Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

AI DECISION SUPPORT FOR RENAL DOSE ADJUSTMENT USING MEDICATION ORDERS AND LABORATORY TRENDS

Juan Perez[1], Ana Gutierrez[1]*, Carlos Lopez[2]

1. ***Department of Intelligent Pharmaceutical Analytics, Faculty of Pharmacy, National Autonomous University of Mexico, Mexico City, Mexico.***
2. ***Department of Computational Drug Sciences, Faculty of Engineering, Monterrey Institute of Technology, Monterrey, Mexico.***

ARTICLE INFO

Received:
29 October 2025
Received in revised form:
16 January 2026
Accepted:
18 January 2026
Available online:
28 February 2026

Keywords: Artificial intelligence, Renal dose adjustment, Clinical decision support, Acute kidney injury, Electronic health records, Medication safety

ABSTRACT

Renal impairment is highly dynamic in hospitalized patients, and many medications require dose adjustment to avoid toxicity or therapeutic failure. Conventional renal dose alerts often depend on static thresholds and may not reflect evolving laboratory trajectories. Existing renal dose decision support tools commonly rely on a single creatinine-derived estimate of kidney function. When alerts are broad, repetitive, or poorly contextualized, clinicians may override them despite potential medication safety risk. This article proposes an AI decision support framework that continuously tracks medication orders and serial renal laboratory results. The system would forecast renal function changes and generate patient-specific dose adjustment recommendations at prescribing or pharmacist review. The proposed framework includes real-time data ingestion from electronic health records, a renal function prediction module, a drug-specific dosing rule engine, an alert prioritization layer, and an embedded CPOE interface. Pharmacist verification is incorporated as a human-in-the-loop safeguard. The system would aim to distinguish stable renal impairment from rapidly declining renal function and suppress alerts when the prescribed regimen is already appropriate. Its recommendations would be accompanied by concise explanations linking recent laboratory trends, medication risk, and renal dosing logic. An AI-augmented renal dosing framework could make renal dose adjustment more timely, individualized, and clinically acceptable. Such a system should be implemented only after careful workflow integration, safety review, and prospective clinical evaluation.

To Cite This Article: Perez J, Gutierrez A, Lopez C. AI Decision Support for Renal Dose Adjustment Using Medication Orders and Laboratory Trends. Pharmacophore. 2026;17(1):22-32. https://doi.org/10.51847/NPcrJjnF9S

Introduction

Inappropriate medication dosing in renal impairment remains a persistent medication safety problem because renal clearance affects exposure, toxicity risk, and therapeutic effectiveness for many commonly prescribed drugs. Studies of renal medication-related decision support have shown that alerts may identify important prescribing discrepancies, yet override behavior remains a major concern when recommendations are poorly targeted or poorly timed [1]. Dose adjustment interventions in hospitalized patients with reduced kidney function have demonstrated the operational appeal of computerized checking, but they also show that system design and clinician response strongly influence whether prescribing behavior changes [2]. Ambulatory and inpatient renal dosing systems therefore need to address both pharmacologic appropriateness and the human factors that determine whether clinicians act on the recommendation [3].

Traditional renal dosing support is often implemented as a rule-based CDSS linked to CPOE, where the ordered medication is compared with a creatinine-based estimate of renal function. Such systems can be useful, but evaluations of renal dose CDSS have shown that alerts based on fixed thresholds may fail to improve downstream clinical indicators when the patient's trajectory and clinical context are not adequately represented [2]. Prescription checking systems for reduced glomerular filtration rate have also highlighted the value of identifying impaired kidney function before medication harm occurs, while underscoring the need for precise, actionable recommendations [4]. These findings suggest that static renal dose rules should be augmented by dynamic data interpretation rather than treated as complete decision-making systems.

AI methods create an opportunity to move renal dose adjustment from reactive threshold checking toward anticipatory decision support. Machine-learning models have been developed to predict acute kidney injury from EHR-derived data, including

Corresponding Author: Ana Gutierrez; Department of Intelligent Pharmaceutical Analytics, Faculty of Pharmacy, National Autonomous University of Mexico, Mexico City, Mexico. E-mail: ana.gutierrez@gmail.com.

laboratory values and clinical context, and these models illustrate how temporal signals could support earlier recognition of renal deterioration [5]. Continuous prediction approaches have further shown that serial clinical data can be transformed into clinically oriented risk forecasts rather than isolated laboratory interpretations [6]. In a renal dosing framework, similar predictive logic could be coupled with medication order parsing so that the system evaluates not only current kidney function but also the likelihood that renal function will decline during the active medication course.

This AIF article proposes an AI-powered decision support system for renal dose adjustment that integrates medication order data with laboratory trends to generate individualized, real-time recommendations. The framework draws on renal dosing CDSS experience, medication-related alert studies, and AI-based kidney function prediction work to define a conceptual architecture rather than an experimental validation study [7]. It also recognizes that alert fatigue can undermine even clinically accurate recommendations, making alert prioritization and workflow embedding central design requirements. The goal is a system that would support prescribers and pharmacists by recommending renal dose changes only when the patient's medication exposure, renal trajectory, and clinical context suggest a meaningful discrepancy.

