Pharmacophore

ISSN-2229-5402

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

DRUG REPURPOSING AS EVIDENCE RECONCILIATION ACROSS DISEASE SIGNATURES, MOLECULAR MECHANISMS, PRIOR FAILURES, AND CLINICAL CONTEXT

Rui Costa[1]*, Emily Watson[2], Karim Benziane[3], Andre Silva[4]

1. ***Department of Drug Repurposing and Evidence Reconciliation, Faculty of Pharmacy, University of Coimbra, Coimbra, Portugal.***
2. ***Department of Disease Signatures and Molecular Mechanisms, Faculty of Pharmacy, University of Glasgow, Glasgow, United Kingdom.***
3. ***Department of Prior Failure Analysis and Clinical Context, Faculty of Sciences, University of Algiers, Algiers, Algeria.***
4. ***Department of Repurposing Decision Support, Faculty of Pharmacy, University of Porto, Porto, Portugal.***

ARTICLE INFO

Received:
12 November 2024
Received in revised form:
19 February 2025
Accepted:
22 February 2025
Available online:
28 February 2025

Keywords: Drug repurposing, Evidence reconciliation, Disease signatures, Molecular mechanisms, Negative evidence, Causal proximity

ABSTRACT

Drug repurposing promises to shorten parts of therapeutic development by redirecting existing compounds toward new disease indications, yet computationally attractive candidates often receive inconsistent support across transcriptomic signatures, molecular networks, genetics, phenotypic assays, observational data, and prior clinical studies. The unresolved problem is not simply insufficient prediction accuracy. It is the absence of a disciplined method for interpreting evidence sources that differ in biological scale, causal proximity, directionality, context, uncertainty, and translational authority. This article develops a proposed therapeutic evidence-synthesis model in which repurposing is treated as evidence reconciliation rather than score aggregation. The model distinguishes disease-state representation, target-mechanism plausibility, phenotypic response, human causal evidence, prior clinical outcomes, and context transfer as separate evidentiary layers. It further proposes that candidate prioritization should depend on compatibility among these layers, explicit accounting for contradictory and negative findings, and staged decision gates that prevent molecular plausibility from being mistaken for therapeutic readiness. The synthesis emphasizes that supportive evidence can converge without being equivalent, while apparent disagreement may reflect differences in cell type, dose, timing, population, endpoint, disease stage, modality, or mechanism. The proposed contribution is conceptual and requires retrospective and prospective evaluation. It cannot establish efficacy, safety, regulatory acceptability, or deployment readiness for any candidate. Its principal implication is that repurposing systems should expose why evidence agrees or conflicts, how close each signal lies to the intended therapeutic mechanism, and which uncertainties must be resolved before a hypothesis advances toward experimental or clinical testing.

To Cite This Article: Costa R, Watson E, Benziane K, Silva A. Drug Repurposing as Evidence Reconciliation across Disease Signatures, Molecular Mechanisms, Prior Failures, and Clinical Context. Pharmacophore. 2025;16(1):103-11. https://doi.org/10.51847/49pN9Gi36y

Introduction

Drug repurposing occupies an unusual position in pharmaceutical science because its apparent efficiency derives from partial prior knowledge rather than from evidentiary completeness. Existing compounds may already have characterized chemistry, manufacturing routes, pharmacokinetic information, safety observations, or clinical experience, but these assets do not automatically establish that the same intervention will produce a therapeutically useful effect in a new disease. Repurposing therefore changes the starting point of development without eliminating the need to establish biological relevance, appropriate exposure, mechanistic direction, patient suitability, and clinical benefit. Contemporary repurposing research spans knowledge-based reasoning, omics-driven signature matching, network medicine, human genetics, phenotypic screening, real-world data, and clinical investigation, each of which captures a different portion of the therapeutic argument [1].

Corresponding Author: Rui Costa; Department of Drug Repurposing and Evidence Reconciliation, Faculty of Pharmacy, University of Coimbra, Coimbra, Portugal. E-mail: rui.costa@uc.pt.

Computational methods have expanded the number and diversity of repurposing hypotheses that can be generated. A drug may be prioritized because it reverses a disease-associated expression profile, lies near a disease module in a molecular network, connects to a target through a knowledge graph, phenocopies a desired cellular state, shares a mechanism with an effective therapy, or is associated with improved outcomes in observational data. These signals can be valuable, but they are not interchangeable. Their predictive outputs may refer to molecular opposition, relational proximity, statistical association, phenotypic rescue, or clinical correlation rather than to the same underlying quantity. Evaluating all such outputs through one ranking metric obscures differences in evidence meaning, validation requirements, and translational consequence [2].

The central challenge is therefore not merely to combine more data. It is to reconcile claims that originate from different scientific contexts and may disagree for legitimate reasons. A transcriptomic reversal observed in one cell line may conflict with genetic evidence suggesting that the relevant target should be activated rather than inhibited. A network score may support a drug–disease relationship while a previous trial failed because the tested population, dose, disease stage, formulation, or endpoint did not match the proposed mechanism. Conversely, a prior negative outcome may reveal a genuine mechanistic contradiction that should strongly reduce enthusiasm for the candidate. Computational repurposing consequently requires explicit treatment of context, contradiction, and inferential distance rather than unqualified evidence accumulation [3].

