



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Closing the Missing Middle between Pharmacogenomic Association, Functional Evidence, Prescribing Recommendations, and Patient-Level Therapeutic Decisions in Routine Care

James O'Leary^{1*}, Aisling Dunne¹, Sean O'Brien²

¹ Department of Pharmacogenomic Translation and Prescribing, School of Pharmacy, Trinity College Dublin, Dublin, Ireland.

² Department of Patient-Level Therapeutic Decisions, School of Pharmacy, University College Cork, Cork, Ireland.

*Email: james.oleary@tcd.ie

ABSTRACT

Pharmacogenomics has generated clinically relevant gene–drug associations, increasingly standardized variant interpretations, prescribing recommendations, and implementation infrastructures. Nevertheless, a persistent translational discontinuity remains between identifying an association and making a defensible therapeutic decision for an individual patient. This discontinuity arises because association credibility, variant and haplotype resolution, functional interpretation, genotype-to-phenotype translation, recommendation formation, patient-context assessment, and decision delivery answer different scientific and clinical questions. Treating any one of these components as sufficient can obscure uncertainty and overstate decision readiness. This Original Evidence-Translation Framework Article proposes the Pharmacogenomic Evidence Bridge as a conceptual architecture for closing this missing middle. The framework represents translation as a sequence of non-substitutable evidence states connected by explicit transition gates. It distinguishes established pharmacogenomic concepts from proposed relationships governing functional resolvability, recommendation translatability, patient-context compatibility, and decision-use readiness. Cross-cutting mechanisms address evidence provenance, version control, uncertainty documentation, ancestry and sex transferability, evidence updating, and conflict resolution. Clinical decision support is positioned as an evidence-delivery and contextualization layer rather than an autonomous therapeutic authority. The framework generates testable propositions concerning evidentiary concordance, phenoconversion, transferability, workflow conditionality, and the validation requirements of patient-specific computational representations. Its evaluation would require analytical, functional, clinical, informatics, implementation, and equity-focused studies conducted in intended populations and therapeutic settings. The proposed architecture does not establish clinical utility, regulatory acceptance, universal applicability, or deployment readiness. Its original contribution is to make the intermediate evidentiary transitions between pharmacogenomic discovery and patient-level therapeutic use explicit, examinable, and governable.

Key words: Pharmacogenomics, Precision pharmacotherapy, Genotype, Phenoconversion, Clinical implementation, Dose individualization

Received: 19 March 2025; **Revised:** 07 July 2025; **Accepted:** 08 July 2025

INTRODUCTION

Pharmacogenomics seeks to explain and use inherited genomic variation that contributes to variability in drug exposure, efficacy, and toxicity. Its scientific development has produced replicated gene–drug associations, functional annotations, genotype-to-phenotype translation systems, prescribing guidelines, implementation

programs, and increasingly computable clinical resources. Yet the evidence linking these achievements to patient-level therapeutic use remains incompletely integrated. Association discovery, clinical implementation infrastructure, and routine decision use are often discussed as parts of one translational continuum even though they involve different evidentiary questions, uncertainties, and responsibilities [1–3]. The central problem is therefore not simply insufficient pharmacogenomic knowledge. It is the absence of a sufficiently explicit account of how distinct forms of knowledge should be connected, challenged, updated, and bounded before they influence an individual therapeutic decision.

Pharmacogenomics, precision pharmacotherapy, and genotype-guided therapy overlap but are not interchangeable. Pharmacogenomics concerns relationships between genomic variation and drug-related phenotypes. Precision pharmacotherapy is broader, integrating genomic evidence with clinical characteristics, environmental exposures, treatment history, therapeutic alternatives, and patient priorities. Genotype-guided therapy describes the use of a genomic result in drug selection or dosing, but it does not establish that genotype is the only relevant determinant or that following a genotype-derived rule necessarily improves outcomes. These distinctions matter because the movement from a population-level association to an individual decision crosses multiple inferential boundaries that are easily compressed into the language of “actionability” [1]. A result may be analytically valid and biologically plausible without being clinically useful, while a guideline recommendation may be scientifically defensible yet inapplicable to a patient whose current phenotype differs from the genetically predicted state.

The phrase “missing middle” is used here to describe the insufficiently specified evidentiary space between association and decision. At one boundary, a reported gene–drug association must be translated into a correctly identified variant, allele, haplotype, or diplotype. At another, that genomic state must be connected to a biological function and then to a clinically interpretable phenotype.

This article develops an original, non-empirical Evidence-Translation Framework for organizing these transitions. It proposes the Pharmacogenomic Evidence Bridge, or PGx-EB, as a sequence of non-substitutable evidence states linked by explicit transition gates and cross-cutting governance requirements. The article does not claim exhaustive coverage of pharmacogenomic scholarship, derive a clinical guideline, evaluate a software system, or report empirical validation. Its objective is to distinguish the types of evidence required at successive stages, clarify which relationships are established and which are newly proposed, identify conditions under which evidence should be returned for reassessment, and define propositions and validation criteria for future investigation. The framework is intended to support more disciplined reasoning about readiness without converting heterogeneous evidence into an unsupported universal score.

