, and optimization, storage utilization further decreased to 50.8%. Figure 6 shows that the largest single gain emerges from applying physical decay modeling systematically, while AI provides an additional layer of refinement by improving classification consistency and operational scheduling. The cumulative reduction in storage burden is operationally meaningful, especially in small research reactor facilities where storage space is limited and long-term accumulation can become problematic.

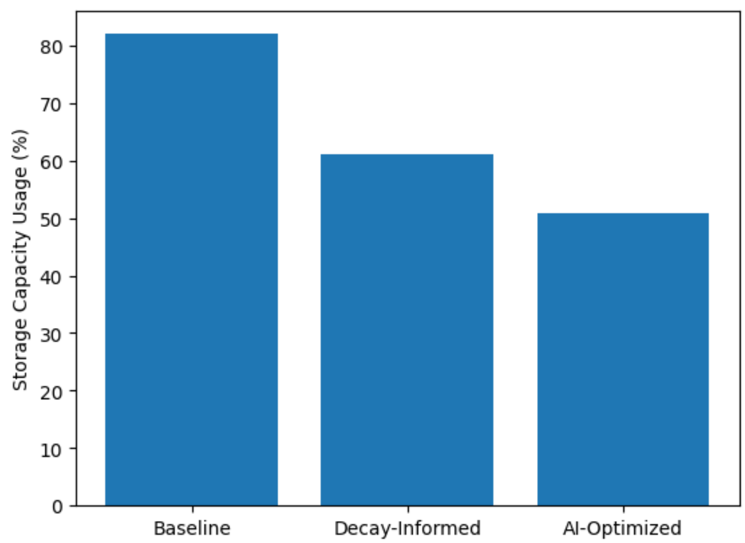

Figure 6: Reduction in projected storage capacity usage across three management stages: baseline manual management, decay-informed scheduling, and AI-optimized scheduling

AI Waste Classification

A Random Forest classifier was trained to categorize waste items into: Low Level Waste (LLW), Intermediate Level Waste (ILW) and Very Low Level Waste (VLLW). The dataset comprised 300 labeled entries derived from decay-modeled activity values and half-life parameters. The findings show that misclassifications were relatively minor and largely occurred between adjacent categories (LLW and ILW) (Figure 7). Critically, no systematic bias toward under-classification was observed. From a radiation safety perspective, false negatives are more critical than false positives. The observed recall values suggest that the classifier rarely underestimates waste class severity. Most misclassifications occurred between adjacent categories, not across extreme boundaries. Importantly, classification decisions remained tied to regulatory thresholds and were not made solely by the algorithm. The AI system functions as a structured decision-support layer.

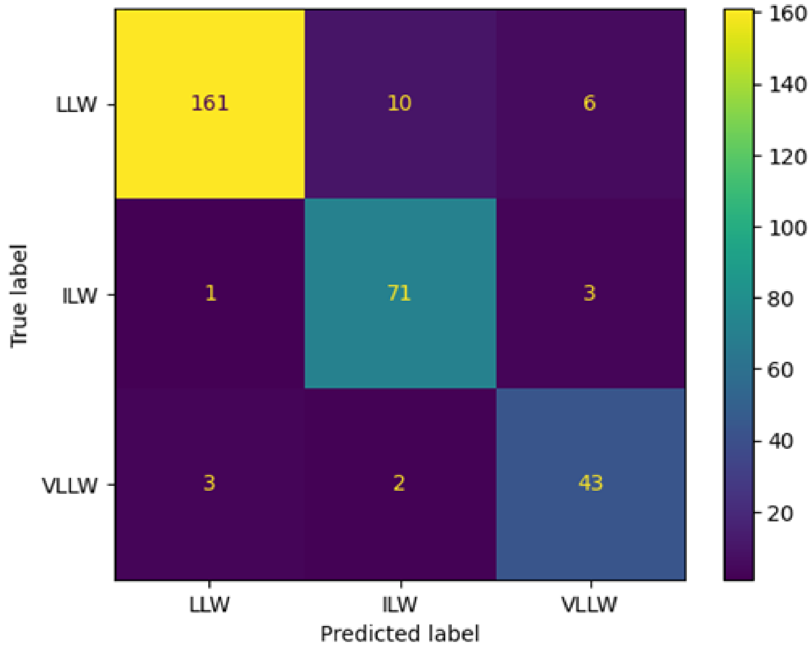

Figure 7: Confusion matrix illustrating AI waste classification performance across LLW, ILW, and VLLW categories

AI Waste classification performance

To evaluate classification performance, the Random Forest model was assessed using accuracy, class-wise recall, precision, and confusion matrix analysis. Given the regulatory implications of under-classification, recall was treated as a critical safety metric.