This article proposes an original therapeutic evidence-synthesis model that treats drug repurposing as a structured reconciliation problem. The model does not replace experimental pharmacology, causal investigation, safety assessment, clinical development, or regulatory review. Instead, it organizes the evidentiary questions that should be answered before a repurposing hypothesis is interpreted as sufficiently coherent to advance. Its central contribution is a distinction among evidence-source identity, biological directionality, causal proximity, contextual compatibility, prior failure, and translational readiness. Candidate prioritization is thereby reframed from "Which drug receives the highest score?" to "Which hypothesis retains a scientifically defensible interpretation after supportive, contradictory, uncertain, and negative evidence have been reconciled?"

Why Repurposing Signals Frequently Conflict Across Evidence Sources

Repurposing signals frequently conflict because different analytical methods operationalize therapeutic relevance in different ways. Even within transcriptomic connectivity analysis, candidate rankings may vary according to the disease signature, perturbation reference set, scoring method, weighting scheme, and definition of reversal. A compound can therefore appear strongly connected under one formulation and weakly or oppositely connected under another without either analysis necessarily containing a computational error. Comparative work on connectivity scoring shows that multiple scores can capture partly different aspects of the relationship between disease-associated and drug-induced expression, making reconciliation necessary when rankings diverge [4]. The conflict is scientific as well as statistical: a score may reflect broad transcriptional opposition while failing to identify whether the affected genes belong to a causally relevant pathway, an adaptive response, or a downstream disease consequence.

Reproducibility further limits the interpretation of isolated repurposing signals. Perturbational signatures may vary with concentration, exposure duration, cell lineage, assay conditions, batch structure, and preprocessing decisions. Disease signatures can likewise change with tissue source, disease stage, patient composition, cell-type abundance, and analytical normalization. Evaluations of connectivity-map-based repositioning have shown that candidate recovery and ranking can be unstable across datasets and analytical choices [5]. This instability does not render signature methods useless; rather, it changes the evidentiary claim they can support. A reproducible, context-matched reversal may justify mechanistic investigation, whereas a single high score from a poorly matched model should remain a weak hypothesis-generating observation.

Knowledge graphs and network models introduce a different class of conflict. Their outputs depend on which entities and relationships are represented, how evidence is curated, whether contradictory relations are encoded, and how missing links are treated. A drug may rank highly because it is connected to many well-studied targets or diseases, while a less frequently studied compound may receive a lower score despite having a more specific biological rationale. Database overlap, literature bias, relation-type ambiguity, temporal inconsistency, and incomplete provenance can all influence apparent network support. Reviews of knowledge-graph repurposing methods show substantial variation in graph construction, included databases, prediction tasks, and evaluation strategies [6]. Consequently, graph support should be interpreted as structured relational evidence, not as direct proof of mechanism or clinical effect.

Genomic repurposing methods also differ in what they infer. Some associate disease-linked loci with candidate genes; others use expression or protein quantitative-trait loci, Mendelian randomization, colocalization, pathway enrichment, or genetic similarity to known drug mechanisms. These approaches vary in their assumptions about causal variants, gene assignment, horizontal pleiotropy, tissue relevance, and the equivalence between lifelong genetic perturbation and pharmacological intervention. Methodological surveys demonstrate that genomic repurposing pipelines can produce different candidate sets because they occupy different positions between disease association and therapeutic mechanism [7]. The proposed model therefore treats disagreement not as noise to be averaged away, but as information about which inferential step remains unresolved. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop why repurposing signals frequently conflict across evidence sources within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Why repurposing signals frequently conflict across evidence sources

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Disease-signature evidence	Does the drug perturbation oppose or normalize a disease-associated molecular state?	Comparison of disease-associated and compound-induced molecular profiles	Generates an initial directionally interpretable hypothesis	Treating transcriptional opposition as therapeutic efficacy	Signature reversal alone does not establish causal mechanism, target engagement, or patient benefit
Perturbation context	Was the drug measured in a biologically relevant cell type, concentration, duration, and state?	Context-dependent molecular and cellular response	Determines whether a signature is transferable to the proposed indication	Generalizing from unmatched models	Evidence is conditional on similarity between the experimental system and the intended disease context
Target-mechanism evidence	Does the compound engage a target or pathway plausibly involved in disease biology?	Pharmacology, target annotation, pathway biology, and mechanistic experiments	Connects a candidate to an interpretable intervention mechanism	Inferring mechanism from association or proximity alone	A plausible target relationship requires confirmation of engagement, direction, and disease relevance
Network or knowledge-graph evidence	Is the candidate relationally connected to disease-relevant genes, pathways, phenotypes, or therapies?	Typed biomedical relationships and network topology	Integrates heterogeneous evidence and exposes possible mechanistic paths	Equating graph proximity with causality or efficacy	Network position is supportive only when relation types, provenance, and inferential distance remain visible
Human genetic evidence	Does human variation support a causal role for the proposed target or pathway?	Genetic association, variant-to-gene mapping, colocalization, and causal inference	Increases or decreases causal plausibility and may inform directionality	Assuming genetic support guarantees pharmacological success	Genetic evidence strengthens a hypothesis but does not resolve modality, dose, timing, tissue, or safety
Phenotypic evidence	Does the compound produce a relevant functional response in a disease model?	Cellular, organoid, tissue, or organism-level functional assays	Tests whether molecular predictions translate into a measurable biological effect	Treating model rescue as clinical benefit	Phenotypic activity remains dependent on model validity, exposure, mechanism, and translatability
Prior clinical evidence	What has already been observed in treated patients or trials?	Trials, observational analyses, safety reports, and treatment histories	Constrains indication transfer and identifies supportive or negative human evidence	Ignoring failure or interpreting uncontrolled associations as efficacy	Clinical evidence must be interpreted in relation to population, dose, timing, comparator, endpoint, and study design
Contradiction structure	Are conflicting findings mechanistically incompatible or contextually explainable?	Cross-evidence comparison and causal interpretation	Prevents automatic averaging of incompatible claims	Downgrading all disagreement equally or selecting only supportive evidence	Contradictions must be classified before they are weighted
Evidence provenance	Where did each claim originate, and how directly does it support the hypothesis?	Source traceability, study design, curation, and versioning	Preserves accountability and enables later reassessment	Combining evidence without source identity	Integrated evidence must remain separable and auditable
Translational boundary	What decision can the available evidence justifiably support?	Alignment between evidence maturity and development action	Restricts conclusions to the appropriate stage	Converting computational prioritization into treatment recommendation	Candidate ranking can support investigation but cannot establish clinical use

