Research Article

Provenance-Preserving Cross-Tissue Transcriptomic Prioritisation of Predicted Metabolite Targets Across The Diabetic Gut-Kidney Axis

Md Faiyaz[1]*, Dr. Yogesh Chand Yadav[2], Alisha Singh[3], Ravi Chandra Kumar[4], Preeti Kumari Dubey[5], Vivek Raj[6], Puja kumari[7], Dr. Amita Singh[8]

[1]PG student, Department of Pharmacology, Faculty of Pharmacy, Uttar Pradesh University of Medical Sciences, Saifai, Etawah, Uttar Pradesh, India, ORCID: 0009-0001-7948-5068, E-mail: drx.mdfaiyaz@gmail.com
[2]Associate Professor, Department of Pharmacology, Faculty of Pharmacy, Uttar Pradesh University of Medical Sciences, Saifai, Etawah, Uttar Pradesh, India, ORCID: 0000-0001-5652-1225, E-mail: drycy31@gmail.com
[3]Assistant Professor, School of Pharmacy, Sandip University, Sijoul, Madhubani, Bihar, India, ORCID: 0009-0001-2081-9676, E-mail: alishabaghel241@gmail.com
[4]Assistant Professor (Pharmacology), S.N.S College of Pharmacy, Motihari, Bihar, India, , E-mail: rckumar1734@gmail.com
[5]Research Scholar, Pharmaceutical sciences (Pharmacology), Uttaranchal University, Prem Nagar, Dehradun Uttarakhand, India, E-mail: preetidubey418pd@gmail.com
[6]Assistant Professor, School of Pharmacy, Sandip University, Sijoul, Madhubani, Bihar, India,, E-mail: vraj2535@gmail.com
[7]Research Scholar, Pharmaceutical sciences (Pharmaceutics), Uttaranchal University, Prem Nagar, Dehradun Uttarakhand, India, E-mail: pujayaduvanshi0730@gmail.com
[8]Veterinary officer & In-charge Animal House, Uttar Pradesh University of Medical Sciences, Saifai, Etawah, Uttar Pradesh, India, E-mail: Amita.vet@gmail.com
Corresponding author: Md. Faiyaz, drx.mdfaiyaz@gmail.com

Abstract

Background: Computational target-prioritisation studies can become misleading when prediction, species-specific transcriptomics, network connectivity and pathway enrichment are merged without preserving the evidential boundary of each layer. We developed a provenance-preserving workflow to prioritise predicted metabolite targets across gut and kidney transcriptomes in diabetes while retaining discordant and null findings.

Methods: An audited human-target prediction workflow was identifier-standardised before independent analysis of three public transcriptomic datasets: human jejunal enteroendocrine-enriched tissue (DS001; GSE132831), supportive mouse ileum (DS002; GSE210876) and human diabetic-nephropathy glomeruli (DS003; GSE96804). Dataset-specific differential-expression evidence was retained without cross-dataset effect-size pooling or combined P values. Direction status was assigned only when at least two datasets robustly supported a target. A predefined 31-target human gut-kidney set was examined using high-confidence STRING functional and physical network contexts and controlled GO, KEGG and Reactome over-representation analysis.

Results: Eight eligible prediction runs yielded 800 preserved records. Deduplication and component-level identifier verification resolved 449 unique approved human genes, all retained as computationally predicted targets. Robust FDR support was observed for 113 targets in DS001, 19 in supportive DS002 and 209 in DS003. Across the 449-target universe, 31 targets were direction-concordant, 22 were direction-discordant and 396 had insufficient multi-dataset support for direction assessment. The 31-target functional STRING network contained five associations connecting nine submitted seeds while 22 remained isolates; the matched physical-network sensitivity query returned no edges. No GO, KEGG or Reactome term survived the prespecified Benjamini-Hochberg FDR threshold.

Conclusions: Cross-tissue evidence can narrow a broad predicted-target universe without converting computational prediction, overlap, network connectivity or nominal enrichment into validation. The resulting target set is hypothesis-generating; sparse network structure and null FDR-controlled enrichment argue for target-level, rather than pathway-level, interpretation.

Keywords: type 2 diabetes; diabetic kidney disease; gut-kidney axis; transcriptomics; target prediction; network pharmacology; evidence integration; reproducibility.

Graphical abstract:

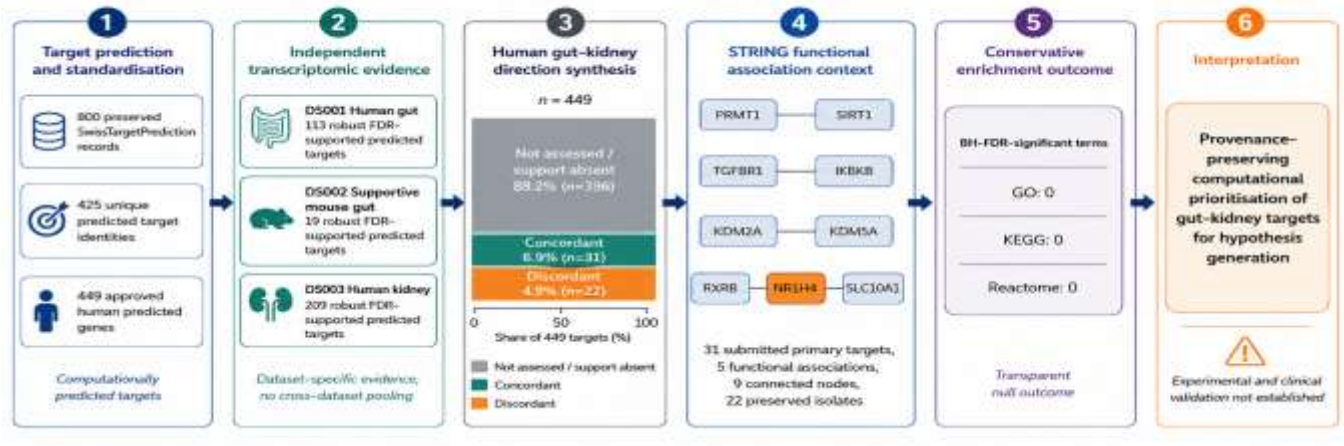

1. Introduction

Type 2 diabetes is accompanied by changes in host metabolism, mucosal physiology and intestinal microbial ecology, yet taxonomic associations are not consistent across cohorts. This heterogeneity is evident in broad reviews of the diabetes microbiome and complicates attempts to infer host mechanisms from microbiota composition alone [1].