Background

Drug Dosing in Renal Impairment

Renal impairment changes drug clearance, half-life, metabolite accumulation, and toxicity risk, which is why many medications require dose reduction, interval extension, avoidance, or therapeutic drug monitoring when kidney function declines. Clinical renal dosing recommendations are often encoded from drug-specific monographs and institutional protocols, but implementation studies show that translating these rules into electronic systems requires reliable identification of kidney function category and prescribed regimen [8]. Online drug information databases and renal dosing references may differ in their recommendations, so a computable dosing engine should preserve drug-specific logic, local formulary governance, and pharmacist review rather than simply applying a generic dose table [9]. For high-risk drugs, including nephrotoxic agents and medications with narrow therapeutic ranges, the system should make conservative, explainable recommendations that reflect both current impairment and expected renal trajectory [10].

Current Renal Dose CDSS and Their Limitations

Current renal dose CDSS typically compare the active medication order with a renal function estimate and generate an alert when the order appears inconsistent with an encoded dosing rule. Evaluations of inpatient renal medication alerts show that commercial EHR systems can produce frequent alerts, but the clinical value of these alerts depends on specificity, timing, and the ability to distinguish meaningful risk from routine exceptions [1]. Dose adjustment systems in renal failure have shown potential to improve prescribing processes, yet their effectiveness may be limited when alerts are not well integrated into prescriber and pharmacist workflows [8]. More broadly, medication-related CDSS research shows that excessive or low-value warnings can normalize overrides, turning safety infrastructure into background noise [11].

Predictive Models of Renal Function

Machine-learning models for acute kidney injury prediction demonstrate how EHR data can be used to detect patterns preceding renal deterioration. Inpatient AKI prediction work has used structured clinical data to support early identification of patients at risk for kidney injury, while emergency department models have shown how prediction can be initiated early in the care pathway [5, 12]. Continuous and time-updated AKI prediction approaches are especially relevant to renal dosing because medication exposure often spans periods during which creatinine and eGFR are changing rather than stable [13, 14]. Deep learning and convolutional approaches further suggest that temporal laboratory sequences could support short-term renal trajectory estimation, although such predictions should be used as decision support rather than autonomous prescribing authority [15, 16].

AI in Medication Safety and Clinical Pharmacy

AI in medication safety can support detection of medication-associated harm, risk stratification, and prioritization of pharmacist attention. Machine-learning approaches have been applied to identify medication-associated acute kidney injury, indicating that medication histories and clinical features can be combined to flag patients whose renal risk may be drug-related [17]. Nephrotoxin-focused studies also suggest that AI models could help identify patients whose drug exposures and clinical context make kidney injury more likely, which is directly relevant to renal dose adjustment and avoidance recommendations [18]. In clinical pharmacy, such models should operate alongside rule-based dosing calculations because learned patterns can suggest risk, while explicit dosing rules remain necessary for transparent dose selection.

Human Factors in Alert Design

Human factors are central to renal dose decision support because clinicians must rapidly decide whether an alert is relevant, safe to follow, or clinically inappropriate. National and institutional studies of medication-related CDSS show that the tradeoff between safety and alert fatigue requires careful attention to alert frequency, severity, and clinical usefulness. Alert appropriateness research emphasizes that decision support should fit workflow and reduce unnecessary interruption rather than simply increase the number of warnings. A renal dosing system should therefore provide concise rationale, allow structured override reasons, and prioritize alerts that are clinically actionable for the prescriber or pharmacist at that moment.

Table 1 shows key human-factor considerations for effective renal dose decision support systems.

Table 1. Human-Factor Principles for Renal Dose Decision Support Systems

Human-Factor Consideration	Description	Implementation in CDSS
Alert Relevance	Clinicians need to determine quickly if an alert applies to the patient's condition	Ensure alerts are tailored to patient-specific parameters (e.g., renal function, comorbidities)
Safety vs. Alert Fatigue	Balance between preventing errors and avoiding excessive interruptions	Prioritize high-severity alerts; reduce low-risk warnings
Clinical Usefulness	Alerts should provide actionable guidance rather than generic warnings	Include concise rationale and clear dosing recommendations
Workflow Integration	Decision support must align with clinical workflow	Design alerts to minimize unnecessary interruptions and fit naturally into tasks
Structured Override Options	Clinicians should document reasons when ignoring alerts	Provide structured override reasons to capture clinical judgment
Actionability	Alerts must be relevant to immediate clinical decisions	Focus on alerts that prescribers or pharmacists can act on at that moment

System Architecture Overview

High-Level Architecture

The proposed system would subscribe to EHR event streams for new medication orders, new laboratory results, patient demographic updates, renal replacement therapy status, and relevant medication discontinuations. A predictive engine would estimate the near-term direction of kidney function by interpreting serial serum creatinine, eGFR, urine output when available, and exposure to nephrotoxic medications, reflecting the broader movement from one-time AKI detection toward continuous renal surveillance [13]. A dosing engine would compare each active or proposed medication order with renal-adjusted recommendations, drawing on the same kind of prescription checking logic used in prior renal impairment systems while adding trajectory-aware interpretation [4]. An alert would be generated only when the predicted renal state and the active order produce a clinically meaningful mismatch that warrants prescriber or pharmacist attention.

Figure 1 illustrates the proposed renal trajectory-to-dose recommendation architecture, showing how medication orders, serial renal laboratory trends, AI forecasting, drug-specific dosing rules, alert prioritization, and pharmacist verification are integrated into CPOE-based renal dose decision support.

