Available online **www.ijpras.com**

International Journal of Pharmaceutical Research & Allied Sciences, 2025, 14(2):1-11
https://doi.org/10.51847/xEyWSpKNjb

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Causal Precision Pharmacotherapy from Genotype to Dose through Phenoconversion, Treatment History, Time-Varying Clinical Context, and Competing Therapeutic Goals

Peter Nilsson[1]*, Eva Johansson[1], Lars Andersson[2]

[1] Department of Causal Precision Pharmacotherapy and Phenoconversion, Faculty of Pharmacy, Karolinska Institute, Stockholm, Sweden.
[2] Department of Time-Varying Clinical Context and Dosing, Faculty of Pharmacy, Lund University, Lund, Sweden.

***Email:** peter.nilsson@ki.se*

ABSTRACT

Pharmacogenomics can identify inherited determinants of drug disposition, response, and toxicity, but genotype alone cannot specify an optimal treatment for a patient whose functional phenotype, disease state, concomitant therapy, prior exposure, therapeutic priorities, and available options change over time. This Original Causal Decision Architecture Article proposes a Causal Precision Pharmacotherapy Decision Architecture that separates germline information from baseline functional phenotype, dynamically updates phenotype through phenoconversion, represents treatment history and current clinical context as a time-indexed patient state, and compares feasible drug-dose-monitoring strategies through counterfactual outcome profiles. The architecture treats efficacy, toxicity, feasibility, treatment burden, and equity as competing rather than interchangeable objectives. It introduces an explicit uncertainty layer through which missing decision-critical information, weak transportability, causal ambiguity, model miscalibration, or unresolved trade-offs may trigger additional measurement, intensified monitoring, restricted use, or abstention. Counterfactual dose and therapy comparison is framed as an estimand-specific process rather than selection of the option with the most favorable prognostic prediction. Validation is distributed across genotype interpretation, dynamic phenotype updating, causal identification, model verification, calibration, temporal and external transportability, clinical utility, human-system interaction, and governance. The original contribution is the integration of these elements into a non-substitutable architecture in which no genomic, pharmacokinetic, computational, or clinical component is sufficient alone. The proposal requires computational, experimental, clinical, implementation, and equity-focused validation. It does not establish a validated predictive model, universal treatment rule, clinical recommendation, regulatory pathway, or deployment-ready pharmacogenomic digital twin.

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

Received: 25 January 2025; **Revised:** 27 March 2025; **Accepted:** 29 March 2025

INTRODUCTION

Pharmacogenomics has established that inherited variation can influence drug metabolism, transport, target interaction, efficacy, and toxicity. Its clinical promise is not that DNA determines treatment independently, but that genomic information can reduce uncertainty within a broader therapeutic decision. Even where a gene-drug relationship is well characterized, the resulting information remains conditional on indication, population, treatment alternatives, organ function, co-medications, and the clinical consequence being considered [1].

The clinical relevance of this distinction is demonstrated by implementation evidence showing that pharmacogenetic information can improve prescribing when it is connected to actionable recommendations and delivered through an organized care pathway. Pre-emptive panel-guided prescribing has reduced clinically relevant adverse drug reactions when genotype results, prescribing recommendations, professional education, and workflow integration operate together [2]. Utility emerges from interaction among a biological marker, an interpretable gene-drug relationship, an appropriate decision, and a functioning implementation environment.

Translation of genomic knowledge into care also depends on informatics infrastructure, standardized interpretation, professional capability, laboratory processes, organizational resources, reimbursement, governance, and knowledge updating. Clinical genomics is consequently not reducible to laboratory testing; it is a distributed decision process connecting evidence, computation, communication, workflow, and accountability [3].

This article develops an original, non-empirical Causal Precision Pharmacotherapy Decision Architecture, termed the CPPDA. The architecture explains how inherited variation may be connected to dose or therapy only through intermediate causal and decision layers: baseline functional phenotype, phenoconversion, treatment history, time-varying patient state, feasible options, competing goals, counterfactual comparison, uncertainty management, and governance. It defines components and relationships, derives testable propositions, and specifies validation boundaries without claiming empirical confirmation, clinical superiority, universal applicability, regulatory acceptance, or deployment readiness.

Why genotype alone cannot determine treatment

The first limitation of a genotype-only model is that genotype is not identical to functional phenotype. A diplotype must be translated through allele-function assignments, activity scores, structural-variant interpretation, laboratory quality controls, and gene-specific rules before a predicted phenotype can be inferred. Consensus work on CYP2D6 demonstrates that even this intermediate mapping requires standardization because laboratories and guideline groups may otherwise assign different phenotypes to the same genotype [4]. Genotype is therefore an inherited input, whereas functional phenotype is an interpreted representation of expected biological activity.