A similar caution applies to diabetic kidney disease (DKD). A meta-analysis of gut microbiome studies in DKD found reproducible group-level trends but also substantial between-study heterogeneity, reinforcing the need to separate descriptive microbial associations from host-tissue molecular evidence [2]. More recent syntheses have likewise emphasised microbial metabolites as plausible mediators of renal injury while noting that mechanistic and clinical translation remains incomplete [3].

Compound-specific evidence can nevertheless establish a biologically credible gut-kidney link. Phenyl sulfate, for example, has been linked experimentally to podocyte injury and albuminuria and was associated with albuminuria progression in patients with diabetes [4]. Clinical intervention evidence targeting the microbiome is still mixed; a review of randomised trials of probiotics and related biotics reported signals in several renal biomarkers but insufficient consistency for a general therapeutic conclusion [5]. Contemporary reviews therefore describe the gut-kidney axis as a bidirectional, metabolite- and inflammation-linked system rather than a single linear pathway [6].

The kidney side of this axis is itself heterogeneous. Earlier compartment-resolved transcriptomic work showed marked glomerular and tubular remodelling in human DKD [7]. Single-cell and spatial transcriptomic studies have since demonstrated cell-type-specific inflammatory, endothelial, fibroblast and epithelial programmes in human diabetic kidneys [8]. These advances support the use of tissue-resolved transcriptomics for target prioritisation while also cautioning that bulk datasets mix cell populations and cannot be interpreted as a single uniform renal response [9].

Network pharmacology provides a useful framework for multi-target hypotheses, but network connectivity is not equivalent to target engagement or therapeutic efficacy [10]. Modern network approaches can integrate drug, disease and protein-association information, yet their biological meaning depends on the source of each edge, the background network and the validation strategy [11].

A recurring weakness in computational target-prioritisation studies is loss of provenance during integration. Technical replicates may be treated as independent samples, prediction outputs may be rebelled as validated targets, animal evidence may be presented as human confirmation and nominal enrichment may be promoted after correction has failed. We therefore asked a deliberately narrower question: can predicted metabolite targets be prioritised across gut and kidney transcriptomes while preserving species, biological-unit, direction, network and multiple-testing boundaries?

The present study uses an audit-driven workflow in which target prediction is kept separate from transcriptomic support, the mouse gut layer is explicitly supportive only, cross-dataset effect sizes are not pooled and negative network or enrichment outcomes remain part of the final result. The objective is not to establish a therapeutic mechanism, but to generate a transparent and testable target-prioritisation framework for subsequent structural and experimental work.

2. Materials and Methods

2.1 Study design and interpretation boundaries

This secondary computational study used a pre-existing metabolite evidence register, locked target-prediction outputs and public transcriptomic datasets. Evidence layers were kept distinct: predicted target, dataset-specific transcriptomic support, cross-dataset direction status, network association and pathway over-representation. No layer was permitted to upgrade a computational prediction to an experimentally established target. No cross-dataset meta-analytic effect size or combined P value was generated. Mouse transcriptomic evidence was retained as supportive mechanistic evidence and was never treated as human validation.

All numerical results were taken from analytical outputs that had been locked before manuscript drafting. The manuscript was written from those locked summaries; no result was recalculated to improve significance, network density or narrative coherence.

2.2 Target prediction and human-gene identifier standardisation

Small molecules eligible for structure-based target prediction were submitted individually to SwissTargetPrediction for Homo sapiens. The platform predicts likely protein targets by comparing query structures with ligands of known bioactivity [12]. In silico target prediction is best treated as a hypothesis-generation step rather than experimental evidence of binding [13], and model performance depends on the chemical and target space represented in the reference data [14]. Accordingly, complete ranked outputs were preserved rather than retaining only favourable predictions.

One very small molecule was rejected by the prediction server because it did not satisfy the input heavy-atom requirement, and a heterogeneous exopolysaccharide was excluded because it did not have a single fixed small-molecule structure. These were retained as eligibility decisions rather than interpreted as negative biology.

Across eligible runs, 800 prediction records were preserved. Target identities were deduplicated while multi-protein complexes were retained and resolved at component level for identifier verification. Approved human symbols were standardised according to HGNC nomenclature principles [15]. The final predicted-target universe contained 449 unique approved human genes. Identifier confirmation was treated as nomenclature quality control; evidence status remained Predicted.

2.3 Transcriptomic data sources and general evidence rule

Transcriptomic data and metadata were obtained from the NCBI Gene Expression Omnibus (GEO), a public archive for functional-genomics datasets [16]. Three datasets were analysed independently and were not pooled at the expression level: DS001 (human gut; GSE132831), DS002 (supportive mouse gut; GSE210876) and DS003 (human kidney; GSE96804). Dataset-specific Benjamini-Hochberg-adjusted evidence was carried into the integration layer as a robust-FDR support state together with the original effect direction where available.

2.4 DS001: human jejunal enteroendocrine-enriched transcriptomics

GSE132831 derives from a study of jejunal enteroendocrine-cell biology in obesity with and without type 2 diabetes [17]. The publication reports RNA-sequencing profiles from 14 obese participants without diabetes and 13 obese participants with diabetes. During row-level GEO reconciliation, deposited sample titles, enrichment labels and subject identifiers did not map cleanly onto that published 14/13 split; the discrepancy was preserved rather than overwritten.

The project reconstruction control contained 224 expression files corresponding to 28 standardised subjects with eight technical files per subject. Fifteen non-diabetic obese and 12 diabetic obese subjects had the required enteroendocrine-enriched context; one additional diabetic subject was represented in a non-enriched context. The primary analysis therefore used the EEC-matched 12-versus-15 biological-subject comparison. Sensitivity analyses retained all 28 subjects and an EEC-matched 26-subject population after exclusion of a prespecified diagnostic candidate. Technical files were aggregated at the biological-subject level before differential expression.

After count reconstruction and quality control, differential expression was analysed with an edgeR quasi-likelihood workflow [18]. Expression filtering and TMM normalisation preceded model fitting. The primary evidence threshold was Benjamini-Hochberg FDR <0.05; a project-defined priority layer additionally required an absolute log2 fold change of at least 1. No sample was automatically removed on the basis of a single quality-control metric.

2.5 DS002: supportive mouse ileum transcriptomics

GSE210876 was generated in a diabetic/NASH mouse study examining intestinal insulin action [19]. The deposited ileum expression files comprised four untreated control-genotype mice, six insulin-treated control-genotype mice and six insulin-treated intestinal epithelial insulin-receptor knockout mice. Two contrasts were retained: insulin-treated versus untreated control genotype (6 versus 4) and knockout versus insulin-treated control genotype (6 versus 6).