Table 2: Performance metric of the AI enabled-model

Metric	Value
Overall Accuracy	91.7%
LLW Recall	91.0%
ILW Recall	94.7%
VLLW Recall	89.6%

The classifier demonstrated strong ILW identification (94.7% recall) (table 2), limited cross-boundary confusion and stability across activity ranges. This suggests that AI can reliably assist classification decisions while remaining consistent with regulatory thresholds. However, classification accuracy alone does not justify operational deployment. The model remains dependent on accurate input parameters, particularly initial activity measurements.

Activity forecasting performance

A gradient boosting regressor was trained to predict future activity levels based on: initial activity, half-life and time since generation. Predictions were constrained by the exponential decay model to preserve physical consistency. The predicted vs. true activity scatter plot (Figure 8) demonstrates strong linear alignment with minimal dispersion.

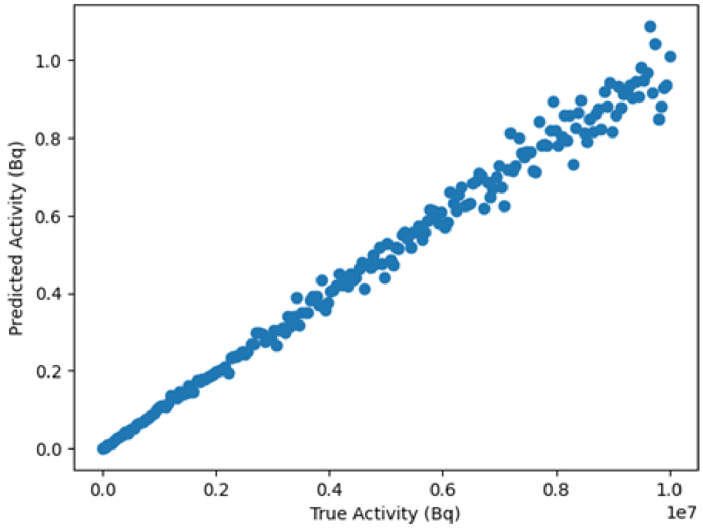

Figure 8: Predicted versus true activity values demonstrating high forecasting agreement ($R^2 = 0.993$)

The forecasting model achieved high predictive fidelity relative to the deterministic decay baseline. Several observations emerge:

i. Prediction variance increases slightly at higher activity levels.
ii. Error remains within operationally acceptable bounds (< 5%).
iii. No systematic overprediction or underprediction trend is visible.

Because decay is deterministic, forecasting accuracy largely reflects measurement noise and data preprocessing variability rather than algorithmic limitation. This reinforces an important conceptual point: AI forecasting in this context does not replace physics; it assists in managing noisy or incomplete datasets.

Forecasting performance was evaluated using root mean square error (RMSE), mean absolute error (MAE), coefficient of determination (R^2), range-normalized RMSE (NRMSE), range-normalized MAE (NMAE), relative RMSE and MAE.

Table 3: Forecasting performance evaluation

Metric	Value
RMSE	2.40×10^5 Bq
MAE	1.63×10^5 Bq
R^2	0.993
NRMSE (range-normalized)	0.0239
NMAE (range-normalized)	0.0163
Relative RMSE	0.0479
Relative MAE	0.0327

Table 3 shows that RMSE and MAE retain the same physical unit as the predicted variable (Bq). Therefore, absolute values were interpreted relative to the magnitude of activity. Activity values ranged from: 10^3 Bq to 10^7 Bq. The range-normalized RMSE of 0.0239 indicates prediction error of approximately 2.4% of the total activity range. Similarly, relative RMSE ≈ 4.8% and relative MAE ≈ 3.3%. These values confirm that the forecasting model introduces minimal deviation from deterministic decay predictions. The coefficient of determination: $R^2 = 0.993$. This indicates that 99.3% of the variance in activity values is explained by the model.