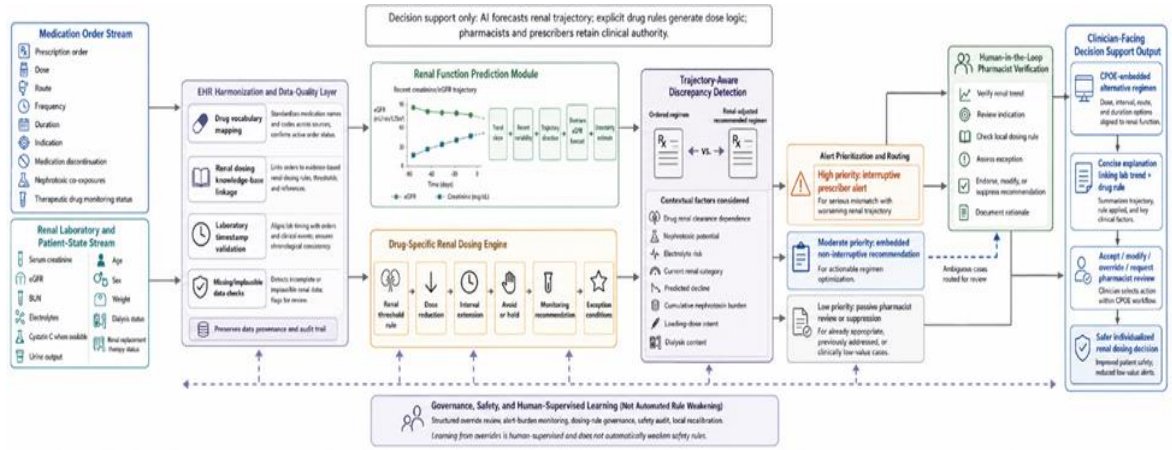

Figure 1. AI Decision Support Architecture for Renal Dose Adjustment Using Medication Orders and Laboratory Trends

Core Data Inputs and Outputs

The system's core inputs would include medication name, dose, route, frequency, duration, indication when available, serum creatinine, calculated eGFR, blood urea nitrogen, electrolytes, cystatin C where implemented, urine output, age, sex, weight, and dialysis status. Prior studies of AKI prediction show that structured EHR variables can support renal risk estimation, while nephrotoxin surveillance research indicates that medication exposure is an essential part of kidney safety assessment [19, 20]. The output would be a recommendation such as dose reduction, interval extension, temporary hold, avoidance, or pharmacist review, with an explanation that links the recommendation to the patient's recent renal trend and the drug's renal dosing rule. For safety, the system would present the recommendation as decision support rather than an automatic medication modification.

Table 2 defines the functional logic of the proposed AI-augmented renal dose adjustment system by linking each architecture layer to its required data elements, analytical contribution, clinical output, and governance safeguard.

Table 2. Functional Logic of the Proposed AI-Augmented Renal Dose Adjustment System

System layer	Primary function	Key data elements	Analytical logic added by the framework	Clinical output	Safety or governance requirement
EHR medication-order stream	Detect active and proposed medication regimens requiring renal review	Medication name, dose, route, frequency, duration, order start/stop time, dose changes, indication when available	Converts medication orders into computable exposure objects that can be compared with renal dosing rules	Active medication list requiring renal dosing assessment	Accurate order-status reconciliation is required to avoid alerts for discontinued, duplicate, or inactive therapies
Renal laboratory and patient-state stream	Represent kidney function as a changing clinical state rather than a single value	Serum creatinine, eGFR, BUN, electrolytes, cystatin C where available, urine output, weight, age, dialysis status, renal replacement therapy	Uses serial timing, slope, variability, and direction of renal change to distinguish stable impairment from acute decline	Current renal category plus recent renal trajectory	Data-quality checks should flag missing, delayed, implausible, or dialysis-confounded laboratory values
Harmonization and mapping layer	Standardize medication and laboratory data for decision-support use	Drug vocabulary mapping, local formulary code, renal dosing knowledge-base link, laboratory timestamps, provenance metadata	Aligns medication identity, renal rule source, and laboratory timing so that downstream recommendations are traceable	Standardized patient-medication-renal state profile	Local pharmacy governance must reconcile differences across renal dosing references before rules are encoded
Renal function prediction module	Estimate short-term direction of renal function	Serial creatinine/eGFR, recent medication exposures, nephrotoxin burden, care setting, comorbidities, urine output where available	Forecasts likely renal deterioration, stability, or recovery over the relevant prescribing window	Short-term renal trajectory forecast with uncertainty	Forecasts should support, not replace, explicit dosing rules and pharmacist judgment
Drug-specific renal dosing engine	Translate renal state into medication-specific dosing logic	Drug renal clearance dependence, renal thresholds, dose-reduction rules, interval-extension rules, avoidance criteria, monitoring needs	Applies explicit, auditable renal adjustment logic to the current or forecasted renal category	Dose reduction, interval extension, avoidance, hold, monitoring, or pharmacist review recommendation	Rules must be version-controlled, locally approved, and reviewed for high-risk medications
Discrepancy detection module	Determine whether the ordered regimen differs meaningfully from the recommended regimen	Ordered dose/frequency, recommended renal-adjusted regimen, predicted renal trajectory, drug risk profile, electrolyte abnormalities	Distinguishes clinically meaningful mismatches from acceptable exceptions or low-value discrepancies	Discrepancy status and preliminary alert priority	Thresholds should be validated against expert pharmacist review rather than alert volume alone
Context-aware alert prioritization layer	Route recommendations according to urgency, harm potential, and workflow relevance	Medication risk, renal trajectory, loading-dose intent, dialysis context, pharmacist documentation, prior override reason	Converts dosing discrepancies into tiered alert actions rather than uniform interruptive warnings	Interruptive alert, embedded recommendation, pharmacist queue item, passive documentation, or suppression	Suppression logic must be transparent, auditable, and reviewed for missed-risk patterns
Human-in-the-loop pharmacist verification	Provide clinical oversight for high-risk, ambiguous, or exception-prone recommendations	Renal trend, indication, dosing rule, TDM availability, exception conditions, prescriber documentation	Combines AI forecast, rule-based recommendation, and clinical judgment before escalation or modification	Endorsed, modified, deferred, or suppressed recommendation with rationale	Pharmacist review should be documented and used for governance, not treated as automatic model retraining
CPOE and pharmacy workflow interface	Present the safest actionable recommendation at the moment of prescribing or verification	Current order, recommended alternative regimen, rationale, acceptance pathway, override options	Embeds the recommendation into existing prescribing and verification tasks to reduce friction	Accept, modify, override, defer, or request pharmacist review	Interface design should minimize unnecessary clicks while preserving clinician authority