The deposited values were processed integer-valued expression measurements and were not independently verified as unnormalised raw counts. They were therefore not analysed with raw-count methods. Genes detected in at least 4 of 16 samples were retained, values were transformed as log2(expression+1), and robust limma-trend modelling with empirical-Bayes moderation was used [20]. The primary threshold was FDR <0.05; the project-defined priority threshold additionally required |log2 fold change| >=0.584963. Leave-one-sample-out analyses were retained as supportive sensitivity checks.

Mouse genes were mapped to approved human orthologues for comparison with the locked 449-gene human target universe. Ensembl resources were used for cross-species identifier resolution [21]. Mapping exceptions and one-to-many relationships were audited rather than silently reassigned. The mouse layer remained supportive only.

2.6 DS003: human glomerular diabetic-nephropathy transcriptomics

GSE96804 contains microdissected glomerular expression profiles from 41 patients with diabetic nephropathy and 20 controls [22]. The controlled workflow retained all 61 samples. Following array preprocessing and approved-gene mapping, a primary limma model and a mandatory sensitivity model were evaluated independently. A target was considered robustly supported only when the locked criteria were met in both model layers.

Validated batch and clinical covariates sufficient for a stable adjusted model were not available in the deposited metadata used for this project. Covariates were therefore not invented. This limitation was preserved in the final analytical lock and carried into interpretation.

2.7 Cross-dataset evidence integration and direction status

Multiple testing within each transcriptomic analysis used the Benjamini-Hochberg false-discovery-rate procedure [23]. Integration occurred only after each dataset had reached its final lock. For every one of the 449 predicted genes, the synthesis matrix retained

presence in the expression universe, dataset-specific FDR support, priority status and effect direction. Effect sizes were not pooled across datasets.

Direction status was calculated only for targets with robust FDR support in at least two datasets. If all supported non-zero effect estimates had the same sign, the target was labelled concordant; opposing signs produced a discordant label. Targets supported in fewer than two datasets were labelled not assessed for direction. Discordance triggered review but did not trigger automatic exclusion.

A predefined 31-target human gut-kidney set was carried forward as the primary network and over-representation query. A separate 36-target mechanistic-context set and a 26-target direction-review set were preserved for sensitivity and annotation. Set membership was descriptive and did not constitute biological rank, target engagement or efficacy.

2.8 STRING network analysis

The 31 primary human gut-kidney targets were mapped to STRING v12.0, which integrates physical interactions and broader functional protein associations [24]. Homo sapiens (taxon 9606) was used with a required interaction score of 700 and zero additional interactors. Two contexts were retained separately: a primary functional-association network and a strict physical-network sensitivity query. Disconnected submitted seeds were preserved in the node register rather than removed to increase apparent network density.

2.9 Over-representation analysis

Controlled over-representation analysis used the 31 primary targets as the main query against the locked 449-target background. The 36-target mechanistic-context set was tested separately against the same background. Gene Ontology annotations were used for functional categories [25]. KEGG was used for pathway annotation [26]. Reactome was used as an independent pathway knowledgebase [27].

The prespecified significance threshold was Benjamini-Hochberg FDR <0.05 with gene-set size limits of 10 to 500. Nominally enriched but FDR-non-significant terms were not promoted into the results narrative.

2.10 Reproducibility, provenance and claim control

Each analytical stage generated a machine-readable output register, read-back audit and integrity information before being treated as final. Upstream locked files were not overwritten during later synthesis. The final 449-target evidence matrix retained explicit fields for dataset support, direction status, predefined-set membership, network context and interpretation boundaries. Manuscript claims were cross-referenced to an evidence register containing internal source files, external references, DOI links and claim limitations.

3. Results

3.1 Audited target prediction produced a 449-gene human target universe

Eligible target-prediction runs produced 800 preserved SwissTargetPrediction records. After deduplication there were 425 unique predicted target identities: 386 single-protein identities and 39 multi-protein complexes. Component-level review of the complexes yielded 66 unique components. Reconciliation across single proteins and complex components produced 449 unique approved human genes. No identifier correction converted a prediction into a known target, and the final evidence status of all 449 genes remained Predicted.

Prediction rows, original target names, ranks, probabilities and classes were retained in the upstream register, whereas the 449-gene list served only as the standardised human background for transcriptomic comparison. This separation prevented nomenclature verification from being misrepresented as independent biological confirmation.

3.2 Independent gut and kidney datasets supported different fractions of the predicted-target universe

DS001 human gut analysis tested 16,769 genes after subject-level reconstruction and filtering. The primary workflow identified 5,867 FDR-significant genes and 576 priority genes. Exact approved-gene mapping yielded 331 predicted-target overlaps, of which 113 met the final robust FDR-support rule and 8 met the priority rule.

DS002 supportive mouse-gut analysis retained 19,845 genes after the prespecified expression filter. Orthologue reconciliation produced 417 approved human genes from the relevant mouse differential-expression evidence and 25 unique exact overlaps with the 449 predicted targets across the two contrast-provenance layers. Nineteen targets met the final robust FDR-support rule used in the cross-dataset synthesis. These observations remained supportive mouse evidence and were not counted as human replication.

DS003 human glomerular analysis contained 70,523 expression features and 22,674 approved-gene rows after mapping. Of the 449 predicted targets, 444 were represented in the expression universe and 5 were absent. Two hundred and nine targets satisfied

the robust FDR-support rule; 20 also met the robust priority rule. Among the 209 robust targets, 90 were upregulated and 119 were downregulated. The dataset-specific robust FDR-supported target counts are summarised in Figure 1 and Table 1.

Figure 1. Dataset-specific robust FDR-supported predicted targets.