Disease Signatures, Target Mechanisms, Phenotypes, and Prior Clinical Outcomes

Disease signatures provide a representation of the molecular state associated with a condition, but they should not be interpreted as the disease itself. Large perturbational resources enable systematic comparison between disease-associated expression patterns and compound-induced responses across many experimental settings [8]. Their value lies in identifying drugs that may oppose, reinforce, or reorganize selected components of a molecular state. Yet a disease signature may contain causal drivers, compensatory responses, treatment effects, changes in cell composition, and nonspecific consequences of tissue injury. The proposed evidence-synthesis model therefore divides signature evidence into at least three interpretive classes: potentially causal features, mechanistically informative correlates, and context-dependent consequences. A candidate should receive stronger support when its perturbation affects the first two classes in a direction consistent with therapeutic benefit, rather than merely producing a large global reversal score.

Cellular resolution is particularly important because bulk disease profiles can combine signals from multiple cell types with different or opposing roles. A compound predicted to reverse the average tissue signature may suppress a protective response in one population while correcting a pathogenic state in another. Single-cell-guided repurposing pipelines demonstrate how drug prioritization can change when disease signatures are decomposed by cell type and when the relative importance of affected populations is considered [9]. In the proposed model, cellular context is not a minor adjustment applied after candidate ranking. It is a primary compatibility condition that influences whether evidence can be transferred from the measured system

to the intended therapeutic setting. Context assessment should consider cell identity, disease state, tissue environment, exposure conditions, and whether the molecular response occurs in the compartment where the mechanism must operate.

Target mechanisms provide an explanatory bridge between molecular reversal and functional intervention. A drug may normalize a disease-associated profile through its intended target, through polypharmacology, through cellular stress, or through a secondary adaptive response. These possibilities have different implications for efficacy, selectivity, safety, and transfer to another indication. Phenotypic profiling of non-oncology compounds in cancer models illustrates that repurposing opportunities can emerge from systematic functional testing, while also showing that activity may be lineage-specific, concentration-dependent, and mechanistically heterogeneous [10]. Phenotypic evidence is therefore stronger than a purely relational prediction when it demonstrates a relevant functional effect, but it remains incomplete until the responsible mechanism, achievable exposure, selectivity, and disease-model validity are understood. The model consequently treats phenotype and mechanism as mutually constraining evidence layers rather than as substitutes.

Prior clinical outcomes occupy the highest human-proximity layer but still require contextual interpretation. Observational treatment histories may reveal associations between drug exposure and disease outcomes, yet they are vulnerable to confounding by indication, treatment selection, adherence, comorbidity, healthcare access, and measurement practices. Target-trial emulation offers a structured way to align observational analyses with a hypothetical randomized comparison by defining eligibility, treatment strategies, follow-up, outcomes, and analytic assumptions [11]. Within the proposed model, prior clinical evidence is examined through a context-transfer profile comprising indication, population, disease stage, dose, route, treatment timing, duration, comparator, endpoint, and co-therapy. A favorable association increases confidence only when these elements are sufficiently compatible with the proposed use and major biases have been addressed. A negative result is likewise not automatically universal: it may invalidate the mechanism, or it may delimit the conditions under which the candidate should not be expected to work. Evidence reconciliation begins by determining which of these interpretations is scientifically defensible.

Evidence-Reconciliation Logic for Candidate Prioritization

The proposed model begins by treating each repurposing hypothesis as a structured claim rather than as a drug name attached to a score. The claim specifies the candidate, proposed indication, intended mechanism, required direction of modulation, relevant biological compartment, anticipated phenotype, and clinical context. Evidence is then attached to this claim according to its source and inferential role. Heterogeneous biomedical networks demonstrate that drugs, diseases, genes, pathways, phenotypes, adverse effects, and established treatments can be integrated without representing every relationship as equivalent [12]. In the proposed model, this heterogeneity is preserved so that a drug–disease association, drug–target interaction, target–pathway relationship, and clinical observation remain distinguishable throughout prioritization.

