

Research Article

Cell-type specific single-cell signatures reveal nephrotoxic drug effects

Aditi Kuchi , Jose Miguel Acitores Cortina, Hongyu Liu , Yasaman Fatapour, Jacob Berkowitz, Nicholas P. Tatonetti *

Cedars-Sinai Medical System, Department of Computational Biomedicine, 700 N. San Vicente Blvd., West Hollywood, CA 90048, USA

ARTICLE INFO

Keywords:

Single cell RNAseq
Nephrotoxicity
Single cell signatures

ABSTRACT

Drug-induced acute kidney injury (AKI) affects about 20 % of hospitalized AKI patients, a significant contributor to morbidity and mortality. The lack of understanding of the kidney system and functioning of nephrotoxic drugs contributes to hospital-acquired AKI cases. AKI is difficult to predict because of its complex injury mechanism and the numerous pathways through which it manifests. Traditional toxicity biomarkers, like elevated creatinine levels, detect AKI only after significant kidney injury has occurred. Concurrently, advancements in single cell RNA sequencing (scRNAseq) have improved our ability to map cellular heterogeneity within tissues, potentially enabling the study of drug effects at a single cell level. We hypothesized that only particular subtypes of kidney cells may be responsible for observed nephrotoxicity and explain prediction challenges. To test this, we generated cellular response scores for 32 kidney cell types from the Human Cell Atlas and estimated drug effects. We identified significant expression differences in 6 cell types (e.g. Indistinct intercalated cell $p = 0.009$, Epithelial Progenitor cell, $p = 0.04$). We also developed an XGBoost model that achieved an AUROC of 0.6 on an external test set, across different kidney cell populations - a significant improvement over using traditional bulk RNA sequencing alone. The single-cell transcriptomic signatures we identified potentially reveal unexplained molecular mechanisms of nephrotoxicity. This work provides both a reproducible computational framework and curated dataset available at <https://doi.org/10.5281/zenodo.15724290> to the research community.

Author Summary

The prediction and early detection of drug-induced kidney injury is a significant clinical challenge since physicians rely on biomarkers that only become elevated after substantial kidney damage has occurred, limiting opportunities for intervention and patient protection. Our investigation utilized kidney single-cell data and available drug toxicity information from FDA labels to examine how individual kidney cell populations respond to potentially harmful medications. We hypothesized that specific kidney cell subtypes are primarily responsible for observed drug toxicity, which may explain the difficulties in predicting drug-induced kidney injury using bulk RNAseq data.

Through comprehensive analysis of 32 distinct kidney cell types, we identified 6 specific cellular populations that demonstrate differential responses to nephrotoxic compounds. We subsequently developed models that demonstrate superior predictive performance compared to analytical approaches using bulk RNA sequencing data. Our methodology represents a substantial advancement in precision medicine approaches to drug safety. These findings have important implications for clinical practice and patient safety. The cellular signatures we identified

may enable earlier detection of kidney injury risk, potentially allowing clinicians to modify treatment regimens before irreversible damage occurs. Our work establishes a foundation for improved drug safety protocols and may contribute to reducing medication-related kidney injury in hospitalized patients.

1. Introduction

Acute Kidney Injury (AKI) is a critical health concern, accounting for 20–60 % of AKI cases in hospitalized patients [1]. The incidence of AKI in hospitalized patients ranges from 5 % to 67 % in patients admitted to intensive care [2]. A significant contributor to AKI is drug-induced nephrotoxicity, which is particularly concerning due to the kidney's crucial role in excreting toxic metabolites and drugs. This risk is further elevated in older age groups who often have coexisting morbidities and are exposed to more diagnostic procedures and therapies [3].

Drug-induced nephrotoxicity operates through complex mechanisms that affect specific cellular components of the kidney. As detailed by Shi et al., [4] nephrotoxic agents can be categorized based on the

* Corresponding author.

E-mail addresses: Aditi.kuchi@cshs.org (A. Kuchi), Nicholas.tatonetti@cshs.org (N.P. Tatonetti).

<https://doi.org/10.1016/j.csbj.2025.11.012>

Received 8 July 2025; Received in revised form 3 November 2025; Accepted 4 November 2025

Available online 6 November 2025

2001-0370/© 2025 The Authors. Published by Elsevier B.V. on behalf of Research Network of Computational and Structural Biotechnology. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

histological component affected, including proximal tubular injury (PTI), tubular obstruction by crystals or drug metabolites, and interstitial nephritis. At the cellular level, drug toxicity manifests through several distinct pathways that damage kidney cells, such as accumulation of nephrotoxic agents in proximal convoluted tubule epithelial cells, causing cellular damage, disruption of ATP production, triggering cell death, induction of oxidative stress and activation of ferroptosis [5] (iron overload) pathways in acute kidney injury cases. These mechanisms represent key cellular processes that contribute to nephrotoxicity across various drug classes.

Current bulk RNAseq studies have mostly focused on animal models to find toxicity biomarkers [6]. Despite major consortiums like the International Life Sciences Institute and the Health and Environmental Sciences Institute (ILSI/HESI) nephrotoxicity working group identifying promising panels of biomarkers through collaborative research, translational challenges persist between clinical outcomes and pre-clinical predictions. While these studies have developed expression biomarkers in rodent models, there remains a significant gap in translating these findings to human patients, with only limited assays being employed by pharmaceutical companies to assess nephrotoxicity before deploying drugs to market. The high levels of discrepancy between the low failure rate of drug candidates due to nephrotoxicity in pre-clinical trials (only 7 %) [7], and the high incidence of AKI in intensive care patients (20–50 %) [8] brings to light the critical need for more detailed and informed analysis on toxicity data (see graphical abstract).