Override learning and governance layer	Monitor system performance and refine rule or alert logic	Structured override reasons, alert frequency, acceptance rate, pharmacist interventions, safety events, suppressed-alert review	Identifies poorly targeted alerts, clinically valid exceptions, and drug-rule combinations needing revision	Updated alert criteria, rule-base review priorities, safety audit reports	Override learning should be human-supervised and should not automatically weaken high-risk safety rules

Design Principles

The architecture should be real-time, actionable, explainable, and respectful of clinical workflow. Real-time design is important because renal function can shift during hospitalization, and AKI prediction studies have shown that time-updated models can provide clinically useful trajectory awareness before injury is fully apparent [14]. Actionability requires that the alert include a specific alternative regimen rather than a generic warning, reflecting lessons from renal dose CDSS implementations in which practical dose adjustment guidance is central to clinician acceptance [3]. Explainability requires showing the relevant laboratory trend, the dosing rule being applied, and the reason an alert is considered high, moderate, or low priority, thereby addressing concerns identified in alert fatigue and appropriateness studies.

Medication Order and Laboratory Data Integration

Real-Time Data Feeds from EHR

The system would consume structured medication orders and laboratory results through an EHR interface such as HL7 FHIR or a comparable institutional integration layer. Each new order, dose change, discontinuation, or laboratory result would trigger an update to the patient state, enabling the renal prediction and dosing modules to re-evaluate whether the current regimen remains appropriate. Real-time surveillance is conceptually aligned with automated continuous AKI prediction models, which demonstrate how repeated EHR updates can support ongoing renal risk assessment rather than a single static calculation [13]. Because implementation studies of renal dosing support show that alert value depends on integration with actual prescribing workflows, the data pipeline should prioritize low latency, accurate medication status, and clear provenance of laboratory inputs [1].

Harmonization of Drug Nomenclature and Renal Dosing Rules

Medication orders would be mapped to a standardized drug vocabulary and linked to a curated renal dosing knowledge base that specifies dose reduction, interval extension, avoidance, or monitoring recommendations. Existing studies of renal dosing databases show that recommendations may vary across information sources, so local governance should review and reconcile dosing logic before it is encoded into the system [9]. The dosing engine should also accommodate specialty contexts, such as anticancer drugs, where renal or hepatic dysfunction may require drug-specific adjustment logic and pharmacist oversight [21]. For each mapped medication, the system would store the relevant renal thresholds, adjustment strategy, exception conditions, and monitoring recommendations in a computable form suitable for CPOE integration.

Laboratory Trend Processing

Serial creatinine values would be processed as a time series rather than as isolated observations, with safeguards for implausible values, dialysis-related shifts, delayed laboratory reporting, and abrupt changes that require clinical interpretation. AKI prediction studies using EHR data support the relevance of temporal laboratory patterns, and deep learning approaches indicate that sequential representations can help identify renal deterioration before a single threshold is crossed [6, 15]. The system would compute current renal function using an institutionally approved equation and then characterize slope, variability, and direction of change so that stable chronic impairment can be distinguished from acute decline. This trend representation would feed the prediction model and would also be displayed to the clinician when it materially influences a dosing recommendation. **Table 3** shows key considerations for processing serial creatinine values in renal dose decision support systems.

Table 3. Processing and Interpretation of Serial Creatinine Values for Renal Decision Support

Consideration	Description	Implementation in CDSS
Time-Series Analysis	Treat serial creatinine measurements as sequential data rather than isolated points	Use temporal models to capture trends in renal function
Safeguards for Implausible Values	Prevent misinterpretation due to lab errors or data entry mistakes	Flag outliers and validate against physiological plausibility
Dialysis-Related Shifts	Recognize sudden changes caused by dialysis rather than renal deterioration	Adjust models to account for post-dialysis creatinine drops
Delayed Laboratory Reporting	Address potential lag between sample collection and result availability	Incorporate timestamp-aware processing to prevent outdated data from influencing decisions
Trend Characterization	Capture slope, variability, and direction of change	Distinguish stable chronic impairment from acute kidney injury

Predictive Modeling	Use sequential patterns to anticipate renal deterioration	Feed trend data into deep learning or time-series prediction models
Clinician Display	Ensure trends inform actionable dosing decisions	Show temporal representation when it materially affects recommendations

Ai-Driven Renal Function Prediction and Dose Calculation

Short-Term eGFR Forecasting Model

The forecasting module would use historical inpatient patterns to estimate the likely direction of eGFR over the next clinical decision window, drawing on serial laboratory values, medication exposures, comorbidities, and current care setting. Machine-learning studies in critically ill and hospitalized patients show that AKI risk can be inferred from routinely collected EHR features, including clinical context and laboratory trajectories [19, 22]. Pediatric and adult prediction work also suggests that renal risk models may need population-specific calibration before they are used to support clinical decisions [23, 24]. In this AIF framework, the model would not prescribe therapy independently; instead, it would provide a forecast and uncertainty estimate that the dosing engine and pharmacist can interpret alongside explicit renal dosing rules.