The translational discontinuity in pharmacogenomics

The translational discontinuity begins with the difference between demonstrating that a gene–drug relationship exists and constructing the infrastructure through which that relationship can affect care. Clinical implementation requires more than laboratory testing. It depends on test selection, allele coverage, result representation, professional education, guideline interpretation, electronic integration, workflow design, patient communication, and mechanisms for updating stored knowledge. The Ubiquitous Pharmacogenomics initiative illustrates this infrastructural character by linking pre-emptive testing with multicountry implementation strategies, clinical decision support, professional preparation, and harmonized therapeutic guidance [4]. Such initiatives establish that pharmacogenomic translation is an organized sociotechnical process. They do not, however, demonstrate that one implementation configuration will be equally effective across populations, health systems, therapeutic areas, or testing platforms.

The distinction between implementation capability and clinical utility is equally important. A health system may successfully return results, store them in the electronic record, and display recommendations without establishing that the intervention improves patient-relevant outcomes. Conversely, absence of a detectable benefit in one implementation design does not invalidate the underlying biological relationship or every possible use of the information. The PREPARE study evaluated a pre-emptive pharmacogenetic panel as a coordinated clinical intervention rather than assuming that the existence of actionable gene–drug pairs was sufficient evidence of benefit [5]. Its design exemplifies the need to evaluate the complete pathway through which testing, interpretation, prescribing, and clinical follow-up interact. Nevertheless, panel-level evidence remains conditional on the included genes and drugs, participating settings, implementation procedures, comparator practices, and outcome definitions.

Randomized evidence further demonstrates why pharmacogenomic translation cannot be reduced to a binary division between “actionable” and “non-actionable.” In TAILOR-PCI, genotype-guided selection of antiplatelet therapy was evaluated within a specified cardiovascular population, treatment comparison, follow-up period, and statistical design [6]. Interpretation of such evidence requires attention to the primary endpoint, treatment crossover, background clinical practice, event frequency, timing of benefit, and the distinction between the direction of an estimate and the certainty attached to it. A context-dependent or statistically nuanced trial result does not erase prior mechanistic and pharmacological evidence, but neither can mechanistic plausibility substitute for direct clinical evaluation. The translational discontinuity therefore includes a design problem: evidence generated at one level may be relevant to another level without being sufficient for it.

The article proposes that this discontinuity be represented as a set of interfaces rather than as a single research-to-practice gap. The association-to-function interface asks whether the implicated genomic state has been identified correctly and whether its biological consequence is sufficiently resolved. The function-to-recommendation interface asks whether mechanistic and clinical evidence justify a treatment rule for a defined context. The recommendation-to-patient interface asks whether that rule remains applicable after considering current physiology, concomitant treatment, ancestry, competing risks, and therapeutic alternatives. The patient-to-workflow interface asks whether the information can be delivered and acted upon without distorting uncertainty or displacing clinician responsibility. This interpretation is an original synthesis. It is consistent with evidence showing that implementation architecture, prospective evaluation, and clinical-context specificity are separate requirements, but it has not been validated as a universal translational model.

Levels of evidence from association to clinical utility

A framework for the missing middle first requires terminological discipline. A genotype describes detected alleles at one or more loci; a diplotype represents the pair of inherited haplotypes assigned for a pharmacogene; and a predicted phenotype translates that genomic configuration into an anticipated functional category. A prescribing recommendation connects the interpreted phenotype to a proposed change in drug selection, dose, or monitoring, whereas clinical actionability concerns whether available evidence and alternatives justify considering such a change. Clinical utility addresses whether using the information produces a meaningful balance of benefits, harms, and resource consequences in a defined setting. Consensus terminology developed for clinical pharmacogenetic test reporting helps separate these concepts and reduces ambiguity in communication among laboratories, guideline groups, informatics systems, and clinicians [7]. Standardized language is therefore a prerequisite for translation, although terminology agreement does not itself establish biological correctness or clinical benefit.

Evidence classification must likewise distinguish the question being answered. Association evidence concerns whether genomic variation is statistically related to a drug phenotype under specified conditions. Functional evidence concerns whether the relevant variant or haplotype alters expression, activity, regulation, transport, binding, metabolism, or another mechanism capable of contributing to the observed phenotype. Clinical-validity evidence concerns whether the genomic state predicts a clinically relevant drug response, while recommendation evidence concerns whether a professional group can translate the accumulated knowledge into an actionable rule. Implementation and utility evidence address whether the rule can be integrated into care and whether its use changes clinically meaningful outcomes. The PharmGKB evidence framework demonstrates how pharmacogenomic knowledge can be organized through structured clinical annotations, variant annotations, evidence levels, and guideline relationships rather than treated as an undifferentiated body of support [8]. Such curation increases transparency but cannot eliminate judgement regarding study quality, relevance, or transferability.

