# AI-Based Clinical Decision Support System for Predicting Renal Impairment Risk in Cancer Patients

Elda Xhumari
*Department of Computer Sciences, Faculty of natural Sciences*
*University of Tirana*
Tirana, Albania
0000-0002-1206-6424

Enrik Dollaj
*Department of Computer Sciences, Faculty of natural Sciences*
*University of Tirana*
Tirana, Albania
enrik.dollaj@fshnstudent.info

Flavio Corradini
*Computer Science - School of Science And Technology*
*University of Camerino*
Camerino, Italy
0000-0001-6767-2184

***Abstract* — Renal insufficiency is a frequent and serious side effect of cancer and frequently is compounded by cytotoxic chemotherapy and nephrotoxic medications. Conventional approaches, such as serum creatinine and eGFR measurement have low sensitivities to detect early changes. Here, we present an Artificial Intelligence (AI) driven Clinical Decision Support System (CDSS) focused on estimating kidney dysfunction risk based on real-world cancer care data. With supervised machine learning such as Random Forest, XGBoost, Support Vector Machines, and Multilayer Perceptrons, we examined the discriminative model for 5,000 patient cases. Preprocessing of raw data included handling the missing values, feature selection, calibrating classes with the help of SMOTE. No model had better performance than XGBoost, whose F1-score was 0.82 and AUC was 0.88. We generated global explanations identifying the most influential features (e.g., eGFR, creatinine, tumor type) using SHAP summary plots, and local patient level explanations illustrating individual risk contributions using LIME and SHAP waterfall plots. Our experimental results demonstrated that AI algorithms can be used in a CDSS to aid with early intervention, improve kidney safety, and personalize the treatment plans for oncology. While the data were single center in origin and there was not external validation, the system demonstrates significant potential for integration into electronic medical records and testing across different clinical backgrounds is needed.**

***Keywords* — *Clinical Decision Support, Renal Impairment, Cancer Patients, Machine Learning, Explainable AI, XGBoost, Electronic Health Records***

## I. INTRODUCTION

In cancer patients, the possibility for them to develop renal impairment is not only common but also serious, which is often aggravated by such treatments as chemotherapy.[1] The impact of renal impairment on treatment outcomes is significant. If the onset of treatment is delayed, there is a greater chance of adverse clinical outcomes and quality of life will suffer accordingly. However, measuring methods like the traditional serum creatinine or estimates of glomerular filtration rate (eGFR), which get used extensively in renal functional assessment, could not show sufficient sensitivity to pick up early or subtle declines in kidney function in cancer patients with already complicated clinical and treatment background.[2]

The study proposes evaluating a Clinical Decision Support System (CDSS) powered by artificial intelligence (AI), utilizing machine learning (ML) models to predict and stratify the risk ofrenal dysfunction in cancer patients. Designed to process a wide range of structured clinical data, including demographic information such as age and gender, laboratory tests (blood biomarkers), medication history; to match the oncology and nephrology experts in order to treat patients for a long time.

A variety of supervised ML algorithms were employed, including Random Forest, XGBoost, Support Vector Machines, and Multilayer Perceptrons, with a particular focus on interpretability by using AI that can be explained. These models were trained and validated in order to assess their performance in binary classification, multiclass risk stratification, regression, and even time-aware modeling tasks.

The primary goal of this study is to formulate a supervised binary classification problem that predicts the risk of renal impairment in cancer patients based on clinical and biochemical features. All this will show how AI can help close the current diagnostic gap for clinicians, providing them with an advanced, interpretable, and data-driven tool to support early intervention; personalization of treatment pathways; and ultimately improved renal safety in oncology.[3] Additionally, the study considers major ethical, regulatory, deployment issues which can help guide responsible integration of AI into real-world clinical activities.

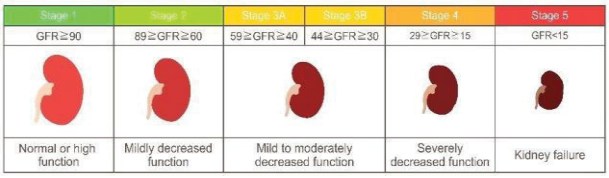


Fig. 1. Five stages of kidney disease (adapted from [5])

## II. RELATED WORKS

The acceptance of AI in Medicine has been rapid and a large number of projects have taken off especially in oncology and nephrology. Machine learning (ML) models have been increasingly applied in a plethora of clinical tasks, ranging from early health monitoring, cancer detection, to optimizing treatment strategies and forecasting adverse outcomes (e.g. Acute Kidney Injury [AKI] and Chronic Kidney Disease [CKD]).