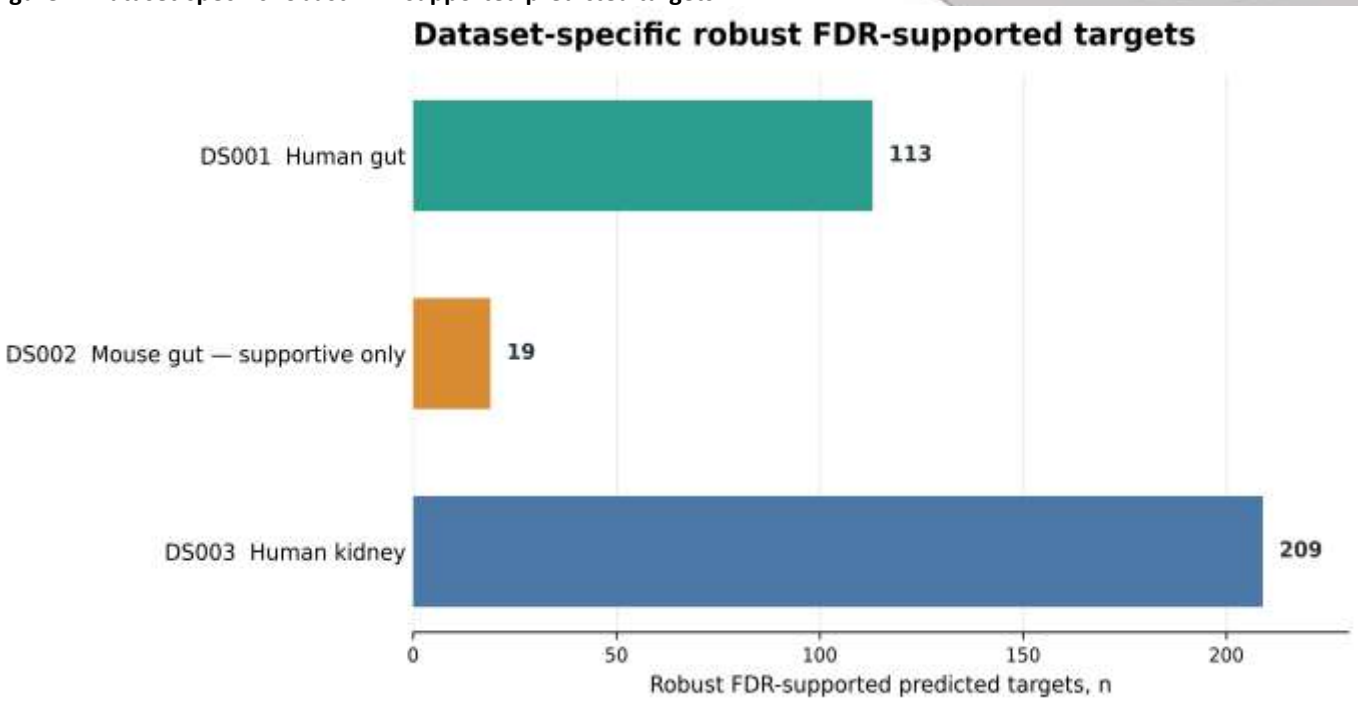

Notes: Robust target-level support is shown separately for DS001 human gut (113), DS002 supportive mouse gut (19) and DS003 human kidney (209). Counts are descriptive and are not pooled across datasets; the mouse layer is supportive only and is not human validation.

Table 1. Transcriptomic datasets and target-level evidence retained in the final synthesis

Dataset	Biological context	Controlled comparison / analysis unit	Statistical framework	Locked robust FDR-supported predicted targets
DS001 / GSE132831	Human jejunal enteroendocrine-enriched context	12 diabetic obese vs 15 non-diabetic obese subjects; sensitivity populations retained	Subject-level edgeR quasi-likelihood; BH FDR	113
DS002 / GSE210876	Supportive mouse ileum	csp vs csn (6 vs 4) and ksp vs csp (6 vs 6)	Robust limma-trend on processed expression values; BH FDR	19
DS003 / GSE96804	Human diabetic-nephropathy glomeruli	41 diabetic nephropathy vs 20 controls	Primary + sensitivity limma models; BH FDR	209

Note: Counts are dataset-specific and are not additive because the same predicted target can be supported in more than one dataset. DS002 is supportive mouse evidence and is not human validation. Source publications: DS001 [17], DS002 [19], DS003 [22].

3.3 Direction-aware synthesis retained concordance, discordance and non-assessable targets

Across the 449-target universe, 53 targets had robust FDR support in at least two datasets and were therefore eligible for direction assessment. Thirty-one were direction-concordant and 22 were direction-discordant. The remaining 396 targets had robust

support in fewer than two datasets and were labelled not assessed rather than assigned an uncertain cross-dataset direction. The distribution is shown in Figure 2.

Figure 2. Human gut-kidney direction-status distribution across the locked 449-target universe.

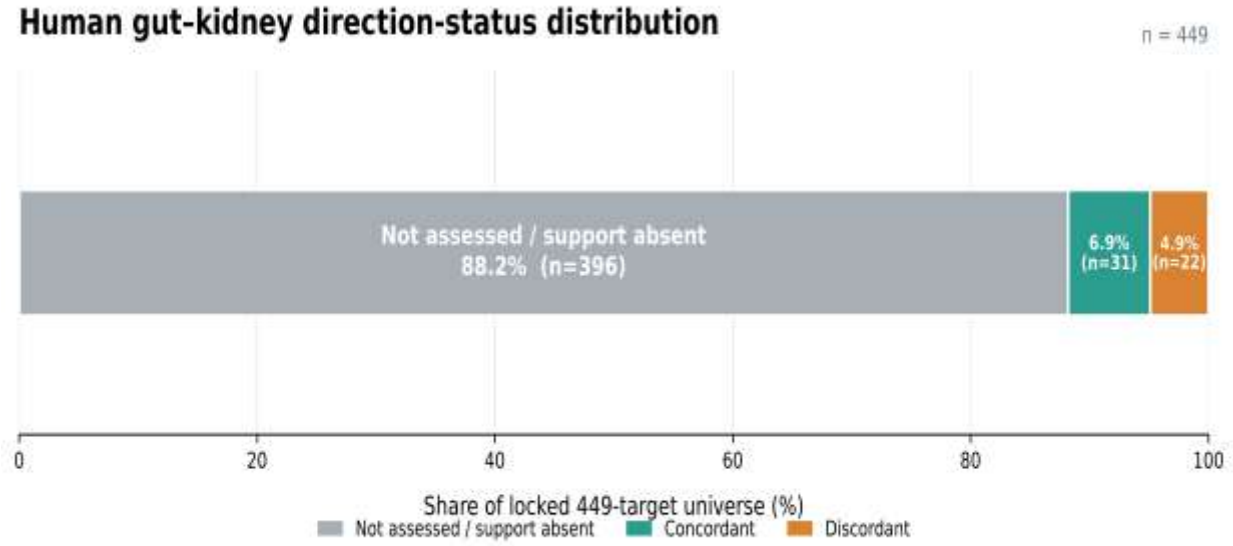

Note: Thirty-one targets were direction-concordant, 22 were direction-discordant and 396 were not assessed because fewer than two datasets provided robust support. Direction status is descriptive and is not a pooled effect estimate.
Direction concordance was not used as a therapeutic ranking, and discordance did not trigger automatic exclusion. This retained tissue- and context-specific responses that would have been hidden by an overlap-only summary. The final synthesis preserved a 31-target primary human gut-kidney set, a 36-target mechanistic-context set and a 26-target direction-review set. Their predefined sizes are shown in Supplementary Figure S1.