Recent advancements in single-cell RNA sequencing (scRNAseq) technology have advanced our understanding of cellular heterogeneity within the human kidney, providing unprecedented resolution in studying drug effects at the cellular level. The abundance of human single-cell data in kidney research has grown exponentially in recent years. For instance, the Human Cell Atlas project [9] has analyzed over 100,000 cells from human kidney tissue, providing a high-resolution map of human kidney cell types and states which we use in this study. This abundance of human cell-level data, combined with comprehensive information on drug targets, presents an opportunity to identify if particular subtypes of kidney cells are responsible for nephrotoxic effects of drugs.

This hypothesis is not without precedent. Recently, Muto et al. [10] used single-cell transcriptomics of human kidney organoids to identify specific proximal tubule cell populations that are particularly vulnerable to cisplatin-induced injury [11], a common consequence of chemotherapy. Similar analyses [12,13], when extended to other drug-induced injuries, could reveal cell type-specific vulnerabilities and resistance mechanisms in human kidney tissue. The integration of drug target and related information with single-cell data also opens up new avenues for predictive drug toxicity in patients [14,15].

In this work, we analyze kidney specific bulk RNAseq data from National Institute of Health, Library of Integrated Network-based Cellular Signatures, commonly referred to as LINCS [16,17], and single-cell RNAseq data from Human Cell Atlas (HCA) [18] in combination with curated datasets with known drug toxicity annotations derived from labels. We used the Ryan reference set [19], American Chemical Society Publications (ACS) Food and Drug Administration (FDA) drug labels [20], Drug Induced Renal Injury List (DIRIL) dataset [21] from the FDA) to show that single cell data 1) has more statistical power than bulk RNAseq dataset to create a toxicity signature for drugs, 2) certain cell types that could be highly affected by nephrotoxic drugs are masked in bulk RNAseq representation, and 3) the presence of such a cell-type specific toxicity signature could help isolate drug candidates whose mechanisms are not yet known to cause nephrotoxicity. Our findings highlight the presence of cell types that are most susceptible to drug induced injury, consistent with other literature in the field [22].

2. Results

2.1. Simulations show bulk RNA data has limited ability to identify nephrotoxicity mediated by cell subtypes

In this section, we explored the use of bulk RNA sequencing for detecting nephrotoxicity signatures under the assumption that the nephrotoxicity is mediated by a kidney cell subtype. To evaluate the data in this context, we simulated scenarios where nephrotoxicity was mediated by Nephron, Endothelium, Immune and Stroma subtypes, representing approximately 82.4 %, 7.8 %, 7.1 % and 2.6 % of kidney cells, respectively. We then created pseudo-bulk data by computing the average across the cellular subtypes (see Methods). We evaluated the ability of this pseudo-bulk RNA sequencing data to differentiate nephrotoxic from non-nephrotoxic effects assuming a hypothetical nephrotoxic biomarker expression change by simulating a 0.01–10-fold change between treated and untreated samples. As shown in Figs. 2 and 3, we found that statistical power was under 60 % even at a maximum fold change of 10X for two cell-type abstraction levels.

Our analysis demonstrates a positive correlation between effect size and statistical power, as illustrated in Fig. 2a, b, c and 3a. Assuming the effect size multiplier represents a directly proportional drug action, we observe larger effect sizes in the simulated data consistently yielded increased statistical power. This supports our hypothesis that single cell level drug effect studies are superior when compared to bulk data analysis for detecting toxicity signatures.

2.2. Nephrotoxicity signatures are obscured in RNAseq of kidney epithelial cells

We performed differential expression analysis on the LINCS L1000¹⁶ dataset to compare 355 nephrotoxic samples (from 19 nephrotoxic drugs) and 16,955 non-nephrotoxic samples (from 857 non-nephrotoxic drugs). The drugs were annotated based on a curated dataset including manually verified FDA drug labels from the Ryan reference set [19], ACS FDA labels [20] and DIRIL [21], identifying toxicity status for kidneys. We identified 100 differentially expressed genes ($p < 0.05$, $|\log \text{fold change}| \leq 1$) (Top 5 shown in Table 1, all in Supplemental Table 1). Independently, we trained three machine learning classifiers and evaluated using five-fold cross validation (Table 2).

2.3. Single-cell analysis identifies novel nephrotoxic signatures

2.3.1. Nephrons are significantly affected by nephrotoxic drugs

We used single cell data from public repositories to identify novel nephrotoxic signatures. We process all datasets using CellRanger [23] and Seurat [24], following standard protocols (see Methods). We annotated the data using SingleR [25,26] with the kidney cell atlas [27] as reference. We derive the drug target information from Drugbank [28] and ChEMBL [29], whereas the toxicity labels are collected from the Ryan reference, ACS FDA labels, and the DIRIL dataset. We map the cells to drugs using their targets. Using drug2cell with our custom database generates a composite score for each cell-drug pair representative of their action.