The first reconciliation operation is evidence typing. Each item is classified as signature, mechanistic, genetic, phenotypic, network, observational, interventional, safety, pharmacological, or contextual evidence. It is also annotated by provenance, study design, model system, population, direction of effect, and distance from the proposed therapeutic outcome. Network-based repurposing linked with population-level analysis illustrates how predictions generated at one evidentiary level may be evaluated against another without implying that the two levels establish the same claim [13]. The model therefore prohibits simple vote counting. Several correlated databases or algorithms may represent one underlying evidence stream, whereas one well-designed human study may provide a qualitatively different constraint.

The second operation is compatibility assessment. Evidence is judged against the proposed target–mechanism–context statement rather than against the candidate in the abstract. Compatibility asks whether the disease process, drug action, tissue, direction, dose, timing, and observed phenotype can coexist within one biologically plausible account. Multilayer knowledge-graph learning shows that drug-repurposing inference can draw on relationships across genes, drugs, and diseases while preserving semantic structure [14]. The proposed model extends this principle conceptually by distinguishing supportive compatibility, unresolved compatibility, contextual incompatibility, and mechanistic incompatibility. A source may therefore support one component of a hypothesis while contradicting another.

The third operation is bounded prioritization. Candidates are assigned an evidence profile rather than a universal scalar value. The profile records evidentiary convergence, causal proximity, directionality agreement, context match, negative evidence, unresolved contradictions, and validation requirements. Network-medicine platforms illustrate the value of integrating disease-module identification, drug prioritization, statistical assessment, and expert input within one analytical environment [15]. The proposed contribution is not the platform itself but the reconciliation logic applied before advancement: candidates move forward because their evidence profile is coherent enough for a specified next decision, not because heterogeneous signals have been collapsed into a deceptively precise aggregate score.

Weighting Directionality, Context, Causal Proximity, and Negative Evidence

Directionality concerns whether the pharmacological action of a candidate moves a disease-relevant process in the therapeutically required direction. A target may be associated with disease without indicating whether activation, inhibition, stabilization, degradation, or partial modulation is appropriate. The druggable-genome literature demonstrates how human genetics can support target identification while leaving separate questions concerning tractability, pharmacological modality, and target validation [16]. Within the proposed model, directional evidence receives weight only when the direction of the

exposure or molecular perturbation can be related plausibly to the intended drug action. A target–disease link without directional information increases relevance but does not establish a therapeutic intervention strategy.

Causal proximity describes the inferential distance between an observed signal and the proposed therapeutic effect. Randomized clinical evidence generally lies closer to patient benefit than database co-occurrence, but proximity is not synonymous with universal superiority because the clinical context may not match the new hypothesis. Human genetic support can strengthen confidence in a drug mechanism, yet analyses of development outcomes show that its association with success depends on how targets, indications, and genetic relationships are defined [17]. The model therefore uses causal proximity as one dimension rather than as an automatic hierarchy. Genetic, mechanistic, phenotypic, observational, and clinical evidence must still be evaluated for relevance to the exact target–indication pair.

Context weighting determines whether evidence generated in one setting can inform another. Tissue, cell type, ancestry, age, disease stage, exposure duration, comorbidity, treatment history, and endpoint can alter the meaning of an otherwise credible signal. Phenome-wide Mendelian-randomization analyses combined with colocalization illustrate both the value of genetically informed target assessment and the possibility that apparently causal associations may be weakened when linked signals are examined more carefully [18]. Context weighting should therefore decrease as the biological or clinical transfer distance increases. It should not be implemented as an arbitrary numerical penalty; it should identify which contextual assumptions must be tested before the evidence is transferable.

Negative evidence includes findings that oppose efficacy, indicate an incompatible mechanism, identify toxicity, expose non-reproducibility, or show failure under specified conditions. Its weight depends on study validity and contextual relevance. Human genetic data can contribute to safety assessment by revealing phenotypes associated with long-term modulation of a therapeutic target, although genetic perturbation is not fully equivalent to drug exposure [19]. The proposed model distinguishes absence of supporting evidence from evidence of absence, and it distinguishes contextual failure from mechanism-level failure. Strong negative evidence cannot be neutralized merely by accumulating additional low-proximity positive predictions. **Figure 1** presents the drug-repurposing evidence reconciliation loom, showing how the article's key components and boundaries are connected within weighting directionality, context, causal proximity, and negative evidence.