Drug-Specific Dosing Engine

For each active or proposed medication, the dosing engine would translate the forecasted renal function category into a recommended regimen using encoded renal adjustment rules. Prior renal CDSS studies show that dose adjustment recommendations are most useful when they are specific to the medication, route, and renal function category rather than expressed as generic warnings [2, 8]. The engine would also consider monitoring context, such as whether therapeutic drug monitoring is available or whether the medication is being used as a loading dose that may justify temporary deviation from maintenance dosing logic. Where the AI forecast suggests impending renal decline, the engine could recommend a safer interval or dose for pharmacist review while clearly indicating the rule and renal trend that drove the suggestion.

Discrepancy Detection and Risk Stratification

The discrepancy module would compare the ordered regimen with the renal-adjusted recommendation and determine whether the difference is clinically meaningful enough to trigger action. Medication-associated AKI research supports combining drug exposure with patient-specific clinical risk, so the system would weigh the drug's nephrotoxic potential, renal clearance dependence, current electrolyte abnormalities, and recent eGFR direction when assigning alert priority [10, 17]. Nephrotoxin exposure studies also suggest that cumulative medication burden should influence risk interpretation, particularly when several kidney-relevant agents are active at the same time [18, 20]. Instead of producing a blanket alert for every mismatch, the system would stratify recommendations into pharmacist review, prescriber-facing alert, passive documentation, or suppression when the order is already consistent with the renal dosing rule.

Alert Prioritization and Context-Aware Decision Support

Tiered Alert Strategy

A tiered alert strategy would classify renal dosing discrepancies by potential harm, urgency, and likelihood of clinical action. High-risk alerts would be interruptive only when the medication, renal trajectory, and ordered regimen together suggest a serious safety concern, while moderate-risk discrepancies could appear as non-interruptive recommendations embedded in the ordering workflow. Low-risk issues could be routed to passive review or documentation queues, reflecting evidence that excessive low-value medication alerts contribute to override normalization and alert fatigue [11]. This approach would preserve clinician attention for cases in which the renal dose recommendation is timely, specific, and plausibly consequential.

Contextual Factors Modifying Alert Delivery

The system should modify alert delivery based on clinical context, including dialysis status, loading-dose intent, palliative goals, temporary renal replacement therapy, recent pharmacist documentation, and structured override history. Medication-related CDSS studies show that overridden alerts may reflect clinically valid exceptions, so a renal dosing framework should not repeat the same recommendation without accounting for the reason it was dismissed [25]. A disease- and medication-aware alert model could further reduce fatigue by learning which alert patterns are repeatedly non-actionable while preserving escalation for severe or newly changed risk states. Contextual suppression should be transparent and auditable so that safety teams can distinguish appropriate filtering from missed medication risk.

Explanation and Clinical Rationale

Each alert should include a concise explanation that connects the medication order, current renal function, recent laboratory trend, forecasted trajectory, and applicable dosing rule. Alert appropriateness research emphasizes that decision support should support clinical workflow rather than interrupt it with vague or poorly justified warnings. A clear rationale would allow clinicians to understand whether the recommendation is driven by chronic renal impairment, rapid acute decline, interacting nephrotoxins, or an electrolyte abnormality that changes the medication's risk profile [10]. The interface should therefore

present the recommended alternative dose or interval, the reason for the recommendation, and a one-click pathway to accept, defer, or request pharmacist review.

Human-In-The-Loop Verification and Override
Pharmacist Review Queue
A pharmacist review queue would provide a safety layer between AI-generated recommendations and prescriber-facing interventions, especially for high-risk or ambiguous cases. Prior renal dosing CDSS evaluations indicate that computerized recommendations are most useful when they complement pharmacist judgment rather than replace clinical review [8]. In this framework, pharmacists could assess renal trends, indication, drug monitoring needs, and patient-specific exceptions before endorsing, modifying, or suppressing the recommendation. This design would also support documentation of the clinical reasoning behind dose changes and create a feedback channel for improving the dosing rule base.

Learning from Overrides
Override reasons should be logged in structured form and reviewed as part of continuous system governance. Studies of medication-related CDSS overrides show that alert burden and clinician response patterns can reveal whether alerts are poorly targeted, clinically inappropriate, or correctly overridden because of patient-specific exceptions [11, 25]. The renal dosing system could use these patterns to refine alert thresholds, update exception logic, and identify drug-rule combinations that require pharmacy or nephrology review. Override learning should not automatically weaken safety rules; instead, it should guide human-supervised revision of the knowledge base and alert prioritization strategy.

Integration Into Cpoe And Pharmacist Workflow
Embedding in the Prescribing Screen
The prescriber-facing interface should embed the renal dose recommendation directly within the CPOE workflow, ideally beside the active dose, route, and frequency fields. Renal medication alert studies in commercial EHR environments show that the presentation and timing of an alert strongly influence whether it is considered clinically useful or becomes another interruptive warning [1]. A CPOE-embedded recommendation could pre-populate an alternative renal-adjusted regimen while preserving the prescriber's ability to accept, modify, or override it with a documented reason. This design would reduce extra clicks and make the safest action easier without removing clinician responsibility for the final order.