Supplementary Figure S1. Predefined analytical target sets.

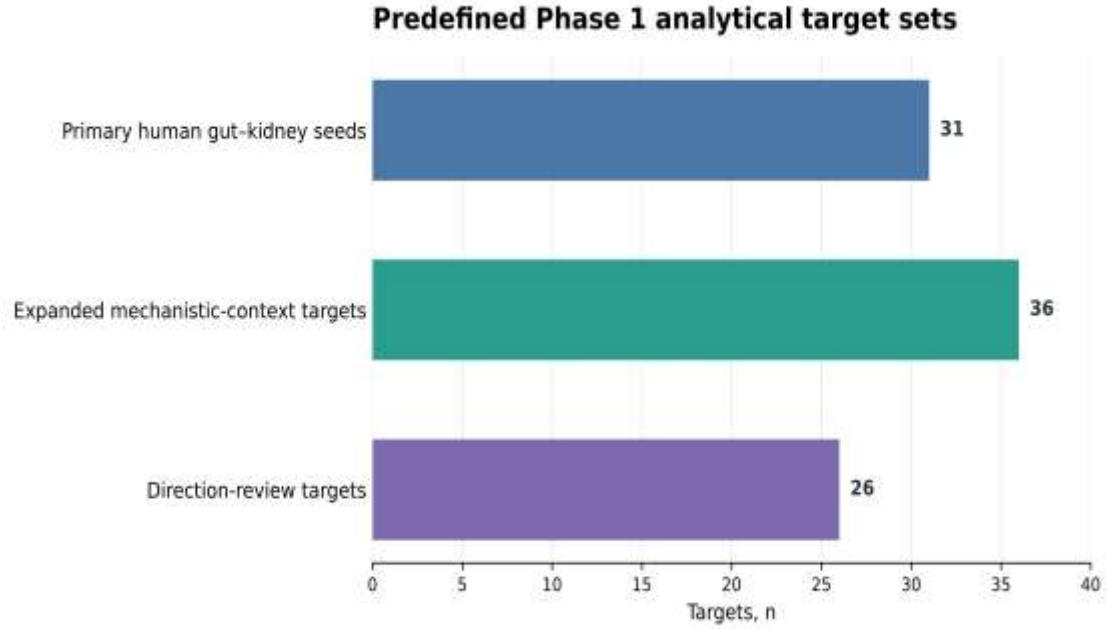

Note: The primary human gut-kidney set (31), expanded mechanistic-context set (36) and direction-review set (26) are controlled analytical subsets; membership does not establish biological rank or efficacy.

3.4 High-confidence STRING analysis revealed sparse functional connectivity and no physical-network edges

At the locked STRING score threshold, the functional-association query retained all 31 submitted primary seeds but returned only 5 associations. Nine seeds participated in at least one association and 22 remained isolates. The connected portion comprised PRMT1-SIRT1, TGFBR1-IKBKB, KDM2A-KDM5A and an RXRB-NR1H4-SLC10A1 component. These are STRING functional associations and were not interpreted as direct binding or regulatory causality. The connected-only nine-node/five-edge view is shown in Figure 3; the 22 isolates remain part of the full submitted network.

Figure 3. Connected component of the primary human gut-kidney STRING functional-association network.

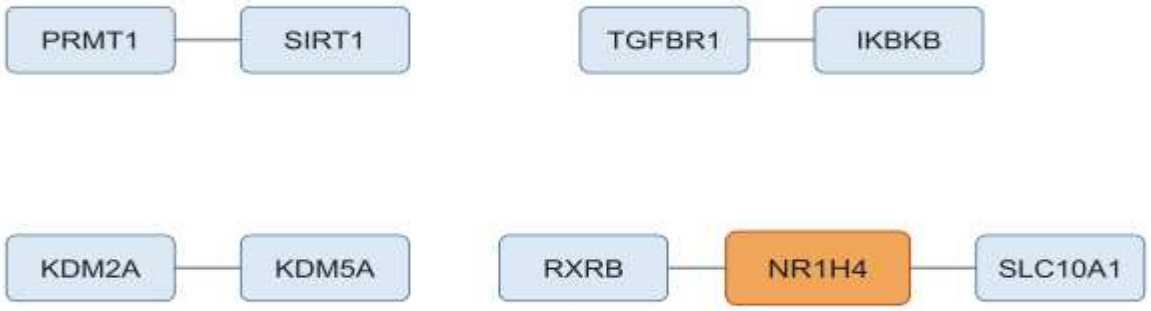

Note: The panel shows the nine connected seeds and five retained functional associations, with NR1H4 highlighted for visual orientation only. The 22 submitted isolates are not displayed in the connected-only panel but remain part of the full 31-node network. Edges represent STRING functional associations and do not establish direct binding, causality or therapeutic efficacy.

The matched physical-network sensitivity query retained the same 31 submitted seeds but returned zero edges. All 31 nodes therefore remained physical-network isolates. The null physical-network result was retained as a threshold-specific observation rather than used to argue universal absence of physical interaction. Functional and physical-network counts are compared in Supplementary Figure S2.

Supplementary Figure S2. Functional and physical-network contexts for the 31 primary submitted seeds.