The existence of multiple credible guideline organizations shows that evidence does not move mechanically from publication to recommendation. CPIC and the Dutch Pharmacogenetics Working Group may evaluate overlapping evidence while differing in gene–drug scope, phenotype definitions, therapeutic alternatives, recommendation categories, and the intended point of use. Their comparison demonstrates that guideline discordance can arise from differences in evidence selection, interpretive rules, clinical objectives, or health-system context rather than simple factual disagreement [9]. Consequently, a recommendation should not be treated as a direct restatement of an association. It is a deliberative translation product that embeds assumptions about evidence sufficiency, clinical consequences, available alternatives, and acceptable uncertainty.

Functional validation of variants and haplotypes

Functional validation occupies a pivotal position because it connects detected genomic variation to an interpretable biological consequence. That connection depends on more than demonstrating an effect in a single experimental system. Clinically interpretable function depends jointly on controlled allele and haplotype nomenclature, adequate detection of structurally complex variants, and standardized genotype-to-phenotype translation [10–12]. PharmVar provides a governed vocabulary for pharmacogene variation, but a stable allele name does not prove that a testing platform can identify every relevant configuration. Similarly, a consensus phenotype term can improve consistency without resolving uncertainty in the underlying genotype or function. Functional validation should therefore begin with specification of the genomic object being evaluated: the individual variant, cis-linked haplotype, complete star allele, diplotype, copy-number configuration, or another structurally defined state.

Analytical incompleteness can appear clinically as biological inconsistency. CYP2D6 illustrates this problem because deletions, duplications, hybrid alleles, tandem arrangements, and other structural configurations can be missed or misassigned by assays optimized for a limited set of variants. More complete interrogation of CYP2D6 structural variation has been shown to improve correspondence between genotype and measured metabolic activity [11]. This evidence supports a general methodological warning rather than a universal numerical rule: apparent discordance between genotype and phenotype may reflect incomplete genomic resolution before it reflects failure of the biological model. The relevance of this warning varies by pharmacogene, platform, ancestry, and allele architecture, and the CYP2D6 example should not be generalized mechanically to genes with different genomic structures.

After the genomic state has been identified, its function must be translated in a reproducible manner. Functional evidence may derive from biochemical assays, expression systems, cellular models, tissue observations, endogenous biomarkers, probe-drug studies, pharmacokinetic measurements, or clinical response patterns. These sources differ in biological proximity, experimental control, susceptibility to confounding, and relevance to the intended therapeutic context. Standardized CYP2D6 genotype-to-phenotype recommendations demonstrate the importance of agreed translation rules, particularly when different combinations of alleles and activity values can otherwise generate inconsistent phenotype assignments [12]. Nevertheless, a standardized rule remains conditional on the accuracy of allele calling, the validity of assigned function, and the absence of major contextual modifiers. It predicts a functional tendency rather than directly measuring the patient's current metabolic state.

Translation into prescribing recommendations

Prescribing recommendations occupy an intermediate position between evidence interpretation and therapeutic decision-making. They translate a defined genomic or predicted phenotypic state into an advised change in drug selection, dose, treatment intensity, or monitoring. The updated CPIC guideline for CYP2C19 and clopidogrel illustrates this translation by linking metabolizer phenotypes with antiplatelet-therapy recommendations in specified cardiovascular settings [13]. The recommendation is not equivalent to the underlying association, because it also incorporates the clinical consequences of reduced active-metabolite formation, the availability of alternative agents, and the therapeutic setting in which the result is intended to be used. Nor is it equivalent to a patient-specific prescription, because contraindications, bleeding risk, indication, co-medications, access, and clinician judgement remain external to the genotype-derived recommendation.

Translation can also involve quantitative dose individualization rather than substitution of one drug for another. DPYD-guided fluoropyrimidine prescribing exemplifies a context in which reduced enzymatic function is translated into an initial dose modification and subsequent clinical management rather than a simple binary decision [14]. This form of recommendation depends on the relationship among genotype, enzyme activity, systemic exposure, toxicity risk, treatment intent, and the possibility of dose titration. It also reveals why functional categories should not be interpreted as universally interchangeable units. The same predicted degree of functional impairment may have different therapeutic implications across drugs, exposure ranges, disease settings, and monitoring environments.

The evidentiary status of a recommendation should remain distinguishable from the clinical utility of using it. The GIFT trial evaluated genotype-guided warfarin dosing through clinical events and anticoagulation control in a defined surgical population [15]. Such trials are important because they test an integrated intervention that includes genotyping, an algorithm, clinical timing, treatment delivery, and follow-up rather than evaluating a variant in isolation. Their findings remain conditional on the population, comparator, dosing strategy, clinical workflow, and outcomes examined. A favorable result does not establish that every genotype-guided dosing

strategy is beneficial, while an uncertain result in another setting would not invalidate the complete mechanistic evidence supporting the gene–drug relationship.