The output of the drug2cell package is a vector for each drug of the drug's estimated effect on each cell. We grouped cells by their annotated type and then tested for differences in the drug2cell scores between scores for nephrotoxic drugs and non-nephrotoxic drugs using a T-test. We performed this evaluation at two levels of cell type abstraction (4 cell types called “abstract”, and 32 cell types called “detailed”). The 4-type abstraction is a superset of the 32-type annotations. This dual-level analysis was conducted to determine whether nephrotoxicity signals were dependent on cell type resolution. The p-values and other statistics in abstract and detailed cell type are reported in Tables 3 and 4 respectively. Supplemental Table 2 reports the p-values of all detailed cell types. We found that Nephrons, Endothelium and Stroma cells




Fig. 1. Workflow diagram – Drug target information is collected from ChEMBL [29] and Drugbank [28] and consolidated. Ryan reference [19], ACS FDA label data [20] and DIRIL datasets [21] are combined to extract toxicity labels for drugs. Logical OR is used to label them to maximize the number of nephrotoxic drugs. To create a unified Drug-Target-Toxicity database the targets and toxicity labels are combined. This database is used in all downstream analyses. For the single-cell analysis, data is collected from HCA and GEO, selecting only healthy kidney samples. We use Cell Ranger to process FASTQ files, Seurat to analyze the data, QC, filter and integrate (see Methods). We annotate the clusters using SingleR, with the Kidney Cell Atlas as reference at two levels of abstraction (4 cell types, called abstract, and 32 cell types called detailed.) We analyze both levels of abstraction at all phases downstream. We then use Drug2Cell package to extract drug scores based on the target genes in the single cell matrix. We log these scores to ensure normal distribution. Our bulk RNAseq analysis extracts kidney-specific LINCS data to identify differentially expressed genes and extractable toxicity signatures. Independently, we simulated drug-affected single-cell expression data and pseudo-bulked it to compare data type power for toxicity predictions. We then used both single-cell and bulk RNAseq toxicity matrices to train machine learning methods for identifying drug toxicity signatures. We include an external test set for thorough validation of the models.




Fig. 2. 4 cell-type results. In figures a), b), c) are power plots for different response rates. Response rates are defined as the percentage of cells within each sub-type that respond to a drug. Power is defined as the number of differentially expressed genes divided by total number of genes, represented on the y axis of the plots. On the x axis is the effect size. Figure d) shows a UMAP of the clustered single cell data coded by color. Many of the cells are Nephrons (82.4 %). Boxplots showing the logged drug scores are in figure e). There is a marked difference in the drug scores between nephrotoxic (dark) and non-nephrotoxic (light) drugs. The nephrotoxic drug scores trend higher.




Fig. 3. Plots showing single cell data having higher nephrotoxicity detection power compared to pseudo-bulked data. The hollow dots represent pseudo-bulked data and are close to the x axis signifying that the power is very low. Box plots show the increased drug scores of nephrotoxic drugs (dark) in 32 cell types.

Table 1
Top 5 differentially expressed genes between nephrotoxic and non-nephrotoxic drugs in kidney specific cell lines (HA1E) from LINCS. Some of these genes (HDAC6, ACKR3) are implicated in AKI and kidney damage pathways. The full list of genes can be found in [Supplemental Table 1](#). When we collapse replicates, there are no differentially expressed genes.

Gene Name	pvalue	log2FoldChange	mean_toxic	mean_nontoxic	FDR
TOP2A	8.53E−08	1.295	0.200	−1.254	6.85E−05
CYTH1	1.91E−07	2.179	0.405	0.011	0.0001
PRSS23	2.07E−07	1.059	0.119	−0.965	0.0001
HDAC6	5.92E−07	1.577	0.427	0.076	0.0002
ACKR3	3.51E−06	1.732	0.373	0.043	0.0006

Table 2
Performance metrics for three machine learning models on kidney-specific cell lines. We used 5-fold cross validation to evaluate the models. The accuracy performance does not cross 0.5, showing the inadequacy of these models using bulk RNAseq data.

Model	Accuracy	F1 score	Precision	Recall	AUC
Logistic regression	0.458	0.5	0.4	0.66	0.416
XGBoost	0.5	0.0	0.0	0.0	0.5
SVM	0.458	0.5	0.4	0.66	0.416

showed significant differences in drug scores between nephrotoxic and non-nephrotoxic drugs (Table 3). We found that Indistinct intercalated cell, MNP-b/non-classical monocyte derived cell, Peritubular capillary endothelium, Distinct proximal tubule, Proliferating Proximal Tubule, and Epithelial progenitor cell types showed significant differences in drug scores in the detailed cell type analysis (shown in Table 4). We found that an XGBoost model trained on the detailed 32 cell types data achieved an AUROC of 0.6.

2.3.2. XGBoost model captures better toxicity signature in single-cell

We evaluated 7 classical machine learning algorithms and a 2-layer multilayer perceptron (MLP) through ROC curve analysis, demonstrating improvement over traditional bulk RNA-seq-based models in Table 2. Individual algorithms (Logistic Regression, XGBoost, Random Forest, and Support Vector Machine, Extra Tree, Gradient Boosting) showed moderate performance with AUROC values ranging from 0.55 to 0.67 (Table 5). To test the models, we evaluate on an independent kidney dataset with diabetic and non-diabetic kidney single cells. We also perform hyperparameter tuning and some feature engineering (detailed in Methods). Table 5 shows the performance of these models on the internal and external test set. The dataset was split into a 80–20 train-test, for which the internal scores are reported. The external dataset shows the performance metrics of the model on the unseen data. XGBoost demonstrated the strongest external validation performance with an AUROC of 0.606 (95 % CI: 0.56 – 0.64), with accuracy improving from 0.71 to 0.72 and specificity of 0.92. This indicates excellent ability to correctly identify non-toxic drugs.

The consistent AUROC values ranging from 0.55 to 0.67 on internal test, and 0.52–0.61 on external test set, coupled with accuracies for 0.63–0.75 show that our single cell approach captures genuine nephrotoxicity signals that generalize beyond the training data.