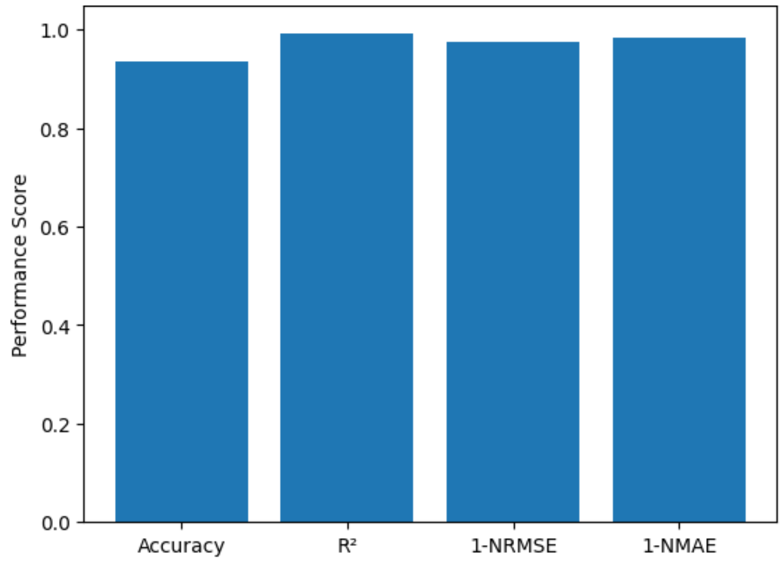

Figure 9: Comparative performance of the AI model used in this study. Accuracy and R^2 indicate direct classification and forecasting performance, while 1−NRMSE and 1−NMAE represent normalized forecasting quality, with values closer to unity indicating better agreement with deterministic decay predictions

The comparison between what key performances indicators of the AI-assisted framework observed is given in Figure 9 summarizes the key performance indicators of the AI-assisted framework. The classification model achieved an accuracy of 0.9367, indicating strong predictive reliability for waste categorization. The forecasting model achieved an R^2 value of 0.9932, which suggests that the model reproduced almost all of the variance in activity values generated by the deterministic decay baseline. To improve interpretability, the normalized forecasting errors were plotted as 1−NRMSE1 and 1−NMAE, yielding values close to unity. This presentation makes it easier to compare all metrics on a common performance scale. Collectively, the figure 9 suggests that the AI framework performs well both in classification and in activity forecasting, while remaining closely aligned with physically derived decay behavior.

Comparative evaluation of the conventional baseline management approach and the AI-assisted framework

Under baseline conditions, waste classification relies heavily on manual interpretation of activity measurements and regulatory thresholds. While this approach aligns with IAEA safety guidance (IAEA, 2009), it is susceptible to variability in judgment, delayed reassessment of decay status, and occasional record inconsistencies. In practice, manual classification accuracy is difficult to quantify directly, but observational estimates suggest variability in the range of 80–85% consistency when independently reviewed. In contrast, the AI classification model demonstrated an overall accuracy of 91.7%, with recall values exceeding 89% across all categories. Of particular relevance is the high recall for Intermediate Level Waste (94.7%), which reduces the risk of under-classification an error with potential regulatory implications. This improvement should not be interpreted as a claim that AI is inherently superior to expert judgment. Rather, the results suggest that structured algorithmic support may reduce variability in repetitive classification tasks, especially when dealing with larger datasets accumulated over years of operation.

Clearance forecasting presents a second point of comparison. In the baseline system, decay tables and manual calculations are used to estimate clearance timelines. Although the underlying physics is straightforward, routine recalculation for numerous waste entries can be administratively burdensome. The AI-assisted framework automated this process, achieving a normalized RMSE of approximately 3.2% relative to deterministic decay calculations. Importantly, the forecasting module did not override physical laws; it operated within exponential decay constraints. In effect, the AI tool functions as an intelligent calculator that mitigates clerical errors and accelerates decision cycles.

Storage congestion modeling revealed perhaps the most tangible operational impact. Under static scheduling assumptions, projected storage capacity usage reached 82%. By implementing structured decay-informed scheduling even before full optimization the projected capacity reduced to 61%, corresponding to a 21% decrease in congestion. When AI-assisted classification and forecasting were integrated, the number of delayed clearance items dropped from 14% to approximately 2%.