So far, ML has been applied to Oncology for survival analysis, predicting the early response to therapy. Clinical notes, histopathology, genomics and other data types, must be integrated frequently for clinical decision making to be effective. In nephrology, for instance, AI-based predictive models have been developed to forecast dialysis requirements and anticipate early worsening of kidney function (as represented in Fig. 1), with the aim to personalize interventions. [4] Previous studies, however, have been organ

specific, and have almost exclusively relied on a single method to detect kidney disease or exposure, which has created a dearth in our knowledge of the integrative feature nature onco-nephrology.

Predictive ability alone tends to be the key (and perhaps only) element of concern for a ML model when trained on these types of datasets, with little to no thought as to how clinically interpretable or explainable a model may or may not be. Further, other predictive approaches are tested on synthetic or benchmark datasets, and therefore they might not be applicable to real-life, heterogenous clinical data limitation consisting on not being grounded on clinical data. [5]

In contrast in this study, we build on top of the state-of-the-art state and apply (idiomatically in vivo on cancer patients with renal risk) real world clinical data for the first time, and directly employ a variety of multiple supervised learning models (Random Forests, XGBoost, SVM etc.), combined with SHAP as well novel limited local explanations. Such explainability techniques allow for both global and local interpretability of model predictions, which in turn, helps clinicians have more confidence in the outcomes. In this paper, we also aim for high predictive accuracy but also interpretability and design of ethical AI, arguably a missing link in prior work. At the nexus of oncology and nephrology, this study not only addresses a major modern research gap but also adds to an emerging field of onco-nephrology percolated by AI (Artificial Intelligence) and emphasizes the importance of transparency and clinical applicability.

Despite recent advancements in AI-based diagnostic and prognostic systems, many initiatives remain limited to algorithmic optimization or experimental demonstration, with minimal implementation in clinical settings. Previous models for predicting renal outcomes in oncology frequently failed to integrate clinical, biochemical, and therapeutic data, resulting in incomplete identification of factors contributing to nephrotoxicity. Although deep neural networks and ensemble learning methods have improved AKIs prediction, their lack of interpretability has hindered regulatory approval and clinician confidence, both essential for integration into critical care. Most existing studies are retrospective and rely on curated datasets, which do not capture the variability and missing data typical of real-world Electronic Health Records (EHRs) systems. Addressing this limitation requires the development of interpretable and explainable AI systems that function effectively within the complex and heterogeneous environment of hospital data.

In the field of onco-nephrology, few studies have modeled kidney function decline as a dynamic process influenced by chemotherapy, comorbidities, and biochemical changes. Most research addresses only AKIs following treatment, neglecting the gradual and subclinical progression of renal impairment that precedes kidney failure. Recent attempts to integrate oncology-specific biomarkers such as neutrophil gelatinase-associated lipocalin (NGAL) and cystatin C with machine learning are increasing, but rarely incorporate explainable models or comparative analyses within a single experimental framework. This study aims to move beyond predictive accuracy by demonstrating that transparent and interpretable AI can facilitate collaboration between nephrology and oncology, supporting precision medicine where algorithmic analysis augments clinical judgment.

## III. Methodology

This study used a quantitative data-driven strategy to develop predictive models for the early identification of MPs-induced RI in cancer patients, using real-world clinical data. This was retrospectively external data collection done as a complement to the ONCONEFRO dataset in partnership with the University of Tirana and the University Hospital Center "Mother Teresa". It included an anonymized set of clinical and biochemical data for more than 5000 oncologic patient cases in were patients, treated and reposed modalities, laboratory values, sex-ratio, age and comorbid databases. Due to ethical and privacy constraints, the dataset is not publicly available and is used solely for academic research under institutional approval.

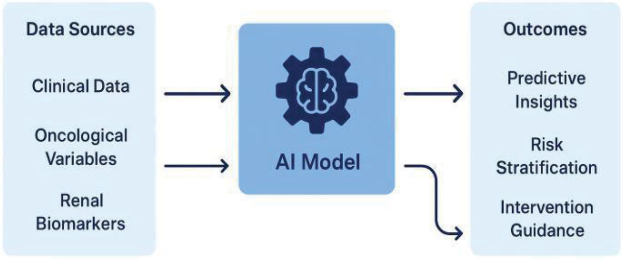


Fig. 2. Conceptional AI model for assessing renal impairment in oncology (adapted from [5])

### A. Dataset Description

This study uses the complete ONCONEFRO dataset, consisting of 5,000 anonymized patient records collected retrospectively at the University Hospital Center "Mother Teresa" (Tirana, Albania) in collaboration with the University of Tirana. The dataset contains structured clinical information routinely recorded during oncology and nephrology practice over the data collection period (2019–2024).