3. Discussion

Our study used single-cell data to investigate the association between

specific kidney cell types and nephrotoxic drugs. The results revealed statistically significant associations between nephrons, endothelium and immune cells with nephrotoxic drugs. Specifically, we identified six detailed cell types that demonstrated significance: Indistinct intercalated cell, MNP-b/non-classical monocyte derived cell, Peritubular capillary endothelium, Distinct proximal tubule, Proliferating Proximal Tubule, and Epithelial progenitor cells. Notably, the majority of these are nephrons, suggesting that nephron-specific cellular mechanisms play a crucial role in drug-induced kidney injury.

The abundance of single-cell data and the inherent heterogeneity of kidney tissue allow for a detailed examination of specific regions that may be more susceptible to drug effects. The kidney is an ideal organ for this analysis due to several important characteristics. It contains discrete cell types that perform distinct functions in drug metabolism and excretion. Different anatomical regions of the kidney are specifically responsible for metabolizing various drugs, making it possible to pinpoint where toxicity may occur. A granular understanding could impact the drug screening process. Since kidneys are responsible for filtration and excretion of drug metabolites, the byproducts are easily measurable in urine and blood, providing convenient biomarkers for toxicity assessment. This approach is also readily extensible to other tissue types and diseases, with liver tissue being another excellent candidate for similar analyses due to its central role in drug metabolism. If we can identify a stronger signature indicative of toxicity in certain cell types, it could lead to the development of more accurate models for screening drug candidates. Moreover, this knowledge could benefit patients undergoing intensive care who are at risk of further complications.

Our analysis reveals critical insights into drug-induced kidney injury that are obscured in bulk RNA sequencing studies. While bulk transcriptomics captures general tissue responses, it masks cell type-specific signals due to averaging across the kidney’s cellular heterogeneity. This "signal dilution" limits detection of toxicity patterns confined to particular microenvironments, or what we term "toxicity niches." In contrast, single-cell RNA sequencing preserves this cellular resolution. Even with a coarse classification into four major cell types, our model demonstrates

Table 3
Abstract (4) cell-type *t*-test results show that Nephrons, Endothelium and Stroma are significant.

Cell Type	T-statistic	Log Fold Change	Fold Change	CI Lower 95 %	CI Upper 95 %	P-value	Corrected P-value
Nephron	2.673	0.378	1.460	1.121	1.902	0.007	0.030
Endothelium	2.283	0.303	1.354	1.059	1.734	0.022	0.033
Stroma	2.246	0.292	1.340	1.058	1.697	0.024	0.033
Immune	1.198	0.171	1.186	0.898	1.565	0.231	0.231

Table 4
Detailed (32) cell-type *t*-tests shows six cells with significant expression between nephrotoxic and non-nephrotoxic drugs.

Detailed Cell Type	T-statistic	Log Fold Change	Fold Change	CI Lower 95 %	CI Upper 95 %	P-value	Corrected P-value	Abstract Cell-Type
Indistinct intercalated cell	−3.639	−0.524	0.592	0.440	0.793	0.0002	0.009	Nephron
MNP-b/non-classical monocyte derived	3.032	0.453		1.184	2.091	0.0024	0.032	Immune
Peritubular capillary endothelium 1	2.972	0.394	1.574	1.161	1.893	0.0030	0.032	Endothelium
Distinct proximal tubule 1	2.754	0.367	1.483	1.136	1.833	0.0060	0.046	Nephron
Proliferating Proximal Tubule	2.689	0.466	1.443	1.154	2.201	0.0072	0.046	Nephron
Epithelial progenitor cell	2.613	0.394	1.594	1.119	1.965	0.0091	0.048	Nephron
			1.483					

Table 5
Performance metrics for various machine learning models using single-cell RNAseq drug scores as features. There is a marked increase in the AUC compared to bulk RNAseq. We performed external validation on an independent dataset not used in model training. CI is confidence interval, int stands for internal, and ext stands for external.

Model	Accuracy		F1 Score		AUROC		Sensitivity		Specificity		Lower - Upper 95 % CI
	Int	Ext	Int	Ext	Int	Ext	Int	Ext	Int	Ext	Ext AUROC
XGBoost	0.715	0.728	0.364	0.184	0.678	0.606	0.326	0.123	0.882	0.921	0.561–0.648
Random Forest	0.634	0.747	0.182	0.018	0.64	0.61	0.163	0.009	0.837	0.968	0.568–0.654
Voting Ensemble	0.68	0.736	0.247	0.089	0.624	0.603	0.209	0.052	0.884	0.953	0.557–0.646
Stacking Ensemble	0.709	0.751	0.167	0.009	0.649	0.556	0.116	0.005	0.953	1	0.512–0.604
MLP	0.669	0.678	0.345	0.291	0.558	0.588	0.349	0.264	0.791	0.814	0.544–0.633
Logistic Regression	0.552	0.618	0.374	0.319	0.563	0.539	0.535	0.359	0.558	0.697	0.491–0.584
Support Vector Machine	0.634	0.673	0.323	0.142	0.578	0.528	0.349	0.109	0.744	0.837	0.483–0.572
Extra Trees	0.698	0.75	0.161	0	0.613	0.544	0.116	0	0.93	1	0.504–0.588

superior predictive power relative to the bulk LINCS dataset, indicating a qualitative shift in our ability to detect nephrotoxicity. This advantage likely reflects the biology of injury itself: nephrotoxic compounds often target cellular processes enriched in specific cell types. For example, proximal tubule cells are vulnerable to mitochondrial toxins due to their high metabolic activity, while podocytes are more susceptible to compounds disrupting cytoskeletal integrity. By resolving these cell type-specific responses, single-cell analysis captures toxicity signatures that bulk approaches miss. This aligns with recent studies showing distinct injury responses across kidney cell types [22] and underscores the importance of selective cell death and fibrosis pathways in the progression from acute kidney injury (AKI) to chronic kidney disease (CKD) [30,31]. Ultimately, maintaining cellular context may offer valuable sensitivity in detecting early or subtle nephrotoxic events, where conventional biomarkers fall short. Recent applications of single-cell RNAseq in AKI research have shown the power of cell-type resolved analysis for understanding renal pathophysiology [32].