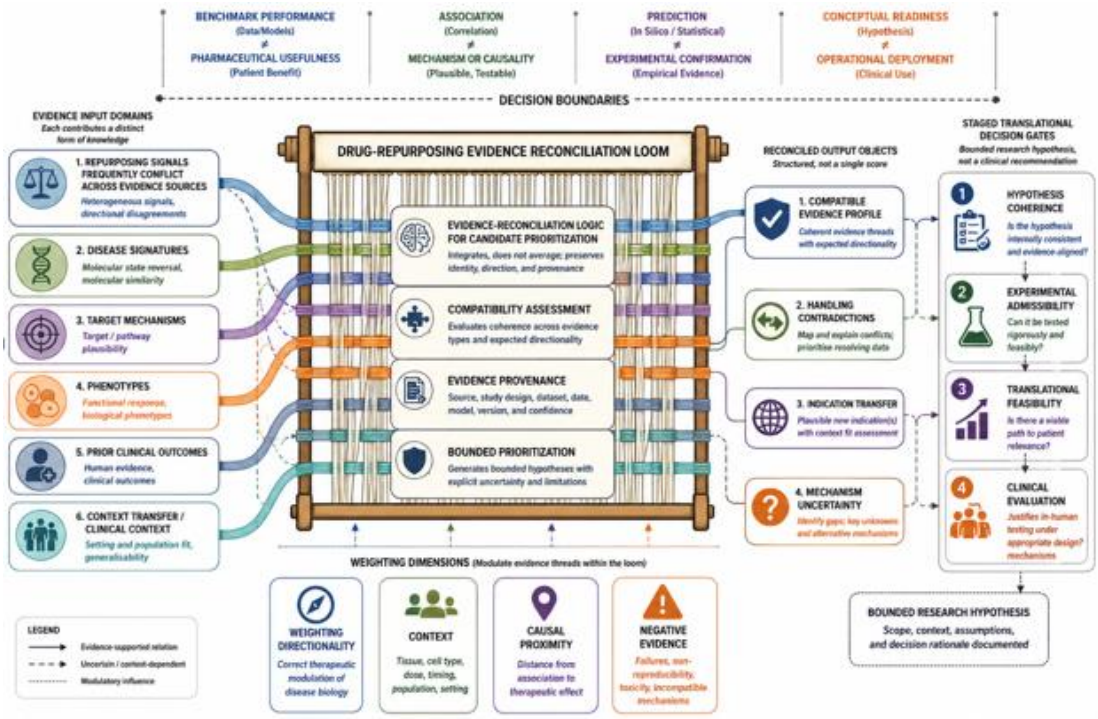

Figure 1. Drug-Repurposing Evidence Reconciliation Loom.

The figure is an original conceptual synthesis that organizes the article's central contribution across repurposing signals frequently conflict across evidence sources, disease signatures, target mechanisms, phenotypes, and prior clinical outcomes, disease signatures, target mechanisms, prior clinical outcomes, evidence-reconciliation logic for candidate prioritization. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Handling Contradictions, Indication Transfer, and Mechanism Uncertainty

Contradictions should first be classified according to the scientific claim they affect. Direct contradictions occur when credible evidence supports opposing directions for the same mechanism in a sufficiently similar context. Apparent contradictions arise when studies measure different outcomes, populations, doses, disease stages, or biological compartments. Representation

contradictions reflect incompatible database labels or relation types, while temporal contradictions arise when evidence predates changes in disease definition, treatment practice, or mechanistic knowledge. Clinician-centered knowledge-graph models illustrate the potential value of exposing multi-hop rationales and predicted contraindications rather than returning drug rankings alone [20]. The proposed model similarly requires each contradiction to remain traceable to the evidence path that produced it.

Indication transfer is evaluated by comparing the original and proposed uses across mechanism, exposure, tissue distribution, patient population, disease biology, endpoint, and treatment setting. Prior approval may reduce uncertainty about selected characteristics of a compound, but it does not make the original dose, formulation, or safety profile automatically suitable for a new disease. Integrated analyses combining gene-expression information with clinical data demonstrate how molecular predictions may be tested against patient-level observations [21]. Such convergence strengthens a hypothesis when the molecular direction and clinical phenotype are aligned, but observational replication does not remove confounding or substitute for controlled prospective evaluation.

Prior clinical failure must be interpreted as structured evidence rather than as a binary property of the molecule. Large randomized evaluations of repurposed antiviral regimens showed that compelling biological rationales and widespread clinical interest did not ensure benefit in hospitalized patients [22]. In the proposed model, a failed study is decomposed into mechanism tested, population, disease stage, exposure, treatment timing, comparator, endpoint, adherence, and study quality. A well-conducted failure under conditions that adequately test the proposed mechanism should strongly reduce priority. A failure under materially different conditions may instead define a boundary beyond which the evidence should not be generalized.

Mechanism uncertainty persists when the observed effect cannot be attributed confidently to the proposed target, when a drug has multiple relevant targets, or when activity appears only at exposures unlikely to be clinically attainable. Randomized evidence on hydroxychloroquine in hospitalized patients illustrates how in vitro activity and early uncontrolled findings can be contradicted by appropriately designed clinical evaluation [23]. The lesson is not that every negative trial invalidates every conceivable use of a compound. It is that mechanistic reinterpretation after failure must be prospective and testable rather than an ad hoc attempt to preserve the hypothesis. The model therefore requires explicit documentation of which mechanistic proposition survived, which was contradicted, and what new evidence would be necessary to justify reconsideration.

Translational Decision Gates for Repurposing Hypotheses

The proposed translational sequence contains gates rather than a single advancement threshold. At each gate, the relevant question is whether the evidence is sufficient for the next form of inquiry, not whether the candidate is generally "promising." Real-world-data studies can generate, prioritize, or evaluate repurposing hypotheses, but their interpretation is sensitive to clinical granularity, data quality, confounding, treatment selection, and outcome definition [24]. Accordingly, observational support may justify deeper causal analysis or prospective study design, but it should not be converted directly into a therapeutic recommendation.