*Data Composition and Variables*

The dataset includes 53 variables covering multiple clinical dimensions, including:

- Demographics: sex, age (derived from date of birth), residence, profession.
- Oncology-related variables: cancer type (solid vs hematologic), diagnosis timing (time since diagnosis), treatment regimen, treatment line, number of chemotherapy cycles.
- Biochemical / laboratory variables: serum creatinine, urea, estimated glomerular filtration rate (eGFR/GFR), potassium, sodium, calcium, phosphorus, albumin, hemoglobin, white blood cells, platelets, and related hemogram values.
- Clinical/vital measurements: blood pressure, heart rate, weight, height, temperature, edema/ascites indicators.
- Comorbidities and risk factors: smoking status, prior kidney disease, nephrotoxic drug exposure, contrast exposure, and other comorbid conditions.

*Outcome Definition*

For supervised learning, renal impairment was operationalized using clinically accepted criteria, primarily based on eGFR/GFR < 60 mL/min/1.73m² and/or clinically significant creatinine elevation from baseline, consistent with nephrology practice.

*Ethics and Access*

All records were anonymized. Due to privacy and ethical constraints, ONCONEFRO is not publicly available and is used exclusively for academic research under institutional approval.

*B. Data Preprocessing*

The preprocessing process aimed to improve quality of data and to deal with the common problems in clinical datasets. The following steps were performed in an orderly manner:

*Preprocessing and Missing Data Imputations*

The incomplete case was mainly missed in laboratory examinations, including serum creatinine, eGFR and hemoglobin. Numerical features were replaced with the median, and categorical variables (e.g., diabetes, hypertension) with their most common category. Data lines with a large number of features that have been unknown or missing were removed to clear any noise and preserve the reliability of the dataset.

*Feature Normalization*

Continuous variables (e.g., biochemical markers, such as patient's age) were normalized to zero mean and unit variance by z-score standardization. This enabled heterogeneity factors to be directly comparable, and it also provided more stable model fitting.

*Categorical Encoding*

Nominal variables such as tumor type of transformation (one-hot encoding), binary clinical features (sex, hypertension, diabetes) were encoded as numbers. This change made the ordinal bias disappear, and it enabled algorithms to efficiently work with categorical data.

*Target Variable Balancing*

The dataset showed class imbalance with less positive cases of renal dysfunction. To address this, synthetic samples of the minority class were created using the Synthetic Minority Oversampling Technique (SMOTE). Because the renal impairment class was rare, SMOTE was applied only to the training set during model development to avoid data leakage and preserve an unbiased test distribution. Additional class-weight adjustments were also performed for some algorithms (XGBoost, Logistic Regression).

*Partitioning Strategy*

The preprocessed data was randomly partitioned into training (70%) and testing (30%) set with stratification, to maintain the same proportion of classes in test set. Furthermore, k-fold cross-validation was employed during training to prevent over-fitting and achieve robust performance estimates.

This clinical based preprocessing pipeline maintained data integrity, addressed imbalance, and optimized the data for the creation of predictive models.

After preprocessing, we conducted data validation and harmonization to ensure consistency between the ONCONEFRO dataset and the University Hospital Center “Mother Teresa” dataset. We reconciled laboratory reference ranges, coding conventions, and measurement units using a standardized transformation schema. Biochemical indicators were converted to common units based on international nephrology standards, and diagnostic codes were mapped to ICD-10 (the classification system deployed in the hospital information system during the data collection period) for interoperability. The unified dataset included longitudinal biochemical data, demographic variables, and comorbidity profiles, structured for machine learning analysis. Outliers were identified using interquartile range analysis. Values outside biological plausibility were corrected through domain-informed interpolation or excluded to prevent bias. This harmonization was critical for reliable model training and for minimizing inconsistencies from integrating data across institutions.

*C. Feature selection and target variable*

Feature selection and target variable Feature candidates were Selection of variables also depended on clinical significance, as well as statistical contribution for predictive performance. The raw data consisted of demographic, clinical, and laboratory data that was collected from electronic medical records. Missing data were imputed in a manner that preserved data structure, including mean substitution for continuous laboratory indices and mode replacement for categorical variables. Repetitive and excessively correlated features were removed to reduce multicollinearity and model variability.

The clinically relevant predictors of renal function were selected. These were serum creatinine, estimated glomerular filtration rate (eGFR), hemoglobin, potassium, and albumin; all commonly available measurements in cancer patients and strongly associated with renal outcomes. Comorbidities (such as hypertension and diabetes) and exposure to nephrotoxic agents (e.g.: cisplatin) were also included in the final model as all have been previously established to be factors involved with renal impairment. Age and type of tumor were included for baseline frailty and treatment-related renal toxicity patterns. Feature importance analysis using SHAP values was then employed to validate these selections and to determine their predictive importance.