One advantage of our approach is that it bypasses the need for complex pathway analysis while still capturing the essential biology of nephrotoxicity. Rather than requiring detailed knowledge of specific molecular pathways disrupted by each potential toxic drug, which varies widely between drug classes, our method leverages the cellular expression patterns and their interaction with the drug targets themselves as an integrated score of toxicity. This represents a robust approach that can detect novel forms of nephrotoxicity even when the precise mechanism is unknown. In the drug toxicity testing pipeline, focusing on these significant cell types will help pharma companies develop more targeted and precise screening methods. The power analysis we conducted specifically shows the value of utilizing single cell

data to reveal nephrotoxicity signatures. Despite the promising results, our study has several limitations. There is scope to improve the simulation of single-cell data, as we do not account for functional relationships between genes. Gene expression of certain significant genes likely varies in combination with others, creating complex patterns that our current models may not fully capture. Another key limitation was the exclusive use of healthy donor kidneys. This restricts our understanding of how drug effects might manifest in diseased states. Although a study based solely on normal kidneys provides valuable baseline information, incorporating data from diseased kidneys would prove highly beneficial, making it more directly clinically relevant. We acknowledge that our analysis uses single-cell data from healthy kidney tissue, which may limit the detection of drug-induced transcriptional changes that happen in the wet-lab/real world. As demonstrated in recent work [33], drug effects in normal tissue can be subtle and challenging to detect via expression profiling. We expect the actual prediction performance to drastically improve when disease-relevant or drug-perturbed single-cell datasets are used, which will better capture the biological response to nephrotoxic drug exposure.

This work lays the foundation for a more nuanced and effective approach to predicting and mitigating adverse drug effects in the kidney. By continuing to refine and validate these methods, we aim to develop tools that can significantly improve drug safety assessment and ultimately enhance patient outcomes through reduced incidence of drug-induced kidney injury. The data and code presented in this paper are immediately available for researchers through the link (<https://doi.org/10.5281/zenodo.15724290>) and our GitHub: <https://github.com/tatonetti-lab/singlecell-nephrotoxicity>). We view this study as establishing a proof of concept that justifies expansion to a multi-tissue, multi-

toxicity atlas. Future work will validate this approach across different organ systems before developing a prediction webtool that can serve the research community.

4. Data and methods

4.1. Data sources and workflow

For this study, we created a curated dataset comprising data from:

- a) DrugBank: We extracted information for approved or withdrawn drugs to ensure clinical relevance.
- b) ChEMBL: This provided experimentally validated data for drug-target pairs.
- c) Nephrotoxicity Dictionaries: We incorporated three specialized nephrotoxicity resources:
 - i. Ryan reference: A manually curated list of drugs with verified nephrotoxicity status.
 - ii. ACS FDA labels: A comprehensive collection of FDA-approved drug labels with extracted nephrotoxicity information.
 - iii. DIRIL dataset: An FDA confirmed nephrotoxicity labeled dataset on 317 single-molecule drugs.

This data resource contains the mappings from common approved drugs names to their targets and their toxicity labels extracted from various validated sources. In some cases where the drug might have conflicting information, we annotate it as toxic if at least one source out of three has labeled it as toxic to kidneys. We will refer to this curated database as the “drug-target-toxicity” database (Fig. 1, pink block).

We collected and processed normal kidney single-cell RNA sequencing (scRNA-seq) data from the Human Cell Atlas (HCA) (DOI: 10.1038/s41586-023-05769-3) and Gene Expression Omnibus (GEO), accession number GSE131685 [34]. Following the workflow shown in Fig. 1 (blue block), we processed and integrated these datasets using Seurat v5.1.2 in R v4.4.2. For samples with available FASTQ files, we used Cell Ranger v1.1.0 (provided by 10X Genomics) to generate gene-barcode matrices. We then converted the generated gene-barcode matrices to count data and then Seurat objects.

We implemented quality control following established best practices for single-cell RNA sequencing data analysis. We filtered cells using the following criteria: (1) minimum UMI count threshold of 500; (2) minimum gene detection threshold of 250 genes per cell; and (3) log10 genes per UMI ratio greater than 0.85, which helps identify and remove potential technical artifacts. We then filtered genes by retaining only those present in at least 10 cells to ensure reliable downstream analysis. Notably, we did not implement specific doublet detection algorithms or upper thresholds for UMI or gene counts. This decision was based on methodological considerations regarding doublet detection tools, which can inadvertently remove cells with intermediate or transitional phenotypes while being optimized for datasets with discrete cell types. Instead, we opted to evaluate potential doublets during later stages of analysis by examining whether marker genes from multiple distinct cell types co-occurred in specific clusters.

We normalized the filtered data using the LogNormalize method. We identified highly variable genes using the variance-stabilizing transformation (VST) approach, selecting the top 2000 genes for downstream dimensional reduction and clustering. Then, we applied scaling to these genes.

We integrated the data based on common cell types and genes between all samples, merging data from 24 samples and creating a comprehensive kidney cell dataset. We followed the standard Seurat integration workflow, but using reciprocal principal component analysis (RPCA) to identify common sources of integration between the datasets.