Table 1 summarizes the evidence states, relationships, uncertainties, and boundaries needed to connect pharmacogenomic association with functional and clinical evidence.

Table 1. Definitions, components, relationships, and intended uses for closing the gap between pharmacogenomic association and functional evidence

Component	Problem addressed	Key inputs	Proposed relationship	Contribution	Evidence required	Main risks	Boundary statement
Association credibility	Gene–drug associations may be unstable or population-specific.	Variant, exposure, phenotype, design, replication, population	Association is the entry state but cannot replace later evidence.	Identifies signals suitable for further evaluation.	Replication, controls, exposure and population definition, and relevant outcomes [1].	Confounding, stratification, phenotype heterogeneity, reporting bias	Association does not establish mechanism, prescribing action, or benefit.
Variant and haplotype specification	Incomplete genomic characterization may cause incorrect phenotype assignment.	Nomenclature, phase, diplotype, copy number, structural variation, assay coverage	Genomic resolution is required for patient-relevant functional interpretation.	Reduces ambiguity in the genomic input.	Governed nomenclature [10] and structural-variant detection [11].	Missed alleles, phasing errors, hybrid genes, assay gaps	Naming an allele does not confirm detection of all relevant configurations.
Functional resolvability	A genomic state may lack a defined biological effect.	Molecular assays, expression, activity, transport, pharmacokinetics, biomarkers	Convergent functional evidence links genomic state to biological effect.	Supports defensible functional interpretation.	Orthogonal, reproducible, biologically relevant, and concordant evidence.	Artificial systems, assay dependence, context effects, conflicting results	Functional alteration alone does not establish clinical utility.
Genotype-to-phenotype translation	The same diplotype may receive different phenotype assignments.	Allele function, activity values, translation rules, terminology	Standardized rules connect genotype to predicted phenotype [12].	Improves consistency across laboratories and guidelines.	Reliable allele calling, governed rules, and reassessment.	Outdated assignments, ambiguous alleles, unmeasured modifiers	Predicted phenotype is not the patient's directly measured phenotype.
Recommendation translatability	Biological evidence does not directly determine treatment.	Phenotype, outcomes, alternatives, treatment context	Guideline reasoning converts evidence into context-specific advice.	Supports selection, dosing, monitoring, or avoidance.	Guideline [13], dosing [14], and integrated clinical evidence [15].	Guideline disagreement, limited populations, changing alternatives	A recommendation is not an individualized prescription or regulatory decision.
Clinical utility	Actionability may be assumed without outcome improvement.	Intervention, comparator, adherence, workflow, outcomes, harms, resources	Utility depends on the full testing-to-treatment pathway.	Determines whether PGx use improves care.	Prospective comparative and implementation studies [5].	Underpowered studies, crossover, contamination, workflow failure	Utility in one setting does not establish universal effectiveness.
Evidence-state separation	Distinct evidence types may be collapsed into one hierarchy.	Association, function, phenotype, recommendation, implementation, utility	Proposed construct: each evidence state answers a	Reveals gaps within the evidence chain.	Comparison of complete and incomplete evidence pathways.	False confidence from combining non-equivalent evidence	These states are not a validated readiness score.

			different question.				
Uncertainty and decision boundary	Concise clinical outputs may omit uncertainty.	Evidence quality, discordance, missing data, population fit, alternatives	Each evidence state should state both its interpretation and limitations.	Preserves uncertainty through translation and implementation.	Auditable provenance and explicit uncertainty representation [8].	Lost caveats, overinterpretation, outdated assumptions	Recording uncertainty does not resolve it.

Patient-level interpretation and contextual modifiers

Patient-level interpretation extends beyond guideline recommendations because genotype-derived phenotype may differ from current metabolic activity. Phenoconversion occurs when clinical or environmental factors alter function relative to genotype prediction; CYP2C19 studies show that drug–drug interactions can produce this divergence [16]. Genotype therefore remains informative but may be insufficient when inhibitors, inducers, inflammation, or other modifiers change metabolism.

Interpretation must also separate established single-gene relationships from emerging polygenic approaches. PRS-PGx methods show that distributed variants can be combined to predict drug response [17], but their calibration, interpretability, ancestry transferability, and added clinical value require treatment-specific evaluation. Polygenic scores should not be assumed to be superior to established gene–drug evidence.

Ancestry further affects evidence transportability because allele frequencies, linkage patterns, structural variation, and pharmacogene combinations differ across populations [18]. This supports explicit transferability assessment but not deterministic conclusions about individuals. Genetic ancestry, race, ethnicity, geography, and social exposure are related but distinct constructs and should be used only with clear scientific justification.