The endpoint was renal insufficiency after (or during) cancer treatment. Renal dysfunction was defined per commonly accepted nephrology definitions with a decrease in eGFR to a clinically relevant level (< 60 mL/min/1.73 m²) or an increase in serum creatinine from preexisting baseline. Patients satisfying these criteria in the observational window were identified as “positive cases”, and patients who preserved the normal renal function as “negative cases”. The binary result was used to build a supervised classification problem that aimed at an early detection of patients at risk.

Alongside selecting features based on clinical knowledge, we used exploratory data analysis to find nonlinear relationships and interactions among predictors. We applied correlation \matrices and mutual information scores to remove variables that were redundant or contributed little. Univariate analyses showed that clinical parameters were statistically significant across different renal outcomes. This step-by-step process helped us choose variables that were both statistically sound and easy to interpret in a clinical context. We considered dimensionality reduction methods like Principal Component Analysis (PCA), but decided not to use them. We chose to keep variables interpretable rather than compress the data, since transparency is important for using predictive models in clinical settings.

In addition to predictive optimization, the feature selection process prioritized alignment with clinical knowledge and ethical explainability. By ethical explainability, we refer to the ability of AI models to provide transparent, human interpretable reasoning that supports clinical accountability and bias mitigation in medical decision-making. Domain expert input was integrated with quantitative validation using SHAP-based importance ranking to ensure both methodological rigor and clinical relevance. This combined approach enabled the identification of subtle but influential predictors, such as mild baseline anemia or slight albumin decline, which may precede renal impairment yet are frequently missed in standard assessments. Incorporating data on treatment-related nephrotoxic exposure enhanced the model's ability to assess risk dynamically. The resulting feature set represents the complex landscape of renal risk, including physiological, biochemical, and therapeutic factors, and provides a clinically interpretable basis for early detection and personalized care.

### *D. Modeling strategy*

A supervised machine learning model based on Random Forest, XGBoost, SVM and MLP algorithms was developed on the preprocessed data.[6] Models were assessed by performance measures such as accuracy, precision, recall, F1-score, and AUC-ROC, with attention given to clinical sensitivity and interpretability. Consistent with the recommendations for transparent reporting of a multivariable prediction model for individual prognosis or diagnosis, our study also used an AI-based explainability framework (explanation on XAI) (e.g., SHAP (SHapley Additive exPlanations) and LIME(Local Interpretable Model-Agnostic Explanations)) to provide interpretation on both global feature importance and local patient-level predictions.

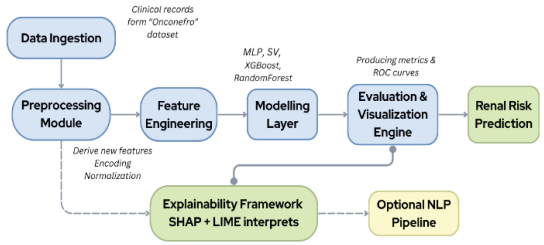


Fig. 3. System workflow (adapted from [5])

## IV. MODEL DEVELOPMENT

We applied several supervised ML models to predict the risk of RI in cancer patients to develop an AI driven Clinical Decision Support System (CDSS). The candidate models have been widely employed in clinical classification tasks considering their state-of-the-art performance, interpretability, capability to deal with high-dimensional data and practical feasibility in medical practice including Random Forest, XGBoost, Multilayer Perceptron (MLP) and Logistic Regression. Each algorithm performed training with the preprocessed dataset which was balanced by Smote method. To motivate the balancing strategy, we first report the class distribution observed after the train/test split. A significant observation from the dataset split was the pronounced class imbalance. The training set contained 3,890 patients with preserved renal function and 110 patients with renal impairment, while the test set consisted of 972 normal functioning cases and 28 impaired cases. This corresponds to approximately 2.75% of the full dataset being positive for renal impairment. Such imbalance is typical in clinical datasets but may bias machine learning models toward the majority class. To mitigate this effect, the Synthetic Minority Over-sampling Technique (SMOTE) was applied exclusively to the training set. This approach generated synthetic samples of the minority class to improve class representation while preserving the original test distribution for unbiased evaluation. Importantly, SMOTE was applied only to the training set to prevent data leakage and ensure that model evaluation reflects real-world class distribution. Demographic, biochemical, and treatment-related clinical variables included age, creatinine, eGFR, hemoglobin (Biochemistry), tumor type (Clinician related) (C) and (related to drugs received for treatment of other conditions or diseases) nephrotoxin exposure, comorbidities. We split the data set into two subsets using stratified sampling to keep class distribution of the two target classes (renal-impaired vs non-impaired). The model was trained and assessed using the scikit-learn, and XGBoost packages in Python. We performed hyperparameter tuning of the change point trough the grid search with cross-validation with the F1-score and ROC_AUC (which is an importance in imbalanced clinical classification problems). These are the metrics which we used to score each model on the test set. We employed explainable AI (XAI) paradigms to our predictive outputs which were not only correct but also interpretable. More concretely, we used SHAP (SHapley Additive exPlanations) to provide a global explanation of feature importance down into the model and also to attribute which clinical features contribute more frequently to predictions.