Pharmacist Oversight and Documentation
The framework should integrate with the pharmacy information system so that pharmacists can review renal dosing recommendations, document their assessment, and communicate with prescribers when a recommendation is modified. Studies of computerized renal dose adjustment and broader medication-related decision support indicate that prescribing safety improves most plausibly when decision support is aligned with medication verification workflows rather than isolated at the point of ordering [2, 26]. Pharmacist documentation should capture the relevant renal trend, dosing rule, clinical exception, and final action taken. This shared record would support continuity of care, auditability, and later review by medication safety committees.

Evaluation Strategy
Alert Performance Metrics
The system should be evaluated against expert pharmacist review to determine whether its renal dosing recommendations identify clinically meaningful discrepancies without overwhelming clinicians. Evaluation should examine conceptual performance domains such as appropriateness, timeliness, explainability, and actionability rather than relying only on whether an alert was fired. Prior CDSS studies suggest that override patterns, alert burden, and clinician response should be interpreted together because a technically correct alert can still be ineffective if it is poorly timed or poorly integrated into workflow. The evaluation should also compare AI-augmented recommendations with existing rule-based renal alerts to understand whether trajectory-aware dosing logic improves clinical usefulness.

Impact on Prescribing and Clinical Outcomes
A pragmatic evaluation could examine whether the system would be associated with more appropriate renal dose adjustment, safer medication monitoring, and fewer medication-related kidney safety concerns. Prior studies of renal dosing support, outpatient prescribing interventions, and adverse drug event prevention provide useful models for evaluating prescribing processes and medication safety outcomes without assuming that alerts alone will improve all downstream clinical indicators [2, 3, 26]. Nephrotoxicity-focused work suggests that medication exposure, renal function, and clinical outcomes should be assessed together when evaluating kidney safety interventions [20]. Because this is a conceptual AIF manuscript, such outcomes should be treated as future evaluation targets rather than reported results.

User Satisfaction and Workflow Integration
User satisfaction evaluation should assess whether prescribers and pharmacists perceive the system as clear, trustworthy, timely, and compatible with routine clinical work. Medication-related CDSS studies show that usability and workflow fit are

essential because clinicians may override or ignore alerts that are difficult to interpret, repetitive, or poorly aligned with task timing [11, 25]. Surveys, interviews, and workflow observation could explore whether the recommendation language, explanation format, and accept-or-override pathway support efficient prescribing and pharmacy verification. The evaluation should also include governance review of suppressed alerts to ensure that fatigue reduction strategies do not unintentionally conceal important renal dosing risks.

Table 4 provides an evaluation and implementation readiness framework for determining whether AI-based renal dose decision support is accurate, actionable, workflow-compatible, explainable, and safe for prospective clinical deployment.

Table 4. Evaluation and Implementation Readiness Framework for AI-Based Renal Dose Decision Support

Evaluation domain	Core question	Suggested assessment approach	High-value success indicator	Failure mode the evaluation should detect	Implementation implication
Renal trajectory validity	Does the model correctly distinguish stable chronic impairment, acute decline, and possible recovery?	Compare model trajectory classifications with pharmacist/nephrologist adjudication and subsequent laboratory evolution	Forecasts identify clinically relevant renal deterioration before static threshold alerts would fire	False reassurance during rapidly worsening kidney function or overprediction during transient laboratory fluctuation	Initial deployment should emphasize high-risk medications and conservative escalation when uncertainty is high
Dosing-rule fidelity	Are renal dosing recommendations consistent with approved local formulary guidance?	Rule-base audit against institutional protocols, drug monographs, and pharmacist review	Recommended dose, interval, hold, or monitoring action matches locally approved renal dosing logic	Conflicting recommendations caused by nonstandard drug mapping, outdated rules, or inappropriate renal thresholds	Rule governance must be version-controlled and reviewed by pharmacy leadership
Medication-safety relevance	Does the system prioritize discrepancies most likely to cause harm?	Expert review of alert cases stratified by drug risk, renal clearance dependence, nephrotoxic potential, and electrolyte context	High-priority alerts concentrate on clinically consequential mismatch scenarios	Equal treatment of trivial and serious dosing discrepancies, leading to alert fatigue	Alert logic should weight medication risk and patient-specific renal trajectory rather than rely only on eGFR category
Alert appropriateness	Are alerts timely, specific, and actionable at the point of decision?	Prescriber and pharmacist case review using appropriateness, actionability, and timing criteria	Alerts provide a clear alternative regimen and rationale during ordering or verification	Vague warnings, repeated alerts after valid override, or alerts delivered after the decision window	CPOE integration should favor embedded alternatives over generic warning messages
Alert-burden reduction	Does context-aware prioritization reduce low-value interruptions without hiding important risk?	Compare interruptive alert rate, acceptance rate, override reasons, and suppressed-alert audit before and after implementation	Fewer low-value alerts with maintained or improved detection of high-risk renal dosing problems	Unsafe suppression of clinically important discrepancies or excessive non-interruptive routing	Suppression rules should be transparent, monitored, and periodically reviewed by medication safety teams
Pharmacist workflow integration	Does the system strengthen pharmacist verification rather than create parallel work?	Workflow observation, queue-volume analysis, time-to-review measurement, pharmacist satisfaction assessment	Pharmacists can efficiently endorse, modify, or suppress recommendations with documented rationale	Duplicative review queues, unclear accountability, or excessive manual reconciliation	Pharmacist-facing design should align with existing verification tasks and documentation practices
Prescriber usability	Do prescribers understand and act on recommendations without excessive cognitive burden?	Usability testing, task-completion analysis, acceptance/override review, interviews	Prescribers understand the renal trend, drug rule, and recommended alternative dose quickly	Confusing explanation language, excessive clicks, or perceived loss of clinical autonomy	Interface should show concise rationale and preserve accept, modify, override, and pharmacist-review options
Explanation quality	Does the system explain why a recommendation was generated?	Review of explanation completeness across medication order, renal trend, forecast, dosing rule, and alert priority	Explanation links recent laboratory trend, forecasted renal state, drug-specific rule, and recommended action	Black-box recommendations that clinicians cannot verify or trust	Every prescriber-facing alert should include a short rationale and traceable dosing rule