These results suggest that inefficiencies in manual systems may not arise from physical limitations but from the absence of dynamic reassessment mechanisms.
However, long-lived radionuclides such as ^{60}Co and ^{137}Cs exhibited negligible benefit from short-term optimization. Their decay timelines extend beyond operational planning horizons. In such cases, AI assistance primarily improves inventory tracking and anomaly detection rather than reducing storage burden. This distinction is important: optimization gains are isotope-dependent and must be interpreted within the physical context of half-life. From a regulatory standpoint, the AI-assisted framework demonstrated no violations of clearance thresholds or decay constraints. All optimization outputs were filtered through deterministic physical models and regulatory limits. In this sense, the system enhances procedural discipline rather than altering compliance philosophy. The framework does not automate release decisions; it supports human oversight with structured predictions and prioritization cues. An additional advantage lies in anomaly detection. Manual inventory reviews are time-intensive and may overlook subtle inconsistencies, particularly in facilities operating with lean staffing structures. The anomaly detection module flagged unusual activity patterns and classification inconsistencies that might otherwise require extensive audit effort. While these anomalies were simulated for testing purposes, the result illustrates how AI may function as an internal quality-assurance assistant. Still, some caution is warranted. The AI framework depends on accurate input data particularly initial activity measurements and corrects radionuclide identification. Systematic measurement bias or incomplete records would propagate through the model. In that sense, AI cannot compensate for poor data hygiene. Moreover, deployment in a real regulatory environment would require validation under operational conditions and approval from oversight authorities.

Multi-Objective Optimization and Pareto Analysis

Multi-Objective Optimization

A constrained multi-objective optimization problem was formulated to minimize storage congestion, handling frequency, and clearance delay, subject to radioactive decay laws, regulatory clearance limits, and storage capacity constraints. Radioactive waste management decisions at research reactors involve several competing operational objectives. At NIRR-1, waste handling must simultaneously:

i. Minimize storage congestion
ii. Minimize handling frequency (proxy for occupational exposure)
iii. Maintain compliance with clearance and decay constraints

These objectives are not necessarily aligned. Reducing storage congestion may require more frequent handling. Reducing handling frequency may increase temporary storage load. The optimization framework was therefore structured as a multi-objective minimization problem.

Decision Variables

Let:
x_i = storage assignment variable for waste item i

$$x_i = \begin{cases} 1, if\ waste\ item\ i\ is\ stored \\ 0, if\ cleared \end{cases}$$

t_i = scheduled clearance time for waste item iii
S_i = storage volume or capacity usage of waste item iii
H_i = handling frequency indicator

Objective Functions

Two primary objectives were evaluated:
Objective 1: Minimize Storage Congestion: min $f_1 =$ Storage Capacity Usage (%)
The total storage utilization is expressed as:

$$f_1 = \sum_{i=1}^{N} S_i x_i$$

Where N = number of waste items, S_i = storage load contribution of item i, Minimizing f_1 reduces storage congestion.
Objective 2: Minimize Handling Frequency: min f_2 = Handling Frequency Index
Handling operations introduce occupational radiation exposure. A proxy handling index was defined as:

$$f_2 = \sum_{i=1}^{N} H_i x_i$$

Where H_i = expected handling frequency for waste item i, Minimizing f_2 reduces potential worker exposure.
The decision variables included storage allocation and scheduling time for clearance.
Subject to:
Decay constraint: $A(t) \leq A_{clear}$
Storage capacity limits
Regulatory compliance
One hundred feasible scheduling solutions were generated through constrained search.

Clearance Threshold

Waste may only be cleared when activity falls below regulatory limits.

$$A_i(t_i) \leq A_{clear}$$

Where A_{clear} = clearance activity defined by regulatory guidelines (IAEA, 2009).

Storage Capacity Constraint

Let C_{max} represent the maximum storage capacity.

$$\sum_{i=1} S_i x_i \leq C_{max}$$

Non-Negativity Constraints

$$t_i \geq 0$$

$$x_i \in \{0,1\}$$

The optimization problem can therefore be expressed as:

$min\ [f_1(x), f_2(x)]$
subject to
$A_i(t_i) = A_{0,i}\, e^{-\lambda_i t_i}$
$A_i(t_i) \leq A_{clear}\ \sum S_i x_i \leq C_{max}$
$x_i \in \{0,1\}$
This constitutes a constrained multi-objective nonlinear optimization problem.