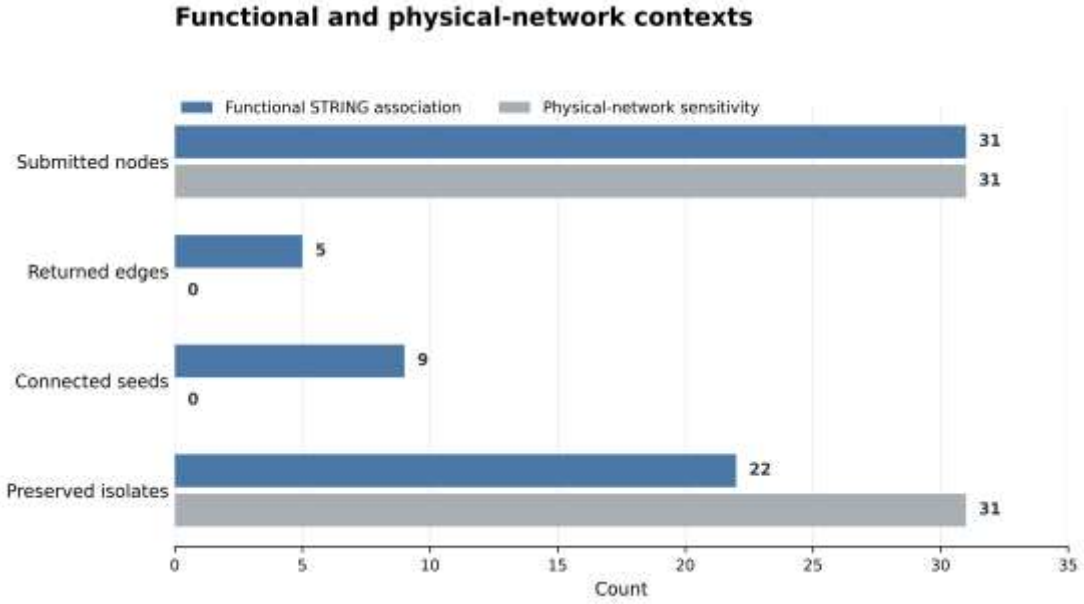

Note: The functional STRING context retained five edges, nine connected seeds and 22 isolates; the matched physical-network sensitivity query retained 31 nodes with zero edges and 31 isolates. The zero-edge physical result is threshold-specific.

3.5 No GO, KEGG or Reactome term survived FDR correction

The controlled over-representation analysis returned zero BH-FDR-significant terms for Gene Ontology, zero for KEGG and zero for Reactome. The null result was preserved in the final package. Nominal or non-significant pathway labels were not substituted for statistically supported enrichment. The locked outcome is displayed in Supplementary Figure S3.

Supplementary Figure S3. Over-representation analysis: locked BH-FDR outcome.

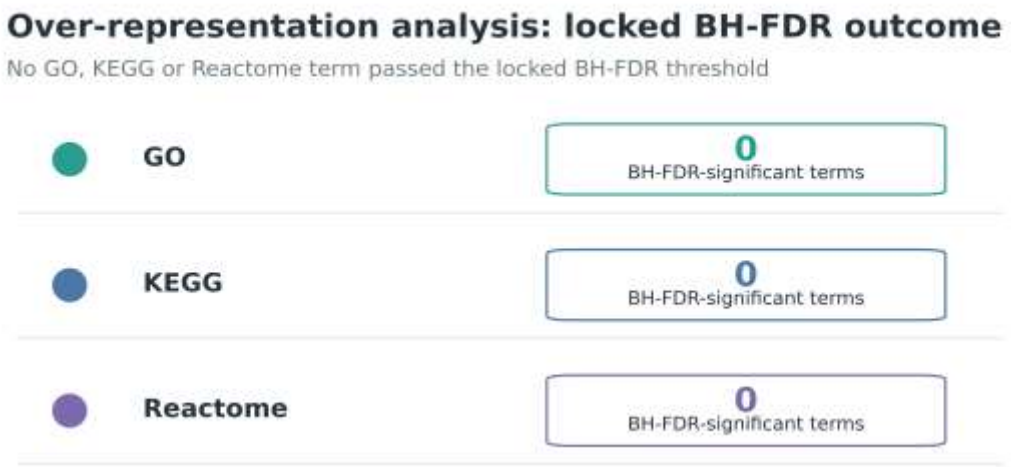

Note: No GO, KEGG or Reactome term passed the prespecified Benjamini-Hochberg FDR threshold. Non-significant or nominal terms were not promoted as pathway evidence.

The evidence architecture and permitted interpretation of each analytical layer are summarised in Table 2. The principal result is therefore a target-level prioritisation structure with sparse network context, not a pathway-enrichment claim.

Table 2. Evidence architecture and interpretation boundaries

Evidence layer	Locked result	Permitted interpretation	Explicitly not established
Predicted target universe	449 approved human genes; all status = Predicted	Standardised background for evidence integration	Known target status or experimental target engagement
Predefined analytical sets	Primary 31; mechanistic context 36; direction review 26	Controlled hypothesis-generating subsets	Therapeutic ranking or causal hierarchy
Direction synthesis	31 concordant; 22 discordant; 396 not assessed (<2 supported datasets)	Cross-dataset direction descriptor where evaluable	Meta-analytic effect or automatic exclusion
Functional STRING	31 nodes; 5 edges; 9 connected; 22 isolates	High-confidence functional association context	Direct binding, master-regulator or causal hub status
Physical STRING sensitivity	31 nodes; 0 edges; 31 isolates	Threshold-specific null sensitivity result	Universal absence of physical interactions
ORA	GO 0; KEGG 0; Reactome 0 BH-FDR-significant terms	Transparent null pathway result	Significant pathway mechanism

4. Discussion

This study demonstrates that predicted metabolite targets can be narrowed by cross-tissue transcriptomic evidence without forcing heterogeneous datasets into a single pooled statistic. The central result is not a single hub or pathway. Instead, it is an

evidence architecture in which 449 predicted human genes were compared independently with human gut, supportive mouse gut and human kidney data; direction was evaluated only when sufficient dataset support existed; and sparse or null network and enrichment outcomes were retained.

The human gut component is relevant because the source study identified a distinct jejunal enteroendocrine-cell transcriptional programme in obesity with type 2 diabetes [17]. Our analysis did not simply reuse the publication's group counts. The GEO files were re-reconciled at the biological-subject level, exposing a mismatch between the published cohort description and the enrichment/subject structure of the deposited files. Aggregating technical files before inference and restricting the primary comparison to the EEC-matched subjects reduce the risk of treating technical measurements as independent biological observations, but the resulting DS001 estimates should be understood as specific to the controlled reconstruction.

The kidney signal was broader, with 209 predicted targets meeting the robust FDR-support rule in DS003. This is consistent with the extensive transcriptional remodelling previously described in glomerular DKD [22] and with more recent cell-resolved evidence that diabetic kidneys contain multiple spatially and transcriptionally distinct disease programmes [8]. The larger kidney count should not, however, be interpreted as evidence that kidney biology is more important than gut biology. The datasets differ in platform, tissue composition, sample size and statistical design, which is exactly why the workflow did not pool their effect estimates.

The mouse ileum layer contributed a third mechanistic context rather than a replication cohort. The source experiment manipulated intestinal insulin signalling in a diabetic/NASH model [19]. Orthologue-mapped mouse expression therefore supplies perturbational information, but species and model differences remain substantial. The 19 robustly supported targets from DS002 were consequently retained as supportive evidence only.