We then performed principal component analysis (PCA) on the integrated data for dimensionality reduction. We determined the optimal number of principal components for downstream analysis using elbow plots. We then used these principal components as input for Uniform Manifold Approximation and Projection (UMAP) to generate two-dimensional visualizations of the cell landscape. We performed clustering using Seurat’s FindClusters function, with resolution parameters adjusted to obtain two levels of granularity in cell type identification (4 and 32 levels of abstraction). To annotate the identified clusters, we employed SingleR, a package for automated cell type annotation, using the Kidney Cell Atlas as a reference [27]. This approach yielded two levels of cluster resolution: a broad classification identifying 4 major cell types (endothelial cells, immune cells, nephrons, and stroma), and a finer resolution distinguishing 32 specific cell subtypes.

4.2. Bulk data processing

We obtained gene expression profiles from the Library of Integrated Network-based Cellular Signatures (LINCS) L1000 dataset, which contains transcriptional responses to thousands of bioactive small molecules across multiple cell lines. For our analysis, we specifically focused on the kidney cell line HA1E to investigate kidney-specific drug responses, as detailed in our results section. This analysis was conducted independently of our single-cell approach. The raw expression data from LINCS was processed to address technical variability between replicates. We implemented a regex pattern matching approach to identify all replicates for each drug and collapsed them by averaging the expression values.

To identify gene expression signatures associated with nephrotoxicity in bulk RNA-seq data, we performed differential expression analysis comparing gene expression profiles between nephrotoxic and non-nephrotoxic drugs. The analysis initially identified 223 differentially expressed genes when using all cell lines, which was reduced to 100 genes when restricting the analysis to kidney-specific cell lines (HA1E). However, after collapsing replicates, no statistically significant differentially expressed genes were detected, highlighting the limitations of bulk transcriptomic approaches for nephrotoxicity prediction (Fig. 1, green block). More detail on the statistical tests and subsequent machine learning can be found in Section 4.5.

4.3. Drug response simulations validate the necessity of cell-level data

Now that we have the processed single cells, we analyze the data to see if there is a marked difference between a baseline dataset (no drug effect) and drug-affected single cells. We randomly sample 10 % of the cells from the kidney normal dataset while preserving the original cell type distributions (prevalence). Drug effects were simulated by applying varying effect sizes to the baseline gene expression values, maintaining the proportion of cell type specific response rates. Effect size, in this experiment, are values ranging from 0.001 to 10 on a log scale. These are added as a multiplier to the original data to change their values by this factor. The percentage of cells that this multiplication factor is applied to are calculated based on the cell type response rate. Cell type response rate is one of 0.8 which is the assumed standard expected rate of response of a drug, 0.95 when the response rate is much higher than usual, and 0.3 to make sure to include drugs that could have very low response rates among different cell types.

The pseudocode to calculate the altered gene expression can be found in Algorithm 1. In a randomly sampled 10 % dataset from the original single cell gene expression: For cell type A, with a prevalence of 80 %, we select response rate (0.8, 0.95, or 0.3) number of cells. To this proportion of cells, we multiply a $1 + \text{effect size}$ [0.001, 10] factor to alter the gene expression.

Algorithm 1. Simulation of Drug Affected Single-Cell data

Input: *data*: Single-Cell Gene expression dataset
cell_types: Cell types to analyze
effect_sizes: Range of effect sizes to test
response_rates: Proportion of cells that respond to the drug
Output: *results*: Simulated Gene Expression dataset for drug affected cells, Power analysis results

1 Function
 SimulateCellExpression(*data*, *cell_types*, *effect_sizes*, *response_rates*):

```

2   results ← []
3   for each combination of (cell_type, effect, rate) do
4       cells ← data filtered by cell_type
5       // Prevalence of cell_types are considered
6       baseline ← gene expressions from cells
7       // Get original expression values
8       num_responding ← response_rate × count(cell_type)
9       responding_cells ← randomly select from num_responding cells
10      // Determine which cells will respond
11      modified_expr ← baseline
12      for each responding cell i do
13          modified_expr[i] ← baseline[i] × (1 + effect)
14          // Apply formula: original × (1 + effect)
15      p_values ← run t-tests for each gene
16      significant ← count where p_values < 0.0000025
17      power ← significant/total genes
18      // Power = proportion of significant genes
19  return modified_expr, power

```

This generates a simulated single cell matrix with the same prevalence make-up as the original data, while modifying gene expression of the drug-affected cells.

To perform the power analysis on the simulated single cell data (Fig. 1, yellow block), we first pseudo-bulk it by averaging the altered gene expression values across all cell types per gene. This results in a gene expression matrix similar to what is expected in a bulk RNAseq dataset, with the cell type annotations masked. By using this data to observe the number of significant genes per effect size, and how the value increases with increasing drug effect (directly translated from the effect size value), we show that lower values of drug effects fail to provide enough statistical power to effectively exhibit patterns to measure nephrotoxicity signatures. It is also vital to note that there could be certain cell types that are highly affected by nephrotoxic drugs that are masked in bulk data, and therefore when averaged, lose their predictive power.

4.4. Mapping drug effects to single cells (Drug2Cell)

Once the cell type annotations from the integrated kidney single cell dataset were obtained, we furthered our analysis by employing Drug2Cell [35], a package designed to integrate drug-target interactions with single-cell expression data as described by Kanemaru et.al. The approach combines a cell-by-gene expression matrix with a drug-by-gene target matrix to evaluate drug-target expression at the single-cell level. Drug2Cell implements a scoring system that quantifies the potential effect of drugs on individual cells based on the expression of their target genes. In our study, we applied this package using our curated drug-target-toxicity database described above. This enables a direct translation from drug targets to a quantitative expression of how a drug affects each cell based on the marker genes of the single cell clusters and their level of expression.

To link drug targets to their nephrotoxicity information, we combined data from DrugBank and ChEMBL (approved drugs only) with the

other three data sources. We obtain 215 nephrotoxic and 641 non-nephrotoxic drugs.