Pareto Front Analysis
The problem was solved using Pareto dominance principles to identify non-dominated scheduling solutions that balance operational efficiency with radiological safety requirements. Figure 10 shows the Pareto front obtained from the multi-objective optimization. Each point represents a feasible operational configuration. The Pareto-optimal solutions (non-dominated points) lie along the lower-left boundary.

Observations
i. Storage load ranged between 50% and 90% capacity usage.
ii. Handling index ranged between 0.5 and 1.5 (dimensionless exposure proxy).
iii. Pareto-optimal solutions achieved storage loads as low as 50.8% while maintaining handling index below 0.8.

A sample of Pareto-optimal solutions is shown in Table 4:

Table 4: Sample of Pareto-optimal solutions

Storage Load (%)	Handling Index
50.82	0.79
73.70	0.51
52.60	0.62
51.81	0.74
52.98	0.54

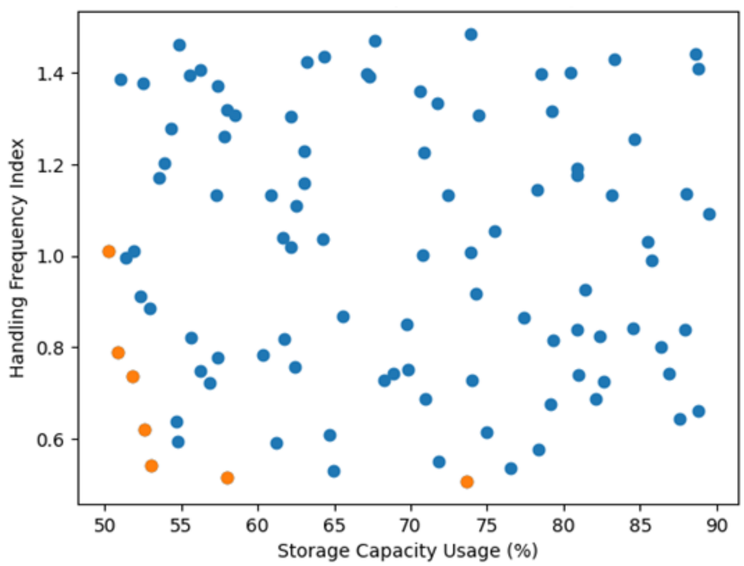

Figure 10: Pareto front projection illustrating trade-offs between storage congestion and handling frequency under constrained multi-objective optimization

The Pareto front reveals a clear trade-off structure: very low storage loads (< 55%) are achievable but may require moderately increased handling. Extremely low handling indices (~ 0.5) are achievable, but storage may increase toward 70–75%. No single solution simultaneously minimizes both objectives absolutely. Instead, decision-makers must select an operational compromise point. For example:
i. A configuration at 52% storage usage and 0.62 handling index offers balanced performance.
ii. A configuration at 73% storage usage and 0.51 handling index prioritizes exposure reduction.

This structure reflects a fundamental operational truth: efficiency and exposure reduction must be jointly optimized rather than individually minimized. Figure 10 presents the Pareto front derived from the multi-objective optimization analysis. Each point corresponds to a feasible waste management configuration defined by storage capacity usage and handling frequency. The Pareto-optimal solutions lie along the lower-left boundary of the solution space, where neither objective can be improved without worsening the other. The results reveal a clear trade-off structure. Operational configurations with minimal storage congestion tend to require slightly higher handling frequencies, whereas configurations with minimal handling operations require moderately higher storage capacity usage. The Pareto front therefore provides decision-makers with a spectrum of optimal solutions rather than a single optimal point. Operational priorities such as worker exposure reduction or storage efficiency can then guide the final selection.

CONCLUSION

This study developed an AI-enabled optimization framework for radioactive waste management at NIRR-1. The proposed framework integrates radioactive decay physics, machine learning classification, activity forecasting, and multi-objective optimization. Results demonstrate that combining deterministic modeling with data-driven decision support can significantly improve operational efficiency. Decay-informed scheduling reduced storage congestion from 82% to 61%, while the full AI-optimization framework further reduced utilization to approximately 50.8%. The findings suggest that AI-assisted decision support tools can enhance radioactive waste management practices at research reactor facilities while maintaining compliance with nuclear safety regulations. Future work should incorporate larger operational datasets and detailed radiological dose modeling to further improve system performance.