In addition, interpretable explanations for models were obtained from LIME (Local Interpretable Model-Agnostic Explanations), which fits a simple model around single predictions of the model. These interpretability techniques enabled physicians to understand how the model was arriving at its decision, and strengthened their confidence in the CDI system's ability to provide reliable suggestions when employed in an oncology context. Whether the system is also well behaved enough (robust model building and interpretability) to consider clinical deployment, as adjuvant supporting diagnostic tool for early renal risk assessment and personalized treatment planning in the field of cancer care.

## V. RESULTS AND COMPARATIVE ANALYSIS

In this study, five ML approach (XGBoosts, Random Forest, SVM, MLP and Logistic Regression) were compared to discover the optimal one for predicting the renal dysfunction in cancer patients through a performance analysis. The performance of each model was assessed based on classification, interpretability, robustness against class imbalance and feasibility in a clinical environment.

### *A. Performance Overview*

The best performance was implemented by XGBoost and we received the final result, being F1-score 0.82 and AUC 0.88 respectively, translated in quite good prediction in all range as well free. Its gradient booster property was particularly good at capturing non-linear relationships between clinical features and renal outcomes. Note that the Random Forest gives a slightly lower score, this behavior is reasonable with the tradeoff between predictive performance and model stability, which is a common consideration in interpretable machine learning models. SVM eventually becomes of good behavior in the balanced data, but is slightly inefficient when dealing with imbalances of classes, which is always the case in realistic

clinical data. Although MLP can represent complex relationships, it often required heavy hyperparameter tuning and was extremely prone to overfitting. Although logistic regression was simple and easy to understand, it showed limited predictive power resulting from its insufficient modeling capability for intricate relationships.

### B. Explainability and Clinical Integration

Interpretability is critical for Medical expert applications. Furthermore, the tree based models like XGBoost and Random Forest conveniently own the feature importance due to its design so they are intrinsically interpretable, aligning with that of new advancements of explainable AI (XAI) techniques) namely SHAP and LIME that make them sound with potential on its use for clinical translational applications. To demonstrate global interpretability, we report SHAP summary plot (Fig. 4) for the XGBoost classifier showing how features influence prediction across the full dataset.

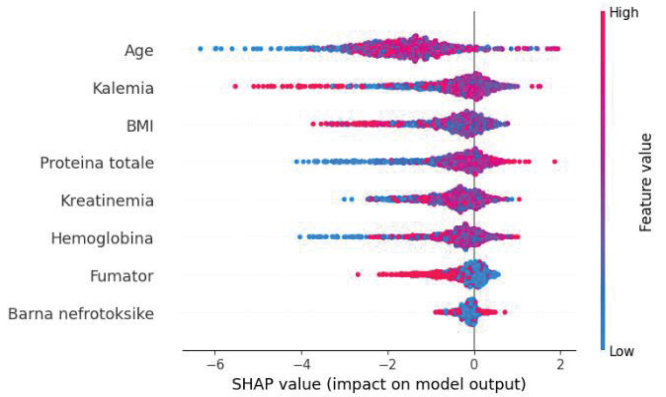


Fig. 4. SHAP summary plot (global interpretability) for the XGBoost classifier, showing feature impact on renal impairment predictions across the dataset (generated by the authors using the ONCONEFRO dataset)

These models can help us understand what kinds of contribution affect the confidence with which predictions are being made and thus help us explain how different features contributed when some medical prediction was made correctly. [7][8] For patient-level explainability, we show a SHAP waterfall plot (Fig. 5) for a representative case, detailing how each feature increases or decreases the predicted risk.