We applied the Drug2Cell algorithm to calculate "drug scores" for each cell type identified in our integrated kidney cell dataset. We created a cell-by-gene expression matrix from our annotated single-cell data, built a drug-by-gene target matrix from our curated drug-target information, and computed the cell-specific drug scores by averaging the expression values of all genes targeted by each drug in each cell. This approach allowed us to predict and quantify potential drug responses across the various kidney cell populations, offering insights into cell type-specific drug sensitivities and potential nephrotoxicity mechanisms. To prepare our data for downstream analyses, we applied a logarithmic transformation to the generated drug scores. This transformation normalized the data distribution of the scores, reducing the impact of extreme values and enhancing data interpretability, allowing for easier comparison of drug effects across different scales. The resulting log-transformed cell-by-drug matrix was used in all downstream analyses.

4.5. Statistical analysis and machine learning

We performed *t*-tests with False Discovery Rate (FDR) correction to find significant differences between nephrotoxic and non-nephrotoxic drug groups. These tests were conducted using the log-transformed drug score matrix as described in Section 4.4. To generate the representative drug score for each cell type used in the statistical tests, we calculated the average of all drug scores pertaining to that specific cell type within our dataset.

Building upon these statistical findings, we extended the analysis by applying machine learning techniques to evaluate whether the cellular response signatures we identified could serve as effective predictors of a drug candidate's potential nephrotoxicity.

We implemented six different algorithms: XGBoost (Extreme Gradient Boosting), Support Vector Classifier (SVC), Logistic Regression and Random Forest, Gradient Boosting and Extra Tree. In addition to these, we trained ensemble models and an MLP. We trained these models on our processed dataset using the log-transformed drug scores in each cell type as features and the known nephrotoxicity status as the target variable.

The logistic regression model served as a baseline linear classifier. Support vector machine was implemented with a radial basis function (RBF) kernel. Tree-based methods (random forest, extra trees, XGBoost) captured non-linear relationships and feature interactions. All models were trained by splitting the original dataset into a train-test 80–20 split. Table 5 shows the performance metrics of this internal test set under the "Int" column. We used an independent test set consisting of kidney cells from diabetic and non-diabetic patients (accessible at: <https://doi.org/10.1038/s41467-022-32972-z>) The same exact processing pipeline was applied to this dataset to extract the drug scores from it. We then run the machine learning models trained on the earlier dataset to get predictions and performance metrics on the new test set as reported in the "Ext" column of Table 5.

5. Conclusion

Our study uses single-cell RNA sequencing to identify cell type-specific signatures that elucidate mechanisms of drug-induced nephrotoxicity, particularly highlighting the vulnerability of Nephrons. These signatures promise enhanced early detection and prediction of nephrotoxicity, potentially improving drug safety screening and design. Despite limitations, like the use of only healthy donor kidneys, there is potential in expanding to larger kidney datasets and incorporating diseased kidney data for extensive in-silico validation. Future work will focus on addressing these limitations and employing advanced machine learning techniques for better predictive capabilities. We will also expand this work to multiple tissue types and toxicities, with a special focus on heart

and liver cells. This research lays a strong foundation for precise strategies in mitigating adverse drug effects on the kidney, with substantial potential to impact patient care and drug development.

Author statement

The authors declare that the submitted manuscript is original and has not been published in its current form anywhere else except as a preprint (DOI: 10.1101/2025.06.17.660070). This manuscript is not under consideration for publication in any other journal or publication venue.

CRediT authorship contribution statement

Nicholas P. Tatonetti: Writing – review & editing, Supervision, Resources, Methodology, Funding acquisition, Conceptualization. **Yasaman Fatapour:** Writing – review & editing, Software, Methodology. **Jacob Berkowitz:** Writing – review & editing, Validation, Software. **Jose Miguel Acitores Cortina:** Writing – review & editing, Software, Methodology. **Hongyu Liu:** Writing – review & editing, Methodology, Investigation. **Kuchi Aditi Sharma:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used Claude, ChatGPT and Adobe Illustrator's generative AI in order to improve readability and generate part of the illustrations. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Acknowledgements

AK, JMA, HL, YF, JB and NPT are supported by National Institute of Health, USA, R35 GM131905.

Conflict of Interest Statement

The authors declare that they have no conflicts of interest.

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at [doi:10.1016/j.csbj.2025.11.012](https://doi.org/10.1016/j.csbj.2025.11.012).

Data availability

The original datasets used in our analysis can be accessed through their respective repositories: the Human Cell Atlas is available through DOI: 10.1038/s41586-023-05769-3, and the GSE131685 dataset can be accessed through the Gene Expression Omnibus (GEO) at <https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE131685>. The independent test set can be accessed here: <https://doi.org/10.1038/s41467-022-32972-z>. Code to implement this analysis can be accessed from GitHub: <https://github.com/tatonetti-lab/singlecell-nephrotoxicity> and the data required can be downloaded from Zenodo: <https://doi.org/10.5281/zenodo.15724290>.

References

- [1] Perazella MA. Pharmacology behind common drug nephrotoxicities. *Clin J Am Soc Nephrol CJASN* 2018;13:1897–908.
- [2] Minja NW, Akrahi H, Yeates K, Kilonzo KG. Acute kidney injury and associated factors in intensive care units at a tertiary hospital in Northern Tanzania. 20543581211027971 *Can J Kidney Health Dis* 2021;8. 20543581211027971.