Override learning	Do structured override data improve governance without eroding safety?	Periodic review of override reasons, repeated alert patterns, valid exceptions, and adverse-event signals	Override patterns identify rule refinements, exception logic, or education needs	Automatic reduction of alert sensitivity based solely on frequent overrides	Override analytics should inform human-supervised governance rather than autonomous rule relaxation
Equity and generalizability	Does the system perform safely across patient groups, services, and care settings?	Stratified evaluation by age group, sex, weight extremes, dialysis status, ICU status, oncology status, and pediatric/adult population where applicable	Comparable recommendation quality and alert appropriateness across major clinical subgroups	Poor calibration in specialty populations or transfer failure across institutions	Expansion beyond initial medication groups should require local validation and subgroup review
Clinical outcome readiness	Is the system ready for prospective evaluation of medication safety impact?	Pragmatic pilot comparing baseline renal dosing processes with AI-augmented decision support	Improved appropriate renal dose adjustment, faster pharmacist intervention, and acceptable alert burden	Reliance on alert firing as the only endpoint without measuring clinical action or safety	Prospective trials should assess prescribing behavior, medication monitoring, adverse drug events, and workflow burden
Governance sustainability	Can the institution maintain the system safely over time?	Review of ownership model, rule-update process, model monitoring plan, audit procedures, and escalation pathways	Clear accountability for dosing rules, model monitoring, alert review, and safety escalation	Model drift, outdated renal rules, unmanaged suppressed alerts, or unclear responsibility	Deployment should include medication safety committee oversight and scheduled recalibration reviews

Limitations

Dependence on Accurate and Timely Data

The system would depend on timely laboratory results, accurate medication order status, reliable patient demographic data, and correct documentation of dialysis or renal replacement therapy. AKI prediction studies based on EHR data show that model outputs are only as reliable as the clinical data streams that feed them, particularly when laboratory timing and care transitions affect interpretation [12, 14]. Missing creatinine values, delayed laboratory posting, inaccurate weights, or unrecorded medication administration changes could cause the system to understate or overstate renal dosing risk. For that reason, the framework should include data-quality checks, uncertainty displays, and pharmacist escalation when essential inputs are incomplete.

Generalizability across Drugs and Patient Populations

Generalizability would be limited by differences in renal dosing rules, institutional formularies, patient populations, and clinical settings. Pediatric AKI prediction work and specialty dosing studies suggest that renal risk models and medication adjustment logic may require population-specific calibration rather than direct transfer across adults, children, oncology patients, and critically ill patients [21, 23]. Drugs with complex pharmacokinetics, sparse renal dosing evidence, or indication-dependent loading strategies may need manual pharmacist review before automated recommendations are enabled. Initial deployment should therefore focus on well-characterized high-risk medications and expand only after governance review, local validation, and clinician acceptance testing.

Conclusion

An AI decision support system for renal dose adjustment could integrate medication orders, serial laboratory trends, and patient-specific clinical context to generate individualized dosing recommendations. By moving beyond static creatinine threshold alerts, the framework would support earlier recognition of renal decline and more timely adjustment of renally cleared or nephrotoxic medications. The system would function as a decision support layer rather than an autonomous prescribing tool.

The main strength of the proposed framework is its combination of real-time prediction, drug-specific dosing logic, context-aware alerting, and pharmacist verification. This design could reduce unnecessary alerts while increasing the clinical value of recommendations that do reach prescribers. Its human-in-the-loop structure would also support accountability and clinical judgment in cases where guideline-based dosing does not fully capture patient complexity.

Important challenges remain before such a system could be deployed broadly. The framework would depend on data freshness, accurate medication order mapping, well-curated renal dosing rules, and careful calibration across institutions and patient populations. It would also require governance mechanisms to monitor overrides, update rules, and ensure that alert suppression does not compromise patient safety.

Future work should prioritize pragmatic clinical trials and implementation studies that evaluate safety, usability, and workflow impact in real care settings. Such studies should determine whether AI-augmented renal dosing support improves medication