Proposed missing-middle evidence framework

Debates about pharmacogenomic implementation frequently reflect different assumptions about what constitutes sufficient evidence. Some stakeholders emphasize mechanistic plausibility and avoidance of preventable toxicity, whereas others require prospective evidence of improved clinical outcomes, cost-effectiveness, or health-system benefit before broad implementation. These positions are not always mutually exclusive because they may address different gene–drug pairs, decisions, risks, and evidentiary thresholds. Examination of implementation evidence has shown that the meaning of “sufficient” depends on the clinical question and intended use [19]. The missing middle persists when these differing evidentiary expectations remain implicit or are collapsed into a single label such as actionable, validated, or ready.

The PGx-EB is proposed as a horizontal architecture comprising association credibility, variant and haplotype resolution, functional resolvability, genotype-to-phenotype translation, recommendation translatability, patient-context compatibility, and decision-use readiness. Each state receives an evidentiary input, produces a bounded interpretation, and identifies residual uncertainty. A transition gate determines whether the output can move forward, should be restricted, or should return to an earlier state for reassessment. The framework does not require every gene–drug pair to possess identical evidence. It requires only that the evidentiary question at each necessary transition be addressed rather than silently substituted by evidence from another state.

Patient-specific computational representations may support this architecture by combining stable genomic information with changing clinical data. Digital twins in precision medicine have been described as potentially dynamic representations requiring mechanistic coherence, multiscale integration, verification, validation, and explicit uncertainty management [20]. A pharmacogenomic digital twin would therefore be conceptually broader than a stored genotype or an automated guideline lookup. It would need to update relevant exposures, disease states, physiological variables, and treatment responses while retaining evidence provenance. The PGx-EB does not claim to constitute such a twin. It specifies evidentiary conditions that any patient-specific pharmacogenomic representation would need to satisfy before its outputs could be considered decision-relevant.

Figure 1 presents the original conceptual synthesis for evidence bridge from pharmacogenomic association to patient-level decision, showing how the article’s principal components, evidence relationships, uncertainties, and decision boundaries are connected.

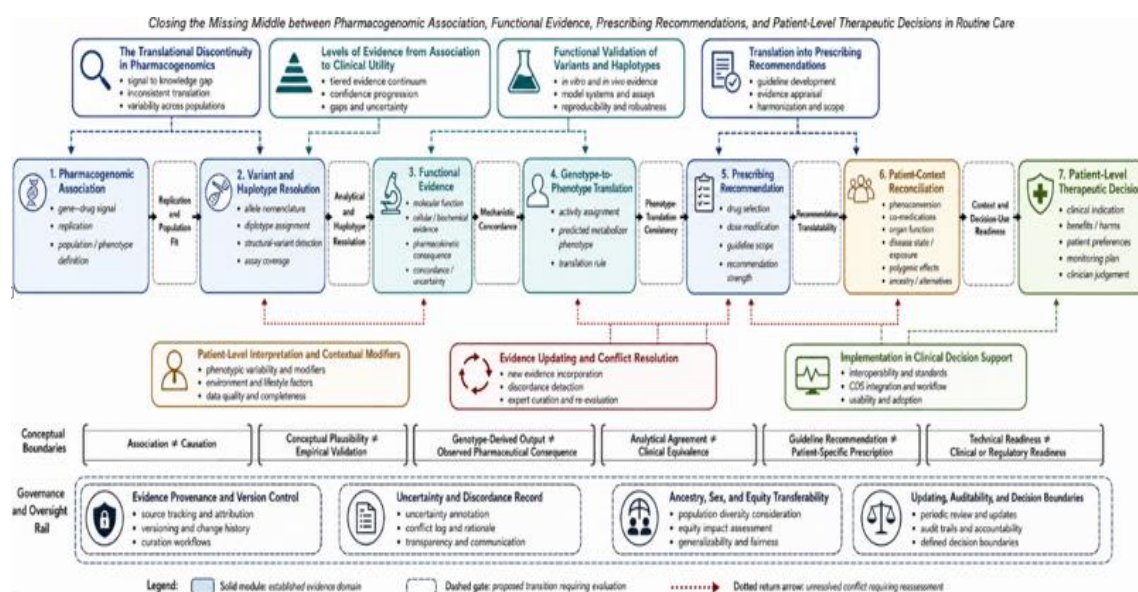


Figure 1. Evidence Bridge from Pharmacogenomic Association to Patient-Level Decision.

The figure is an original conceptual synthesis developed for “Closing the Missing Middle between Pharmacogenomic Association, Functional Evidence, Prescribing Recommendations, and Patient-Level Therapeutic Decisions in Routine Care.” Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

Drug–drug–gene models provide one example of the reconciliation required inside the bridge. Quantitative phenotype models can combine genotype-derived activity with inhibition or other exposure-related effects to estimate a context-dependent metabolic state [21]. Their usefulness depends on the accuracy of interaction parameters, concentration assumptions, phenotype definitions, and external validation. They can organize mechanistic information without automatically establishing the correct drug or dose for a patient. Within the PGx-EB, such models are treated as conditional evidence-processing tools whose assumptions, uncertainty, and applicable populations must remain visible.