- [3] Awdishu L, Atilano-Roque A, Tuey S, Joy MS. Identification of novel biomarkers for predicting kidney injury due to drugs using ‘omic’ strategies. *Pharm Pers Med* 2020;13:687–705.
- [4] Shi Y, Hua Y, Wang B, Zhang R, Li X. In silico prediction and insights into the structural basis of drug induced nephrotoxicity. *Front Pharm* 2022;12:793332.
- [5] Li J, et al. Ferroptosis: past, present and future. *Cell Death Dis* 2020;11:88.
- [6] Lee S, et al. Human nephrotoxicity prediction models for three types of kidney injury based on data sets of pharmacological compounds and their metabolites. *Chem Res Toxicol* 2013;26:1652–9.
- [7] Fuchs TC, Hewitt P. Biomarkers for drug-induced renal damage and nephrotoxicity—an overview for applied toxicology. *AAPS J* 2011;13:615–31.
- [8] Case J, Khan S, Khalid R, Khan A. Epidemiology of acute kidney injury in the intensive care unit. *Crit Care Res Pract* 2013;2013:479730.
- [9] Stewart BJ, Ferdinand JR, Clatworthy MR. Using single-cell technologies to map the human immune system — implications for nephrology. *Nat Rev Nephrol* 2020;16:112–28.
- [10] Muto Y, et al. Single cell transcriptional and chromatin accessibility profiling redefine cellular heterogeneity in the adult human kidney. *Nat Commun* 2021;12:2190.
- [11] Sears SM, Orwick A, Siskind LJ. Modeling cisplatin-induced kidney injury to increase translational potential. *Nephron* 2023;147:13–6.
- [12] Muhl L, et al. Single-cell analysis uncovers fibroblast heterogeneity and criteria for fibroblast and mural cell identification and discrimination. *Nat Commun* 2020;11:3953.
- [13] Kuppe C, et al. Decoding myofibroblast origins in human kidney fibrosis. *Nature* 2021;589:281–6.
- [14] Sun Y-Y, et al. scDrug+: predicting drug-responses using single-cell transcriptomics and molecular structure. *Biomed Pharm* 2024;177:117070.
- [15] Fustero-Torre C, et al. Beyondcell: targeting cancer therapeutic heterogeneity in single-cell RNA-seq data. *Genome Med* 2021;13:187.
- [16] Duan Q, et al. LINCS Canvas Browser: interactive web app to query, browse and interrogate LINCS L1000 gene expression signatures. *Nucleic Acids Res* 2014;42:W449–60.
- [17] Evangelista JE, et al. SigCom LINCS: data and metadata search engine for a million gene expression signatures. *Nucleic Acids Res* 2022;50:W697–709.
- [18] Regev, A. *et al.* The Human Cell Atlas White Paper. Preprint at <https://doi.org/10.48550/ARXIV.1810.05192> (2018).
- [19] Ryan PB, et al. Defining a reference set to support methodological research in drug safety. *Drug Saf* 2013;36:33–47.
- [20] Hunter FMI, et al. Drug safety data curation and modeling in ChEMBL: boxed warnings and withdrawn drugs. *Chem Res Toxicol* 2021;34:385–95.
- [21] Connor S, Li T, Qu Y, Roberts RA, Tong W. Generation of a drug-induced renal injury list to facilitate the development of new approach methodologies for nephrotoxicity. *Drug Discov Today* 2024;29:103938.
- [22] Soo JY-C, Jansen J, Masereeuw R, Little MH. Advances in predictive in vitro models of drug-induced nephrotoxicity. *Nat Rev Nephrol* 2018;14:378–93.
- [23] Zheng GXY, et al. Massively parallel digital transcriptional profiling of single cells. *Nat Commun* 2017;8:14049.
- [24] Hao Y, et al. Dictionary learning for integrative, multimodal and scalable single-cell analysis. *Nat Biotechnol* 2024;42:293–304.
- [25] Dvir Aran, Aaron Lun, Daniel Bunis, Jared Andrews, Friederike Dündar. Single R. *Bioconductor* <https://doi.org/10.18129/B9.BIOC.SINGLER>.
- [26] Aran D, et al. Reference-based analysis of lung single-cell sequencing reveals a transitional profibrotic macrophage. *Nat Immunol* 2019;20:163–72.
- [27] Lake BB, et al. An atlas of healthy and injured cell states and niches in the human kidney. *Nature* 2023;619:585–94.
- [28] Knox C, et al. DrugBank 6.0: the DrugBank Knowledgebase for 2024. *Nucleic Acids Res* 2024;52:D1265–75.
- [29] Zdravil B, et al. The ChEMBL Database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic Acids Res* 2024;52:D1180–92.
- [30] He M, Liu Z, Li L, Liu Y. Cell-cell communication in kidney fibrosis. *Nephrol Dial Transplant Publ Eur Dial Transpl Assoc Eur Ren Assoc* 2024;39:761–9.
- [31] Sanz AB, Sanchez-Niño MD, Ramos AM, Ortiz A. Regulated cell death pathways in kidney disease. *Nat Rev Nephrol* 2023;19:281–99.
- [32] Cui C, Cui F, Zou Q, Zhang Z, Jia L. Progress and applications of single-cell RNA sequencing and spatial transcriptome technology in acute kidney injury research. *Mol Ther Nucleic Acids* 2025;36:102583.
- [33] Ianevski A, et al. Single-cell transcriptomes identify patient-tailored therapies for selective co-inhibition of cancer clones. *Nat Commun* 2024;15:8579.
- [34] Su C, et al. Single-cell RNA sequencing in multiple pathologic types of renal cell carcinoma revealed novel potential tumor-specific markers. *Front Oncol* 2021;11:719564.
- [35] Kanemaru K, et al. Spatially resolved multiomics of human cardiac niches. *Nature* 2023;619:801–10.