The hypothesis-coherence gate asks whether the candidate, indication, mechanism, direction, and context form a scientifically consistent claim. The experimental-admissibility gate asks whether the mechanism and phenotype can be tested in an appropriate model at relevant exposure. The translational-feasibility gate addresses formulation, pharmacokinetics, tissue access, dose rationale, safety, intellectual-property constraints, manufacturing, and the feasibility of generating confirmatory evidence. Contemporary analyses of therapeutic innovation in repurposing emphasize that an integrated evidence base, credible dose rationale, and clear development plan are needed early rather than after a candidate has attracted extensive attention [25].

The mechanistic-phenotypic gate requires evidence that the candidate produces a relevant functional effect and that the effect can be connected sufficiently to the proposed therapeutic mechanism. Phenotypic discovery literature shows that functionally useful compounds may emerge before their precise targets are fully understood, but it also emphasizes target deconvolution, model relevance, and translational interpretation [26]. The proposed model therefore does not require complete mechanistic certainty before every experiment. It requires the degree of uncertainty to be visible and proportional to the proposed decision. A candidate with strong phenotype but uncertain mechanism may advance to deconvolution studies, not directly to clinical efficacy claims.

The clinical-evaluation gate asks whether residual uncertainty can be addressed through an ethically and scientifically justified human study. This includes a defensible patient population, treatment window, comparator, dose, endpoint, monitoring plan, and evidence that the mechanism can operate under the proposed conditions. Multidimensional pharmaceutical-development frameworks demonstrate the importance of aligning target biology, tissue exposure, safety, patient selection, and therapeutic differentiation [27]. The proposed repurposing model adapts this logic by requiring prior failures and indication-transfer assumptions to be reviewed before clinical advancement. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop translational decision gates for repurposing hypotheses within the article's central argument.

Table 2. Evaluation, implementation, and research priorities arising from Drug Repurposing as Evidence Reconciliation across Disease Signatures, Molecular Mechanisms, Prior Failures

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement

Hypothesis specification	Candidate, indication, proposed target, mechanism, direction, tissue, population, and endpoint	Converts an unbounded candidate into a testable therapeutic claim	Versioned hypothesis statement	Initial mechanism or context may be incompletely known	Independent expert review and consistency assessment
Evidence typing	Signature, network, genetic, phenotypic, pharmacological, observational, trial, and safety evidence	Preserves the scientific meaning and provenance of each source	Typed evidence inventory	Sources may be correlated or incompletely reported	Source-level audit and reproducible extraction
Directionality assessment	Disease-associated direction, drug action, genetic effects, pathway response, and phenotype	Tests whether modulation is aligned with intended benefit	Directionality-compatible or unresolved classification	Genetic and pharmacological perturbations may differ	Target-engagement and perturbation studies
Context compatibility	Cell type, tissue, disease stage, ancestry, dose, timing, route, population, and endpoint	Determines whether evidence can transfer to the proposed use	Context-transfer profile	Relevant contextual variables may be unavailable	Replication in matched biological and clinical settings
Causal-proximity assessment	Association, network path, genetic inference, mechanism, phenotype, observational outcome, and trial result	Locates evidence along the path from signal to therapeutic effect	Causal-proximity profile	High proximity does not ensure contextual fit	Triangulation using independent designs
Contradiction classification	Opposing directions, discordant phenotypes, database conflicts, and prior failures	Distinguishes mechanistic incompatibility from contextual disagreement	Resolved, bounded, unresolved, or disqualifying contradiction	Selective reporting may conceal negative findings	Prespecified contradiction adjudication
Negative-evidence layer	Failed replication, toxicity, contraindication, null phenotype, observational harm, and trial failure	Prevents supportive evidence from overwhelming credible adverse findings	Negative-evidence register	Absence of evidence may be confused with evidence of absence	Validity and relevance assessment for each negative result
Experimental-admissibility gate	Coherent mechanism, suitable model, feasible assay, achievable exposure, and measurable phenotype	Determines whether laboratory evaluation is justified	Experimental study recommendation or de-prioritization	Model systems may not reproduce human disease	Reproducibility across complementary models
Translational-feasibility gate	Pharmacokinetics, formulation, tissue exposure, safety, manufacturing, and development constraints	Tests whether biological plausibility can become a viable development program	Translational feasibility statement	Existing data may derive from another dose or indication	Fit-for-purpose pharmacology and safety studies
Clinical-evaluation gate	Human rationale, dose, population, timing, comparator, endpoint, monitoring, and equipoise	Determines whether prospective clinical evaluation is scientifically defensible	Trial-ready research hypothesis, not a clinical recommendation	Residual confounding and mechanism uncertainty may remain	Prospectively registered and appropriately controlled evaluation
Post-evidence updating	New experimental, observational, safety, or clinical findings	Revises the hypothesis without erasing prior evidence	Versioned evidence profile and decision history	Updating may be biased toward preserving favored candidates	Independent reassessment using prespecified rules

Limitations and Boundary Conditions

The proposed model is a conceptual synthesis and has not been empirically validated as a prioritization system. It does not establish an optimal mathematical weighting function, universal evidence hierarchy, or numerical threshold for candidate advancement. The relevance of each evidentiary layer varies across disease areas, modalities, mechanisms, and development stages. A genetic signal may be highly informative for one target yet limited for another, while a phenotypic assay may closely reproduce one disease process but poorly model another. The framework can organize such differences, but it cannot determine their importance without domain-specific scientific judgment.