Evidence updating and conflict resolution

Pharmacogenomic evidence is not static. Allele definitions are revised, structural configurations are discovered, functional assignments change, and previously uncertain variants may acquire additional evidence. PharmVar functions as a global repository for maintaining and communicating pharmacogene variation and nomenclature [22]. Such repositories support consistency across research, laboratories, guidelines, and clinical systems, but repository inclusion is not equivalent to evidence of clinical actionability. A versioned allele definition may improve traceability while leaving functional or therapeutic significance unresolved.

Guideline knowledge also requires recurring review and governance. CPIC’s development illustrates how clinical pharmacogenomic translation depends on structured evidence evaluation, standard terminology, guideline revision, dissemination, and computable resources [23]. Updating is not merely the addition of new publications. It may require reassessment of allele function, phenotype translation, recommendation strength, therapeutic alternatives, or the populations to which an interpretation applies. Because different components may change at different times, an update to one layer should trigger assessment of dependent downstream interpretations.

International growth further increases the importance of versioning and harmonization. Contemporary CPIC activities demonstrate continuing efforts to advance clinical pharmacogenomics through updated guidance, implementation resources, and global collaboration [24]. Nevertheless, expansion does not eliminate disagreement among laboratories, professional groups, health systems, or regulatory contexts. Conflict may concern the identity of an allele, its functional effect, phenotype translation, recommendation wording, or the sufficiency of clinical evidence. The PGx-EB proposes that such disagreement be represented as a structured state requiring adjudication rather than concealed through automatic selection of one source.

Conflict resolution should begin by identifying the level at which disagreement arises. Analytical conflict concerns whether the genomic state was called correctly. Functional conflict concerns the biological consequence

of that state. Translational conflict concerns phenotype assignment or recommendation formation. Contextual conflict concerns whether a recommendation applies to the current patient. Resolution may involve reanalysis, orthogonal testing, additional functional evidence, re-evaluation of guideline scope, specialist review, or retention of an explicitly unresolved status. No general rule should force discordant evidence into a single categorical answer when the basis for reconciliation is inadequate.

Implementation in clinical decision support

Clinical decision support is the operational layer through which pharmacogenomic evidence reaches the point of care. Its foundation is structured data representation. Results stored only in free-text reports may be difficult to retrieve, update, or link to future prescriptions. EHR implementation studies show how genomic indicators can be stored as reusable data for downstream decision logic [25]. However, structured storage does not ensure that an interpretation is current, complete, clinically relevant, or delivered at the right time.

Effective integration requires alignment among laboratory results, terminology, phenotype translation, medication data, guideline logic, and workflow. Experience with EHR genomics modules illustrates the mapping, governance, and maintenance needed for pharmacogenomic CDS [26]. A simple alert may therefore depend on a complex chain of assumptions. Changes in allele definitions, phenotype rules, or medication interactions may require reassessment of stored interpretations and active recommendations. Technical functionality alone does not confirm that these dependencies are managed correctly.

Human factors also shape CDS use. Clinician surveys highlight the importance of usability, perceived usefulness, workflow fit, and clarity [27]. These features influence acceptance, dismissal, and override but should not be confused with clinical utility. An acceptable alert may contain incomplete evidence, while a valid alert may fail if it appears too late, disrupts workflow, or gives unclear instructions.

Within the PGx-EB, decision-use readiness is the final gate before evidence informs treatment. It asks whether the result is timely, accurate, linked to the correct patient and medication, based on the current evidence version, understandable, and accompanied by uncertainty, alternatives, monitoring, and follow-up responsibilities. CDS can deliver and contextualize evidence, but it cannot resolve underlying genomic or functional uncertainty or replace the clinician's final judgment.

Limitations and research agenda

The framework is limited by biases in the evidence base. Polygenic models trained mainly in European-ancestry populations may perform less well elsewhere and could worsen disparities if transferred without validation [28]. Pharmacogenomic evidence also contains ancestry and sex biases affecting discovery, allele characterization, functional interpretation, and implementation [29]. These problems cannot be corrected simply by adding demographic variables. Research should address assay coverage, allele architecture, linkage patterns, environmental context, treatment access, and the distinction between genetic ancestry and socially structured determinants of health.

The PGx-EB also simplifies relationships that may be iterative, drug-specific, and biologically complex. Functional findings may reshape associations, clinical outcomes may revise phenotype rules, and unexpected responses may reveal new variants or interactions. The bridge preserves inferential order rather than implying that discovery is strictly linear. Gate relevance may also differ across severe toxicities, preventive prescribing, rare variants, established monogenic relationships, and emerging polygenic models. Empirical studies should determine which transitions are essential, which can be combined, and when uncertainty should block or only qualify progression.