safety while remaining acceptable to prescribers and pharmacists. Broad deployment should follow only after rigorous validation, transparent governance, and evidence that the system improves care without increasing clinical burden.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Shah SN, Amato MG, Garlo KG, Seger DL, Bates DW. Renal medication-related clinical decision support (CDS) alerts and overrides in the inpatient setting following implementation of a commercial electronic health record: implications for designing more effective alerts. J Am Med Inform Assoc. 2021;28(6):1081-7.
2. Wada R, Takeuchi J, Nakamura T, Sonoyama T, Kosaka S, Matsumoto C, et al. Clinical decision support system with renal dose adjustment did not improve subsequent renal and hepatic function among Inpatients: the Japan adverse drug event study. Appl Clin Inform. 2020;11(5):846-56.
3. Li YJ, Lee WS, Chang YL, Chou YC, Chiu YC, Hsu CC. Impact of a clinical decision support system on inappropriate prescription of glucose-lowering agents for patients with renal insufficiency in an ambulatory care setting. Clin Ther. 2022;44(5):710-22.
4. Sonoda A, Kondo Y, Iwashita Y, Nakao S, Ishida K, Irie T, et al. In-hospital prescription checking system for hospitalized patients with decreased glomerular filtration rate. Kidney360. 2022;3(10):1730-7.
5. Koyner JL, Carey KA, Edelson DP, Churpek MM. The development of a machine learning inpatient acute kidney injury prediction model. Crit Care Med. 2018;46(7):1070-7.
6. Tomašev N, Glorot X, Rae JW, Zielinski M, Askham H, Saraiva A, et al. A clinically applicable approach to continuous prediction of future acute kidney injury. Nature. 2019;572(7767):116-9.
7. Sonoda A. A clinical decision support system promotes the appropriate use of drugs in hospitalized patients with kidney impairment. J Pharm Health Care Sci. 2025;11(1):26.
8. Desmedt S, Spinewine A, Jadoul M, Henrard S, Wouters D, Dalleur O. Impact of a clinical decision support system for drug dosage in patients with renal failure. Int J Clin Pharm. 2018;40(5):1225-33.
9. Sancar M, Altiparmak O, Altan A, Demirtunc R, Izzettin FV, Okuyan B. Evaluation of medication dose adjustments in patients with impaired renal function using different online drug information databases. Eur J Hosp Pharm. 2022;29(5):255-8.
10. Martin M, Wilson FP. Utility of electronic medical record alerts to prevent drug nephrotoxicity. Clin J Am Soc Nephrol. 2019;14(1):115-23.
11. Nanji KC, Seger DL, Slight SP, Amato MG, Beeler PE, Her QL, et al. Medication-related clinical decision support alert overrides in inpatients. J Am Med Inform Assoc. 2018;25(5):476-81.
12. Martinez DA, Levin SR, Klein EY, Parikh CR, Menez S, Taylor RA, et al. Early prediction of acute kidney injury in the emergency department with machine-learning methods applied to electronic health record data. Ann Emerg Med. 2020;76(4):501-14.
13. Chiofolo C, Chbat N, Ghosh E, Eshelman L, Kashani K. Automated continuous acute kidney injury prediction and surveillance: a random forest model. Mayo Clin Proc. 2019;94(5):783-92.
14. Sandokji I, Yamamoto Y, Biswas A, Arora T, Ugwuowo U, Simonov M, et al. A time-updated, parsimonious model to predict AKI in hospitalized children. J Am Soc Nephrol. 2020;31(6):1348-57.
15. Rank N, Pfahringer B, Kempfert J, Stamm C, Kühne T, Schoenrath F, et al. Deep-learning-based real-time prediction of acute kidney injury outperforms human predictive performance. NPJ Digit Med. 2020;3(1):139.
16. Le S, Allen A, Calvert J, Palevsky PM, Braden G, Patel S, et al. Convolutional neural network model for intensive care unit acute kidney injury prediction. Kidney Int Rep. 2021;6(5):1289-98.
17. Abdullah SS, Rostamzadeh N, Sedig K, Lizotte DJ, Garg AX, McArthur E. Machine learning for identifying medication-associated acute kidney injury. Informatics. 2020;7(2):18.
18. Griffin BR, Mudireddy A, Horne BD, Chonchol M, Goldstein SL, Goto M, et al. Predicting nephrotoxic acute kidney injury in hospitalized adults: a machine learning algorithm. Kidney Med. 2024;6(12):100918.
19. Shawwa K, Ghosh E, Lanius S, Schwager E, Eshelman L, Kashani KB. Predicting acute kidney injury in critically ill patients using comorbid conditions utilizing machine learning. Clin Kidney J. 2021;14(5):1428-35.
20. Griffin BR, Wendt L, Vaughan-Sarrazin M, Hounkponou H, Reisinger HS, Goldstein SL, et al. Nephrotoxin exposure and acute kidney injury in adults. Clin J Am Soc Nephrol. 2023;18(2):163-72.

21. Rolvink J, Gerards AL, Kater AP, Swart EL, Becker ML. The effect of a Clinical Decision Support System on the frequency of dose adjustments of anticancer drugs in case of renal or hepatic dysfunction. J Oncol Pharm Pract. 2022;28(5):1063-9.
22. Karway GK, Koyner JL, Caskey J, Spicer AB, Carey KA, Gilbert ER, et al. Development and external validation of multimodal postoperative acute kidney injury risk machine learning models. JAMIA Open. 2023;6(4):ooad109.
23. Fragasso T, Raggi V, Passaro D, Tardella L, Lasinio GJ, Ricci Z. Predicting acute kidney injury with an artificial intelligence-driven model in a pediatric cardiac intensive care unit. J Anesth Analg Crit Care. 2023;3(1):37.
24. Li X, Wang P, Zhu Y, Zhao W, Pan H, Wang D. Interpretable machine learning model for predicting acute kidney injury in critically ill patients. BMC Med Inform Decis Mak. 2024;24(1):148.
25. Wong A, Wright A, Seger DL, Amato MG, Fiskio JM, Bates D. Comparison of overridden medication-related clinical decision support in the intensive care unit between a commercial system and a legacy system. Appl Clin Inform. 2017;8(3):866-79.
26. Hak F, Guimarães T, Santos M. Towards effective clinical decision support systems: A systematic review. PLoS One. 2022;17(8):e0272846.