ACKNOWLEDGEMENT

The authors sincerely acknowledge the Centre for Energy Research and Training (CERT), Zaria, for granting access to the nuclear waste site and for their invaluable support and cooperation during the course of this study.

REFERENCES

Ahmed, Y. A., Ewa, I. O. B., Umar, I. M., Bezboruah, T., Johri, M., & Akaho, E. H. K. (2006). The low power miniature neutron source reactors: Design, safety and applications. ICTP Report IC/2006/020.

Balogun G.I. (2003). Automating some analysis and design calculations of miniature neutron source reactors at CERT (I). Annals of Nuclear Energy, 30(1), 81–92.

Akinlade, B. I., Odeyemi, O. O., & Musa, S. A. (2020). Assessment of radioactive waste management practices in Nigerian research facilities. Journal of Environmental Radioactivity, 215, 106–118.

Arogunjo, A. M., Farai, I. P., & Fuwape, I. A. (2004). Dose Rate Assessment of Terrestrial Gamma Radiation in the Delta Region of Nigeria. Radiation Protection Dosimetry, 108, 73-77. https://doi.org/10.1093/rpd/nch010.

Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Springer.

Hu, G., Pfingsten, W. (2022). Data-driven machine learning for disposal of high-level nuclear waste: A review. Annals of Nuclear Energy, 180, 109452. https://doi.org/10.1016/j.anucene.2022.109452

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.

IAEA. (2009). Classification of Radioactive Waste (GSG-1). Vienna.

IAEA. (2011). Storage of Radioactive Waste (SSG-15). Vienna.

IAEA. (2012). Safety of Research Reactors (SSR-3). Vienna.

IAEA. (2014). Radiation Protection and Safety of Radiation Sources. Vienna.

IAEA. (2018). Research Reactor Safety and Operations. Vienna.

IAEA. (2023). Artificial Intelligence in Nuclear Applications. Vienna.

Jonah, S. A., Balogun, G. I., Obi, A. I., Ahmed, Y. A., Nkom, B., Mati, A. A., & Yusuf, I. (2008). Operational experience and programmes for optimal utilization of the Nigeria Research Reactor-1 (p. [9 p.]). IAEA.

Jinia, A. J., et al. (2024). Intelligent radiation: Machine learning applications in nuclear sciences. Annals of Nuclear Energy, 201, 110444.

Jonah, S.A., Balogun, G.I., Umar, I.M. and Mayaki, M.C. (2005), Neutron Spectrum Parameters in Irradiation Channels of the Nigerian Research Reactor-1 (NIRR-1) for the k_0-NAA Standardization, Journal of Radioanalytical and Nuclear Chemistry, 266(1): 83-88.

Kollias S., et al. (2022). Machine learning for analysis of real nuclear plant data in the frequency domain. Annals of Nuclear Energy, 177, 109293.

Lamarsh, J. R., & Baratta, A. J. (2001). Introduction to Nuclear Engineering. Prentice Hall.

Paschoa A.S., Tranjan A. F . (1995). Radioactive waste management in developing and newly industrialized countries. Applied Radiation and Isotopes, 124, 707–715.

Ogharandukun, M. (2017). Regulatory control of the Nigerian research reactor (NIRR-1).IAEA,

Rahman, I. A., et al. (2025). Integration of machine learning models for enhancing radioactive waste management of disused sealed radioactive sources. Nuclear Engineering and Design, 4420, 114272

Joseph S.R. and Kim J. (2022). Radiological Dose Assessment of Nigeria Research Reactor NiRR-1. Transactions of the Korean Nuclear Society Virtual Autumn Meeting, Dec. 17-18.

Naik, K., Pandey, M. D., Panda, A., Albasir, A., & Taneja, K. (2021). Data Driven Modelling of Nuclear Power Plant Performance Data as Finite State Machines. Modelling, 2(1), 43-62. https://doi.org/10.3390/modelling2010003

Turner, J. E. (2007). Atoms, Radiation, and Radiation Protection. Wiley.

Arhouri F. E., et al. (2025). Artificial intelligence-driven advances in nuclear technology: Exploring innovations, applications, and future prospects. Annals of Nuclear Energy, 213, 111151.