CONCLUSION

The Pharmacogenomic Evidence Bridge reframes the movement from pharmacogenomic association to patient-level therapeutic decision as a series of distinct, non-substitutable evidentiary transitions. It separates association credibility, variant and haplotype resolution, functional interpretation, genotype-to-phenotype translation, prescribing-recommendation formation, patient-context compatibility, and decision-use readiness while adding cross-cutting requirements for provenance, updating, conflict resolution, uncertainty, and equity. Its principal contribution is not a new clinical rule but an original architecture for identifying where translation can fail and what evidence would be required to examine those failures. The framework preserves the difference between a genomic association, a functional inference, a guideline recommendation, a computational output, and a

therapeutic decision made for an individual patient. It remains a proposed conceptual synthesis requiring analytical, experimental, clinical, implementation, and equity-focused validation. It does not establish clinical utility, authorize automated prescribing, determine regulatory acceptability, or claim universal applicability. Closing the missing middle will depend not on treating genotype as sufficient, but on making every transition between evidence and decision visible, testable, updateable, and appropriately bounded.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Roden DM, McLeod HL, Relling MV, Williams MS, Mensah GA, Peterson JF, et al. Pharmacogenomics. *Lancet*. 2019;394(10197):521-32. doi:10.1016/S0140-6736(19)31276-0
2. Volpi S, Bult CJ, Chisholm RL, Deverka PA, Ginsburg GS, Jacob HJ, et al. Research directions in the clinical implementation of pharmacogenomics: An overview of US programs and projects. *Clin Pharmacol Ther*. 2018;103(5):778-86. doi:10.1002/cpt.1048
3. Lauschke VM, Ingelman-Sundberg M. Emerging strategies to bridge the gap between pharmacogenomic research and its clinical implementation. *NPJ Genom Med*. 2020;5:9. doi:10.1038/s41525-020-0119-2
4. van der Wouden CH, Cambon-Thomsen A, Cecchin E, Cheung KC, Dávila-Fajardo CL, Deneer VHM, et al. Implementing pharmacogenomics in Europe: Design and implementation strategy of the Ubiquitous Pharmacogenomics Consortium. *Clin Pharmacol Ther*. 2017;101(3):341-58. doi:10.1002/cpt.602
5. Swen JJ, van der Wouden CH, Manson LEN, Abdullah-Koolmees H, Blagec K, Blagus T, et al. A 12-gene pharmacogenetic panel to prevent adverse drug reactions: An open-label, multicentre, controlled, cluster-randomised crossover implementation study. *Lancet*. 2023;401(10374):347-56. doi:10.1016/S0140-6736(22)01841-4
6. Pereira NL, Farkouh ME, So D, Lennon RJ, Geller N, Mathew V, et al. Effect of genotype-guided oral P2Y12 inhibitor selection vs conventional clopidogrel therapy on ischemic outcomes after percutaneous coronary intervention: The TAILOR-PCI randomized clinical trial. *JAMA*. 2020;324(8):761-71. doi:10.1001/jama.2020.12443
7. Caudle KE, Dunnenberger HM, Freimuth RR, Peterson JF, Burlison JD, Whirl-Carrillo M, et al. Standardizing terms for clinical pharmacogenetic test results: Consensus terms from the Clinical Pharmacogenetics Implementation Consortium. *Genet Med*. 2017;19(2):215-23. doi:10.1038/gim.2016.87
8. Whirl-Carrillo M, Huddart R, Gong L, Sangkuhl K, Thorn CF, Whaley R, et al. An evidence-based framework for evaluating pharmacogenomics knowledge for personalized medicine. *Clin Pharmacol Ther*. 2021;110(3):563-72. doi:10.1002/cpt.2350
9. Bank PCD, Caudle KE, Swen JJ, Gammal RS, Whirl-Carrillo M, Klein TE, et al. A comparison of the guidelines of the Clinical Pharmacogenetics Implementation Consortium and the Dutch Pharmacogenetics Working Group. *Clin Pharmacol Ther*. 2018;103(4):599-618. doi:10.1002/cpt.762
10. Gaedigk A, Ingelman-Sundberg M, Miller NA, Leeder JS, Whirl-Carrillo M, Klein TE. The Pharmacogene Variation Consortium: Incorporation of the Human Cytochrome P450 allele nomenclature database. *Clin Pharmacol Ther*. 2018;103(3):399-401. doi:10.1002/cpt.910
11. Dalton R, Lee SB, Claw KG, Prasad B, Phillips BR, Shen DD, et al. Interrogation of CYP2D6 structural variant alleles improves the correlation between CYP2D6 genotype and CYP2D6-mediated metabolic activity. *Clin Transl Sci*. 2020;13(1):147-56. doi:10.1111/cts.12695
12. Caudle KE, Sangkuhl K, Whirl-Carrillo M, Swen JJ, Haidar CE, Klein TE, et al. Standardizing CYP2D6 genotype to phenotype translation: Consensus recommendations from the Clinical Pharmacogenetics Implementation Consortium and Dutch Pharmacogenetics Working Group. *Clin Transl Sci*. 2020;13(1):116-24. doi:10.1111/cts.12692