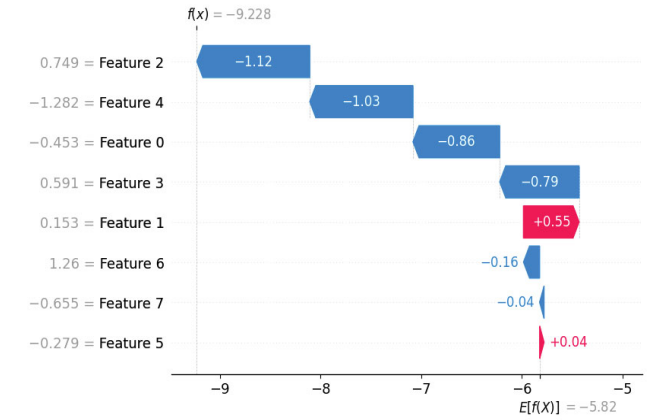


Fig. 5. SHAP waterfall plot (local explainability) for a representative patient, illustrating individual feature contributions to the predicted renal impairment risk (generated by the authors using the ONCONEFRO dataset)

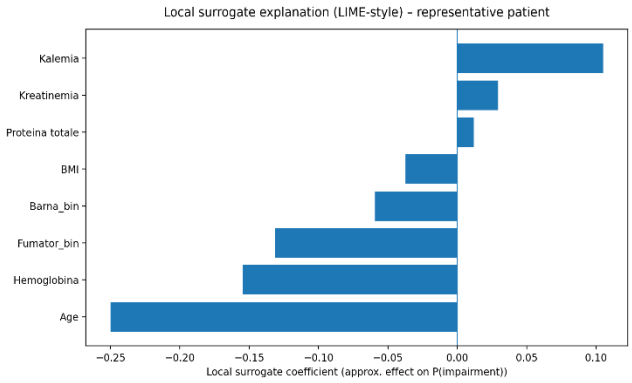


Fig. 6. LIME-style local surrogate explanation for a representative patient prediction, showing the most influential features in the local neighborhood (generated by the authors using the ONCONEFRO dataset)

A more complex model like MLP and SVM are relatively more complex to be explained due to their denser internals and limited model-agnostic interpretability. Logistic Regression is still pretty reasonable to interpret but it doesn't predict well or generalize at all.

### C. Summary Table

TABLE I. PERFORMANCE OF MACHINE LEARNING MODELS

| Model | Accuracy | Recall | F1-Score | AUC | $R^2$ | Interpretability | Clinical Integration Potential |
|---|---|---|---|---|---|---|---|
| XGBoost | 95% | 52% | 53% | 0.63 | - | High (SHAP-compatible) | Very High |
| Random Forest | 95% | 52% | 52% | 0.68 | - | Moderate | High |
| Neural Network | 90% | 56% | 52% | 0.61 | - | Low | Moderate |
| SVM (RBF Kernel) | 76% | 58% | 47% | 0.58 | - | Low | Moderate |
| XGBoost Regression | GFR Regression | - | - | - | -0.035 | Low | Low |
| XGB+Time feature | GFR Regression | - | - | - | 0.002 | Low | Very Low |

Explainability scores represent qualitative assessments based on model transparency and compatibility with SHAP/LIME rather than a quantitative metric. The time feature used in the models, was derived as the time elapsed since cancer diagnosis, computed from the recorded diagnosis date and follow-up statement.

### D. Results

Comparing machine learning models, XGBoost emerges as potentially the best approach to predict renal dysfunction in cancer patients. At 95% accuracy and an F1-score of 53%, XGBoost was trading precision and recall well. Even if its sensitivity (recall = 52%) was modest, its compatibility with SHAP analysis allowed a very good interpretability of the decision-making process, as the clinicians could see which clinical features (creatinine, eGFR and tumor type, for instance) were the most critical in the generated predictions. Such a transparency renders XGBoost especially appropriate to be adopted by clinical decision support systems, where explainability is a fundamental requirement for implementation.

Random Forest performed as well as XGBoost with an accuracy of 95% again and relatively higher number of AUC at 0.68. It was less interpretable, but the recall was 52%, meaning it failed to recognize a substantial proportion of high risk patients. Its clinical applicability potential was rated lower than for XGBoost despite the model remaining a good candidate.

The sensitivity of neural networks was slightly higher (recall = 56%) than RF, which is more desirable in detection cases for patients at risk of developing renal impairment. Yet, this benefit was achieved at the cost of exposure: accuracy fell to 90% and AUC to 0.61. These strengths and weaknesses, together with the limited interpretability of neural networks, impede their clinical adoption as a primary-ease prediction tool or offer them a subsidiary position.

The model SVM (with RBF kernel) was given the highest recall (58%), which was the proportion of the actual impaired subjects that were classified as impaired. Despite its relatively low precision (76%) and F1-score (47%) indicating a high false positive rate, its AUC (0.58) implied a performance of slight better than random guessing. Therefore, although sensitive, the model did not have adequate accuracy and stability for clinical utilization.

Efforts to predict kidney functional deterioration by the means of regressive approaches (XGBoost Regression and XGB+Time Feature) failed. Both the models even expressed negative or near zero $R^2$ values i.e. very less predictive power. Lack of their interpretability and clinic integration potential also confirmed their inappropriateness for the situation.