The model is also limited by the evidence available to populate it. Repurposing databases may contain incomplete target annotations, publication bias, duplicated evidence, inconsistent disease terminology, and outdated clinical information. Negative experiments and abandoned programs may be less visible than positive findings. Transcriptomic, network, genetic, phenotypic, and observational sources can share hidden dependencies, creating an appearance of independent convergence. Reconciliation cannot correct these problems automatically; it can only make provenance, dependence, uncertainty, and missingness explicit. Poor-quality evidence does not become reliable merely because it has been carefully categorized.

The framework cannot establish efficacy, safety, benefit–risk acceptability, regulatory suitability, or clinical readiness. It is not a substitute for experimental confirmation, pharmacological characterization, toxicology, clinical-trial design, or patient-centered decision-making. Its use as an automated prescribing or operational development tool would exceed its intended scope. The appropriate output is a bounded statement about what the current evidence supports, what it contradicts, and which study should occur next. Human scientific oversight remains necessary at every gate, particularly when evidence is sparse, contradictory, or consequential.

Research Agenda and Implementation Priorities

The first priority is to evaluate whether explicit reconciliation improves decision quality compared with conventional ranking. Retrospective studies should reconstruct completed repurposing programs without using information that became available after the original decision. Candidate rankings based on score aggregation could then be compared with profiles that preserve

directionality, context, causal proximity, contradiction, and negative evidence. Evaluation should examine whether the proposed model identifies known failure mechanisms, avoids unsupported escalation, and produces stable explanations under reasonable changes in data source or analytical method. Such work should assess decision calibration rather than only candidate-retrieval accuracy.

A second priority is the development of interoperable evidence representations. Repurposing resources require controlled descriptions of disease state, target action, molecular direction, exposure, tissue, population, endpoint, provenance, and evidentiary status. Contradictory claims should be stored without forcing premature resolution, and historical versions should remain accessible when knowledge changes. Reporting standards should require authors to distinguish database prediction, mechanistic inference, experimental confirmation, observational association, and clinical evaluation. Candidate lists should be accompanied by explicit applicability boundaries and major negative findings rather than by performance metrics alone.

The third priority is staged prospective validation. Early studies should test individual components, including contradiction classification, context-transfer assessment, and negative-evidence weighting. Later studies could embed the complete model within multidisciplinary repurposing programs and record how it changes experiments, de-prioritization, and escalation decisions. Implementation should begin as a transparent deliberative aid rather than an automated decision engine. Advancement toward operational use would require reproducibility across therapeutic areas, evidence that the model improves decisions without concealing uncertainty, and governance mechanisms for expert challenge, documentation, updating, and accountability.

Conclusion

Drug repurposing is not adequately represented as the search for a single persuasive signal or the optimization of one predictive score. It is an evidence-reconciliation problem in which disease signatures, molecular mechanisms, phenotypes, human genetics, clinical histories, negative findings, and contextual constraints make different claims about the same therapeutic hypothesis. The proposed Drug-Repurposing Evidence Reconciliation Loom preserves these differences while organizing directionality, causal proximity, context compatibility, contradiction, and prior failure into bounded candidate profiles and staged translational gates. Its contribution is conceptual rather than validated: it can help clarify why a candidate should be investigated, de-prioritized, or reconsidered, but it cannot establish efficacy, safety, clinical utility, regulatory acceptability, or deployment readiness. The central requirement is not universal evidentiary agreement, but a transparent account of which findings converge, which remain uncertain, which conflict, and what evidence is required before the hypothesis can responsibly advance.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Pushpakom S, Iorio F, Eyers PA, Escott KJ, Hopper S, Wells A, et al. Drug repurposing: Progress, challenges and recommendations. Nat Rev Drug Discov. 2019;18(1):41-58. doi:10.1038/nrd.2018.168
2. Tanoli Z, Fernández-Torras A, Özcan UO, Kushnir A, Nader KM, Gadiya Y, et al. Computational drug repurposing: Approaches, evaluation of in silico resources and case studies. Nat Rev Drug Discov. 2025;24(7):521-42. doi:10.1038/s41573-025-01164-x
3. Sosa DN, Altman RB. Contexts and contradictions: A roadmap for computational drug repurposing with knowledge inference. Brief Bioinform. 2022;23(4). doi:10.1093/bib/bbac268
4. Samart K, Tuyishime P, Krishnan A, Ravi J. Reconciling multiple connectivity scores for drug repurposing. Brief Bioinform. 2021;22(6). doi:10.1093/bib/bbab161
5. Lim N, Pavlidis P. Evaluation of Connectivity Map shows limited reproducibility in drug repositioning. Sci Rep. 2021;11(1):17624. doi:10.1038/s41598-021-97005-z
6. Perdomo-Quinteiro P, Belmonte-Hernández A. Knowledge graphs for drug repurposing: A review of databases and methods. Brief Bioinform. 2024;25(6). doi:10.1093/bib/bbae461
7. Wang L, Lu Y, Li D, Zhou Y, Yu L, Mesa Eguiagaray I, et al. The landscape of the methodology in drug repurposing using human genomic data: A systematic review. Brief Bioinform. 2024;25(2). doi:10.1093/bib/bbad527
8. Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu X, et al. A next generation Connectivity Map: L1000 platform and the first 1,000,000 profiles. Cell. 2017;171(6):1437-52. doi:10.1016/j.cell.2017.10.049
9. He B, Xiao Y, Liang H, Huang Q, Du Y, Li Y, et al. ASGARD is a single-cell guided pipeline to aid repurposing of drugs. Nat Commun. 2023;14(1):993. doi:10.1038/s41467-023-36637-3