13. Lee CR, Luzum JA, Sangkuhl K, Gammal RS, Sabatine MS, Stein CM, et al. Clinical Pharmacogenetics Implementation Consortium guideline for CYP2C19 genotype and clopidogrel therapy: 2022 update. *Clin Pharmacol Ther.* 2022;112(5):959-67. doi:10.1002/cpt.2526
14. Amstutz U, Henricks LM, Offer SM, Barbarino J, Schellens JHM, Swen JJ, et al. Clinical Pharmacogenetics Implementation Consortium guideline for dihydropyrimidine dehydrogenase genotype and fluoropyrimidine dosing: 2017 update. *Clin Pharmacol Ther.* 2018;103(2):210-16. doi:10.1002/cpt.911
15. Gage BF, Bass AR, Lin H, Woller SC, Stevens SM, Al-Hammadi N, et al. Effect of genotype-guided warfarin dosing on clinical events and anticoagulation control among patients undergoing hip or knee arthroplasty: The GIFT randomized clinical trial. *JAMA.* 2017;318(12):1115-24. doi:10.1001/jama.2017.11469
16. Abouir K, Exquis N, Gloor Y, Daali Y, Samer CF. Phenoconversion due to drug-drug interactions in CYP2C19 genotyped healthy volunteers. *Clin Pharmacol Ther.* 2024;116(4):1121-9. doi:10.1002/cpt.3378
17. Zhai S, Zhang H, Mehrotra DV, Shen J. Pharmacogenomics polygenic risk score for drug response prediction using PRS-PGx methods. *Nat Commun.* 2022;13(1):5278. doi:10.1038/s41467-022-32407-9
18. Karamperis K, Katz S, Melograna F, Ganau FP, Van Steen K, Patrinos GP, et al. Genetic ancestry in population pharmacogenomics unravels distinct geographical patterns related to drug toxicity. *iScience.* 2024;27(10):110916. doi:10.1016/j.isci.2024.110916
19. Luzum JA, Petry N, Taylor AK, Van Driest SL, Dunnenberger HM, Cavallari LH. Moving pharmacogenetics into practice: It's all about the evidence! *Clin Pharmacol Ther.* 2021;110(3):649-61. doi:10.1002/cpt.2327
20. De Domenico M, Allegri L, Caldarelli G, d'Andrea V, Di Camillo B, Rocha LM, et al. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *NPJ Digit Med.* 2025;8(1):37. doi:10.1038/s41746-024-01402-3
21. Viviani R, Berres J, Stingl JC. Phenotypic models of drug-drug-gene interactions mediated by cytochrome drug-metabolizing enzymes. *Clin Pharmacol Ther.* 2024;116(3):592-601. doi:10.1002/cpt.3188
22. Gaedigk A, Casey ST, Whirl-Carrillo M, Miller NA, Klein TE. Pharmacogene Variation Consortium: A global resource and repository for pharmacogene variation. *Clin Pharmacol Ther.* 2021;110(3):542-5. doi:10.1002/cpt.2321
23. Relling MV, Klein TE, Gammal RS, Whirl-Carrillo M, Hoffman JM, Caudle KE. The Clinical Pharmacogenetics Implementation Consortium: 10 years later. *Clin Pharmacol Ther.* 2020;107(1):171-5. doi:10.1002/cpt.1651
24. Caudle KE, Whirl-Carrillo M, Relling MV, Hoffman JM, Donnelly RS, Haidar CE, et al. Advancing clinical pharmacogenomics worldwide through the Clinical Pharmacogenetics Implementation Consortium (CPIC). *Clin Pharmacol Ther.* 2025;118(6):1512-22. doi:10.1002/cpt.70005
25. Caraballo PJ, Sutton JA, Giri J, Wright JA, Nicholson WT, Kullo IJ, et al. Integrating pharmacogenomics into the electronic health record by implementing genomic indicators. *J Am Med Inform Assoc.* 2020;27(1):154-8. doi:10.1093/jamia/ocz177
26. Hall BT, Eken E, Cavallari LH, Duarte JD, Wiisanen KW, Cicali EJ, et al. Implementing pharmacogenomics clinical decision support: A comprehensive tutorial on how to integrate the Epic Genomics Module. *Clin Pharmacol Ther.* 2025;117(6):1670-81. doi:10.1002/cpt.3599
27. Grant CW, Marrero-Polanco J, Joyce JB, Barry B, Stillwell A, Kruger K, et al. Pharmacogenomic augmented machine learning in electronic health record alerts: A health system-wide usability survey of clinicians. *Clin Transl Sci.* 2024;17(10). doi:10.1111/cts.70044
28. Martin AR, Kanai M, Kamatani Y, Okada Y, Neale BM, Daly MJ. Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat Genet.* 2019;51(4):584-91. doi:10.1038/s41588-019-0379-x
29. Corpas M, Siddiqui MK, Soremekun O, Mathur R, Gill D, Fatumo S. Addressing ancestry and sex bias in pharmacogenomics. *Annu Rev Pharmacol Toxicol.* 2024;64:53-64. doi:10.1146/annurev-pharmtox-030823-111731