In conclusion, the findings place XGBoost as the most pragmatic model that could reach a compromise between prediction accuracy and interpretability. Neural networks and SVMs, albeit having an increased sensitivity, suffer from reliability problems that prevent their use in real-world application.

## VI. Discussion and clinical implications

The performance of machine learning-based CDSS, which offers reliable and highly accurate pre-symptomatic prediction for renal impairment risk among cancer patients, has important clinical implications. The system helps in timely intervention early identification of individuals based on structured clinical and biochemical data to identify at risk people and fine-tune treatment. A point particularly critical in oncology, given the potential to compromise treatment continuance and effectiveness through nephrotoxicity.

Incorporation of explainable AI (XAI) techniques like SHAP and LIME in the model provides interpretability and transparency to some extent, which is necessary for clinical applicability. It is possible to use these tools to run any given prediction for an individual back through the model predict and get out the original clinical indicators that led each specific prediction down a specific path, aligning outputs with medical reasoning and increasing trust in automated recommendations.

Critically, it also demonstrates the ability to integrate the CDSS in EHRs systems. The modular system architecture, structured data format independent and explainable outputs makes the model an operational and scalable solution for real-time renal risk monitoring in cancer care units. This would allow continuous risk assessment which could replace mainly delayed laboratory or manual risk scoring tools.

Nonetheless, some limitations have to be noted. This study was carried out using data from a single medical center to both train and validate the model, which may limit its generalizability to wider or more diverse patient populations. The models may perform differently due to differences in demographics, clinical protocols, or laboratory practices across institutions. Furthermore, although simulated data was used to augment the dataset to account for class imbalance, external validation on multi-center datasets is necessary to validate generalization ability and clinical safety prior to live use.

While we would not say that these limitations cripple the power of this work, it instead serves to offer a platform for future investigations and system polishing. Our results underscore the utility of AI models to support decision making in onco-nephrology and suggest that further evaluation is warranted via multicenter partnerships, with a focus on scaling, validating, and fine-tuning this tool for broader clinical use. Compared to traditional threshold based renal risk assessment using creatinine and eGFR alone, the proposed AI models achieve improved risk stratification by incorporating multidimensional clinical information, although at the cost of moderate predictive performance.

In addition to forecasting outcomes, the system provides practical support for oncology teams. By continuously reviewing patient data such as lab results, demographics, and treatments, it can alert clinicians early to potential kidney issues that might otherwise go unnoticed. The model can also help optimize therapy by suggesting dose adjustments, hydration strategies, or kidney protective measures when risk increases. Looking ahead, this clinical decision support system could integrate with chemotherapy scheduling and infusion monitoring, enabling real-time adjustments based on each patient's kidney health. This data driven approach aims to prevent treatment delays, reduce hospital stays, and improve patients' quality of life, supporting both personalized medicine and cost-effective care.

The framework also paves the way for research and collaboration. Because the system's design is easy to interpret, researchers can spot patterns and risk factors across different cancer groups, inspiring new ideas about kidney side effects or possible biomarkers. If adopted across hospitals and research networks, the decision support system could become a learning tool that improves over time by sharing knowledge—without risking patient privacy. This kind of growth supports the goal of building an AI-powered kidney safety registry, which would help create evidence-based guidelines and policies for cancer care. Ultimately, this approach highlights how AI can do more than assist with diagnosis; it can help bridge the gap between clinical intuition and insights from large scale data.

## VII. Conclusion and Future Works

*Conclusion*

In this paper an experimental evaluation of the use of machine learning for the prediction of renal failure in cancer patients using both classification and regression models is presented. In the range of models which we have assessed, XGBoost proved to provide a good deal between accuracy and interpretability, and between clinical integration potential. SHAP enabled transparent feature importance analysis, which is specifically important in the area of real-world medical decision support. On the other hand, though Random Forest also obtained similar accuracy, it did not have similar interpretability with

Sparse Binary Logistic Regression, and Neural Networks and SVMs demonstrated higher sensitivity but lower specificity. The regression models were not able to make useful predictions of continuous GFR decline indicating the complexity of modelling kidney function trajectories.

The results of this study indicate that interpretable classification methods are potentially valuable in improving the patient management decisions in oncology and nephrology. Through alerting high risk patients for renal dysfunction, the ICx-based systems may help oncologists and nephrologists with the implementation of early interventions that will allow enhanced patient safety and optimized treatment outcome.

Beyond evaluating model performance, this research demonstrates how artificial intelligence can support risk prediction in oncology by aligning algorithms with clinical reasoning. By comparing tree-based, regression, and deep learning models, we clarify the trade-offs among accuracy, interpretability, and computational efficiency that affect real-world use. The results show that a balanced approach, such as XGBoost with explainability, can build clinician trust and speed adoption compared to black-box models. Identifying feature importance with SHAP also enhances transparency and can help generate clinical hypotheses by highlighting key biomarkers or therapy-related factors linked to renal impairment. This integration of predictive analytics and clinical expertise marks a significant step forward for AI-driven decision support by connecting quantitative models with medical knowledge.