Direction-aware integration added information that a Venn overlap would miss. Twenty-two targets were supported in at least two datasets but changed in opposing directions. Such discordance can arise from tissue-specific regulation, species differences, cellular composition, disease stage or model-specific perturbations. Automatic exclusion would discard potentially meaningful context dependence, whereas treating all overlaps as uniformly dysregulated would overstate coherence. The review category is therefore an intentional uncertainty label rather than a failed result.

The network results also caution against dense-hub narratives. STRING functional associations are broader than direct physical interactions [24], and network-based drug/disease models are sensitive to the underlying relation types and background graph [11]. At the prespecified high-confidence threshold and without added interactors, only five functional associations were returned among the 31 primary seeds, while 22 seeds remained isolated. Removing those isolates from the evidence register would have exaggerated connectedness. The strict physical-network query returned no edges, which is a threshold-specific null observation rather than evidence that no physical interactions exist in biology.

The zero-FDR enrichment result is equally important. Over-representation analysis depends on the query, background universe, database composition and multiplicity correction; seemingly small choices can alter significance [28]. Practical guidance therefore recommends predefining the background and avoiding post hoc promotion of nominal terms [29]. Benchmark reviews further show that pathway definitions, gene-set collections and analytical assumptions can produce divergent results even from the same gene list [30]. In this context, the absence of FDR-significant GO, KEGG or Reactome terms means only that the locked 31-target query did not support a statistically robust pathway-level conclusion under the prespecified analysis. It does not mean that the genes lack biological functions.

The study's emphasis on provenance is also substantive. Reproducibility in bioinformatics depends not only on rerunning software but on preserving enough information to verify the biological interpretation of the output [31]. Frameworks for FAIR and reproducible analytics likewise emphasise machine-readable provenance for data, software and computational results [32]. Here, locked source files, read-back audits, explicit claim boundaries and a manuscript evidence register were used to reduce post hoc drift between analysis and reporting.

Finally, computational target prediction has an inherent applicability domain. Comparative evaluations of drug-target interaction methods show that performance can degrade for compounds or targets that are poorly represented in training knowledge [33]. That limitation reinforces the central reporting rule used here: the 449 genes remain predicted targets even when they are supported by transcriptomic evidence or connected in STRING. The present analysis prioritises hypotheses for later structural or experimental testing; it does not substitute for target-engagement experiments.

The contribution is therefore methodological rather than an absolute first-ever claim. The workflow combines complete target-prediction provenance, component-level identifier quality control, biological-unit-aware human gut analysis, species-aware supportive mouse mapping, human kidney transcriptomics, non-pooled direction logic, high-confidence network sensitivity and transparent null enrichment. This combination produces a smaller and more defensible hypothesis space while keeping uncertainty visible.

5. Strengths and limitations

The principal strengths are preservation of the complete prediction record; component-level handling of predicted protein complexes; approved-gene standardisation; subject-level treatment of technical files; independent dataset-specific statistics; explicit species separation; no cross-dataset effect pooling; preservation of isolates and discordant targets; and retention of zero-FDR pathway outcomes. The use of locked outputs and a claim-evidence register further reduces the opportunity for result drift during manuscript preparation.

Several limitations constrain interpretation. First, the 449 genes are computational predictions and do not demonstrate direct metabolite-protein binding. Second, DS001 contained a documented mismatch between the published group description and row-level deposited metadata, so the primary population is a controlled reconstruction rather than a verbatim reproduction of the publication's 14-versus-13 cohort. Third, DS002 used processed expression values rather than independently verified unnormalised raw counts, limiting modelling choices. Fourth, mouse-to-human orthologue mapping cannot remove species-specific regulation. Fifth, validated covariates sufficient for adjusted modelling were unavailable for DS003. Sixth, the three datasets differ in platform, tissue context and design and were therefore unsuitable for naive effect-size pooling. Finally, no prospective experimental target-engagement or functional validation is reported in this paper.

6. Conclusion

A conservative evidence-integration workflow prioritised predicted metabolite targets across human gut and kidney transcriptomes while preserving supportive mouse evidence, directional discordance, sparse network structure and null pathway enrichment. The analysis identifies target-level hypotheses rather than a validated therapeutic mechanism. That distinction is the practical value of the framework: it reduces a broad computational target universe without converting prediction, overlap or network connectivity into evidence that the data do not contain.

Data availability

The transcriptomic datasets analysed in this study are publicly available in NCBI GEO under accession numbers GSE132831, GSE210876 and GSE96804. The submission package includes a claim-evidence register containing DOI, PubMed/PMC, GEO/database and locked internal-source links. The full 449-target evidence matrix, 31-target primary table and provenance records should be supplied as supplementary data from the final locked project package.

Ethics statement

No new human participants or experimental animals were enrolled, sampled or handled for this secondary computational study. Ethics and animal-governance information for the source datasets is reported in the corresponding publications [17], [19] and [22].

References

1. Gurung M, Li Z, You H, Rodrigues R, Jump DB, Morgun A, et al. Role of gut microbiota in type 2 diabetes pathophysiology. EBioMedicine. 2020;51:102590. doi:10.1016/j.ebiom.2019.11.051.
2. Han S, Chen M, Cheng P, Zhang Z, Lu Y, Xu Y, et al. A systematic review and meta-analysis of gut microbiota in diabetic kidney disease: comparisons with diabetes mellitus, non-diabetic kidney disease, and healthy individuals. Front Endocrinol (Lausanne). 2022;13:1018093. doi:10.3389/fendo.2022.1018093.
3. Yu JX, Chen X, Zang SG, Chen X, Wu YY, Wu LP, et al. Gut microbiota microbial metabolites in diabetic nephropathy patients: far to go. Front Cell Infect Microbiol. 2024;14:1359432. doi:10.3389/fcimb.2024.1359432.
4. Kikuchi K, Saigusa D, Kanemitsu Y, Matsumoto Y, Thanai P, Suzuki N, et al. Gut microbiome-derived phenyl sulfate contributes to albuminuria in diabetic kidney disease. Nat Commun. 2019;10(1):1835. doi:10.1038/s41467-019-09735-4.
5. Paul P, Kaul R, Chaari A. Renal health improvement in diabetes through microbiome modulation of the gut-kidney axis with biotics: a systematic and narrative review of randomized controlled trials. Int J Mol Sci. 2022;23(23):14838. doi:10.3390/ijms232314838.
6. Chu C, Behera TR, Huang Y, Qiu W, Chen J, Shen Q. Research progress of gut microbiome and diabetic nephropathy. Front Med (Lausanne). 2024;11:1490314. doi:10.3389/fmed.2024.1490314.
7. Woroniecka KI, Park ASD, Mohtat D, Thomas DB, Pullman JM, Susztak K. Transcriptome analysis of human diabetic kidney disease. Diabetes. 2011;60(9):2354-2369. doi:10.2337/db10-1181.
8. Chen D, Shao M, Song Y, Ren G, Guo F, Fan X, et al. Single-cell RNA-seq with spatial transcriptomics to create an atlas of human diabetic kidney disease. FASEB J. 2023;37(6):e22938. doi:10.1096/fj.202202013RR.
9. Mao ZH, Gao ZX, Liu Y, Liu DW, Liu ZS, Wu P. Single-cell transcriptomics: a new tool for studying diabetic kidney disease. Front Physiol. 2023;13:1053850. doi:10.3389/fphys.2022.1053850.