10. Corsello SM, Nagari RT, Spangler RD, Rossen J, Kocak M, Bryan JG, et al. Discovering the anticancer potential of non-oncology drugs by systematic viability profiling. Nat Cancer. 2020;1(2):235-48. doi:10.1038/s43018-019-0018-6
11. Zang C, Zhang H, Xu J, Zhang H, Fouladvand S, Havaldar S, et al. High-throughput target trial emulation for Alzheimer's disease drug repurposing with real-world data. Nat Commun. 2023;14(1):8180. doi:10.1038/s41467-023-43929-1
12. Himmelstein DS, Lizee A, Hessler C, Brueggeman L, Chen SL, Hadley D, et al. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. Elife. 2017;6. doi:10.7554/eLife.26726
13. Cheng F, Desai RJ, Handy DE, Wang R, Schneeweiss S, Barabási AL, et al. Network-based approach to prediction and population-based validation of in silico drug repurposing. Nat Commun. 2018;9(1):2691. doi:10.1038/s41467-018-05116-5
14. Bang D, Lim S, Lee S, Kim S. Biomedical knowledge graph learning for drug repurposing by extending guilt-by-association to multiple layers. Nat Commun. 2023;14(1):3570. doi:10.1038/s41467-023-39301-y
15. Sadegh S, Skelton J, Anastasi E, Bernett J, Blumenthal DB, Galindez G, et al. Network medicine for disease module identification and drug repurposing with the NeDRex platform. Nat Commun. 2021;12(1):6848. doi:10.1038/s41467-021-27138-2
16. Finan C, Gaulton A, Kruger FA, Lumbers RT, Shah T, Engmann J, et al. The druggable genome and support for target identification and validation in drug development. Sci Transl Med. 2017;9(383). doi:10.1126/scitranslmed.aag1166
17. King EA, Davis JW, Degner JF. Are drug targets with genetic support twice as likely to be approved? Revised estimates of the impact of genetic support for drug mechanisms on the probability of drug approval. PLoS Genet. 2019;15(12). doi:10.1371/journal.pgen.1008489
18. Zheng J, Haberland V, Baird D, Walker V, Haycock PC, Hurle MR, et al. Phenome-wide Mendelian randomization mapping the influence of the plasma proteome on complex diseases. Nat Genet. 2020;52(10):1122-31. doi:10.1038/s41588-020-0682-6
19. Carss KJ, Deaton AM, Del Rio-Espinola A, Diogo D, Fielden M, Kulkarni DA, et al. Using human genetics to improve safety assessment of therapeutics. Nat Rev Drug Discov. 2023;22(2):145-62. doi:10.1038/s41573-022-00561-w
20. Huang K, Chandak P, Wang Q, Havaldar S, Vaid A, Leskovec J, et al. A foundation model for clinician-centered drug repurposing. Nat Med. 2024;30(12):3601-13. doi:10.1038/s41591-024-03233-x
21. Wu P, Feng Q, Kerchberger VE, Nelson SD, Chen Q, Li B, et al. Integrating gene expression and clinical data to identify drug repurposing candidates for hyperlipidemia and hypertension. Nat Commun. 2022;13(1):46. doi:10.1038/s41467-021-27751-1
22. WHO Solidarity Trial Consortium. Repurposed antiviral drugs for Covid-19—interim WHO Solidarity Trial results. N Engl J Med. 2021;384(6):497-511. doi:10.1056/NEJMoa2023184
23. RECOVERY Collaborative Group. Effect of hydroxychloroquine in hospitalized patients with Covid-19. N Engl J Med. 2020;383(21):2030-40. doi:10.1056/NEJMoa2022926
24. Tan GSQ, Sloan EK, Lambert P, Kirkpatrick CMJ, Ilomäki J. Drug repurposing using real-world data. Drug Discov Today. 2023;28(1):103422. doi:10.1016/j.drudis.2022.103422
25. Recino A, Rayner MLD, Rohn JL, Della Pasqua O; UCL Repurposing TIN Committee. Therapeutic innovation in drug repurposing: Challenges and opportunities. Drug Discov Today. 2025;30(7):104390. doi:10.1016/j.drudis.2025.104390
26. Vincent F, Nueda A, Lee J, Schenone M, Prunotto M, Mercola M. Phenotypic drug discovery: Recent successes, lessons learned and new directions. Nat Rev Drug Discov. 2022;21(12):899-914. doi:10.1038/s41573-022-00472-w
27. Morgan P, Brown DG, Lennard S, Anderton MJ, Barrett JC, Eriksson U, et al. Impact of a five-dimensional framework on R&D productivity at AstraZeneca. Nat Rev Drug Discov. 2018;17(3):167-81. doi:10.1038/nrd.2017.244