The methodological framework established in this study, consists of data preprocessing, supervised learning, class balancing, model evaluation, and explainability analysis using SHAP and LIME. It provides a basis for both translational scalability and regulatory compliance. The use of a structured, interoperable data schema and explainable outputs enables integration with EHRs systems and supports alignment with the European Union Artificial Intelligence Act and Medical Device Regulation (MDR) standards for clinical artificial intelligence tools. These compliance-oriented design principles are critical for maintaining data security, model accountability, and traceability within hospital settings. The flexible nature of this approach also allows for adaptation to additional risk domains, such as cardiotoxicity or hepatotoxicity, through retraining on relevant biomarkers. This flexibility demonstrates the versatility and long term applicability of interpretable machine learning systems as clinical infrastructures increasingly adopt data driven precision medicine.

*Future Works*

As positive as these outcomes are, many areas are still open problems for future research and development:

- Expansion of Datasets

The present study was based on single center data. Further study with broader population and multi center data will be needed to enhance the generalization of the model.

- Temporal and Longitudinal Modeling

Integration of time series data, as trends in creatinine or eGFR over treatment cycles, may allow better prediction of renal decline. One could investigate using recurrent neural network or transformer-based architecture for this.

- Genomic and Imaging Integration

Integration of molecular platforms and radiological imaging with structured clinical datasets can increase the predictive potential and the mechanistic understanding of multifactorial factors underlying renal dysfunction.

- Real-World Clinical Validation

Prospective clinical validation of a system is important to determine its reliability, usability, and clinical relevance. oncology center pilot implementations could test model relevance in actual practice.

- Model Interpretability Enhancements

SHAP enjoys a good interpretability power for tree-based models, this work leaves room for the future development of hybrid methods that combine causal inference with ML for generating explanations that better adhere to medical reasoning.

- Design of real-time CDSS

Long term, the predictive model may be integrated within EHRs platforms to offer real-time risk estimation, decision support to oncologists and nephrologists and facilitate clinical uptake.

In summary, this study is an advancing step forward to the application of AI in personalized care for oncology, which shows that interpretable ML models, especially XGBoost, can offer interpretable and clinically applicable assistance in early prediction of RI risk. It will be critical for the next steps along these avenues to transform this science into a useful and effective clinical modality that can improve the care of the patient.

## REFERENCES


[1] Hanai, K., Takahashi, H., Uojima, H., et al. (2020). Risk of AKI in patients receiving chemotherapy with cisplatin: a systematic review and meta-analysis. Cancer Chemotherapy and Pharmacology, 85(4), 695–706. https://doi.org/10.1007/s00280-020-04025-2

[2] KDIGO Clinical Practice Guidelines (2012). Clinical Practice Guidelines for AKI. Kidney International Supplements, 2(1), 1–138. https://doi.org/10.1038/kisup.2012.1

[3] Y. Wang, Z. He and D. Wang, "Predicting treatment response in cancer using AI-enhanced multimodal models," Scientific Reports, vol. 14, p.

[4] Article 77638, 2024

[5] Zhou, X., Yang, G., Liu, Q., & Wang, J. (2022). Artificial Intelligence in Nephrology: Current Status and Future Directions. Frontiers in Medicine, 9, 837861. https://doi.org/10.3389/fmed.2022.837861

[6] E. Xhumari and E. Dollaj, "Review of AI Applications in Oncology: Current Tools, Challenges, and Opportunities for Clinical Decision

[7]

[8] Support," in 6th International Conference on Recent Trends in Applications of Computer Science and Information Technology (RTA-CSIT), Tirane, Albania, 2025.

[9] Sidey-Gibbons, J. A. M., & Sidey-Gibbons, C. J. (2019). Machine learning in medicine: a practical introduction. BMC Medical Research Methodology, 19(1), 64. https://doi.org/10.1186/s12874-019-0681-4

[10] Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). https://doi.org/10.1145/2939672.2939785

[11] A. Rajkomar, J. Dean and I. S. Kohane, "Machine learning in medicine," New England Journal of Medicine, vol. 380, no. 14, p. 1347–1359, 2019.

[12] S. M. McKinney, M. Sieniek, V. Godbole, J. Godwin, N. Antropova, H. Ashrafian, T. Back, M. Chesus, G. S. Corrado, A. Darzi and J. Dean, "International evaluation of an AI system for breast cancer screening," Nature, vol. 577, no. 7788, p. 89–94, 2020