10. Hopkins AL. Network pharmacology: the next paradigm in drug discovery. Nat Chem Biol. 2008;4(11):682-690. doi:10.1038/nchembio.118.
11. Rintala TJ, Ghosh A, Fortino V. Network approaches for modeling the effect of drugs and diseases. Brief Bioinform. 2022;23(4):bbac229. doi:10.1093/bib/bbac229.
12. Daina A, Michielin O, Zoete V. SwissTargetPrediction: updated data and new features for efficient prediction of protein targets of small molecules. Nucleic Acids Res. 2019;47(W1):W357-W364. doi:10.1093/nar/gkz382.
13. Byrne R, Schneider G. In silico target prediction for small molecules. Methods Mol Biol. 2019;1888:273-309. doi:10.1007/978-1-4939-8891-4_16.
14. Mathai N, Chen Y, Kirchmair J. Validation strategies for target prediction methods. Brief Bioinform. 2020;21(3):791-802. doi:10.1093/bib/bbz026.
15. Bruford EA, Braschi B, Denny P, Jones TEM, Seal RL, Tweedie S. Guidelines for human gene nomenclature. Nat Genet. 2020;52(8):754-758. doi:10.1038/s41588-020-0669-3.
16. Barrett T, Wilhite SE, Ledoux P, Evangelista C, Kim IF, Tomashevsky M, et al. NCBI GEO: archive for functional genomics data sets-update. Nucleic Acids Res. 2013;41(Database issue):D991-D995. doi:10.1093/nar/gks1193.
17. Osinski C, Le Gleau L, Poitou C, de Toro-Martin J, Genser L, Fradet M, et al. Type 2 diabetes is associated with impaired jejunal enteroendocrine GLP-1 cell lineage in human obesity. Int J Obes (Lond). 2021;45(1):170-183. doi:10.1038/s41366-020-00694-1.
18. Robinson MD, McCarthy DJ, Smyth GK. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. Bioinformatics. 2010;26(1):139-140. doi:10.1093/bioinformatics/btp616.
19. Soeda K, Sasako T, Enooku K, Kubota N, Kobayashi N, Ikushima YM, et al. Gut insulin action protects from hepatocarcinogenesis in diabetic mice comorbid with nonalcoholic steatohepatitis. Nat Commun. 2023;14(1):6584. doi:10.1038/s41467-023-42334-y.
20. Ritchie ME, Phipson B, Wu D, Hu Y, Law CW, Shi W, et al. limma powers differential expression analyses for RNA-sequencing and microarray studies. Nucleic Acids Res. 2015;43(7):e47. doi:10.1093/nar/gkv007.
21. Harrison PW, Amode MR, Austine-Orimoloye O, Azov AG, Barba M, Barnes I, et al. Ensembl 2024. Nucleic Acids Res. 2024;52(D1):D891-D899. doi:10.1093/nar/gkad1049.
22. Pan Y, Jiang S, Hou Q, Qiu D, Shi J, Wang L, et al. Dissection of glomerular transcriptional profile in patients with diabetic nephropathy: SRGAP2a protects podocyte structure and function. Diabetes. 2018;67(4):717-730. doi:10.2337/db17-0755.
23. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. J R Stat Soc Series B Stat Methodol. 1995;57(1):289-300. doi:10.1111/j.2517-6161.1995.tb02031.x.
24. Szklarczyk D, Kirsch R, Koutrouli M, Nastou K, Mehryary F, Hachilif R, et al. The STRING database in 2023: protein-protein association networks and functional enrichment analyses for any sequenced genome of interest. Nucleic Acids Res. 2023;51(D1):D638-D646. doi:10.1093/nar/gkac1000.
25. Gene Ontology Consortium. The Gene Ontology knowledgebase in 2023. Genetics. 2023;224(1):iyad031. doi:10.1093/genetics/iyad031.
26. Kanehisa M, Furumichi M, Sato Y, Kawashima M, Ishiguro-Watanabe M. KEGG for taxonomy-based analysis of pathways and genomes. Nucleic Acids Res. 2023;51(D1):D587-D592. doi:10.1093/nar/gkac963.
27. Milacic M, Beavers D, Conley P, Gong C, Gillespie M, Griss J, et al. The Reactome Pathway Knowledgebase 2024. Nucleic Acids Res. 2024;52(D1):D672-D678. doi:10.1093/nar/gkad1025.
28. Li W, Shih A, Freudenberg-Hua Y, Fury W, Yang Y. Beyond standard pipeline and $p < 0.05$ in pathway enrichment analyses. Comput Biol Chem. 2021;92:107455. doi:10.1016/j.compbiolchem.2021.107455.
29. Chicco D, Agapito G. Nine quick tips for pathway enrichment analysis. PLoS Comput Biol. 2022;18(8):e1010348. doi:10.1371/journal.pcbi.1010348.
30. Mubeen S, Kodamullil AT, Hofmann-Apitius M, Domingo-Fernandez D. On the influence of several factors on pathway enrichment analysis. Brief Bioinform. 2022;23(3):bbac143. doi:10.1093/bib/bbac143.
31. Suetake H, Fukusato T, Igarashi T, Ohta T. A workflow reproducibility scale for automatic validation of biological interpretation results. Gigascience. 2023;12:giad031. doi:10.1093/gigascience/giad031.
32. Levinson MA, Niestroy J, Al Manir S, Fairchild K, Lake DE, Moorman JR, et al. FAIRSCAPE: a framework for FAIR and reproducible biomedical analytics. Neuroinformatics. 2022;20(1):187-202. doi:10.1007/s12021-021-09529-4.
33. Jung YS, Kim Y, Cho YR. Comparative analysis of network-based approaches and machine learning algorithms for predicting drug-target interactions. Methods. 2022;198:19-31. doi:10.1016/j.ymeth.2021.10.007.