The second limitation is that a dose is determined by more than metabolic classification. Model-informed precision dosing illustrates the broader information requirements of dose individualization: pharmacokinetic and pharmacodynamic behavior, age, body size, organ function, disease severity, measured concentrations, response, toxicity, adherence, and observation timing may all alter the defensible action. Quantitative models can integrate these determinants, but clinical use depends on validated models, reliable data, computational support, interpretable recommendations, and compatible workflows [5].

The third limitation concerns transferability. Genetic effects and polygenic scores can perform differently across populations because allele frequencies, linkage disequilibrium patterns, environmental exposures, healthcare structures, and representation in discovery datasets differ. Clinical use of predictors derived predominantly from one ancestry may therefore increase disparity when uncertainty is ignored or performance is assumed to be universal [6]. It requires transparent evaluation of ancestry-related transportability, dataset composition, population-specific uncertainty, and access to testing. Genotype-guided prescribing is consequently a starting constraint rather than a complete treatment rule because genomic interpretation, dynamic dosing, and clinical implementation each contribute non-substitutable information [1, 4, 5].

A genotype-only formulation also fails to distinguish among decision levels. Whether a variant changes enzyme activity is mechanistic; whether it predicts exposure is pharmacokinetic; whether an alternative dose changes benefit or harm is causal; and whether that alternative can be implemented equitably is clinical and organizational. The CPPDA therefore treats inherited variation as an upstream constraint whose relevance must be propagated through explicit intermediate layers. This interpretation does not diminish established gene-drug relationships; it defines when those relationships may or may not be sufficient for a particular decision.

Proposed causal precision-pharmacotherapy architecture

The CPPDA is organized around a causal question: among treatment strategies feasible at a defined decision time, which strategy would be expected to produce the most defensible balance of outcomes for this patient under stated assumptions? Prognostic models estimate an outcome conditional on observed information, whereas causal comparison concerns potential outcomes under alternative actions. Causal machine-learning methods can support treatment-effect estimation, but require a defined estimand, assumptions about assignment and confounding,

appropriate outcome timing, and validation addressing treatment-effect heterogeneity rather than predictive accuracy alone [7].

The architecture represents the patient through a scoped, time-indexed state rather than a static genomic profile. This state contains inherited variation, the current functional phenotype estimate, phenoconversion determinants, treatment history, prior response and toxicity, measured concentrations, organ function, disease activity, adherence information, environmental exposure, and uncertainty about each element. Digital-twin scholarship emphasizes explicit computational scope, implementation conditions, and regulatory context; such representations should not be treated as complete replicas of the person [8].

Figure 1 presents the original conceptual synthesis for causal path from genotype through dynamic phenotype to dose, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

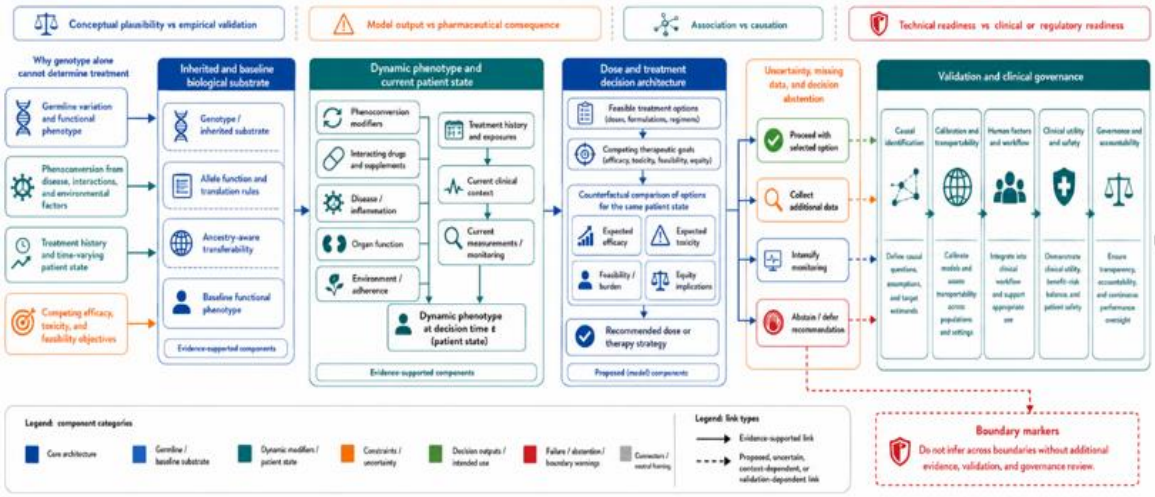

Figure 1. Causal Path from Genotype through Dynamic Phenotype to Dose

The first layer contains the inherited substrate and its interpretation. Germline variants constrain a baseline functional phenotype, accompanied by uncertainty from allele calling, structural variation, translation rules, ancestry transferability, and incomplete mechanistic knowledge. The third combines the updated phenotype with treatment history and the current patient state. Model-informed precision-dosing scholarship supports connecting changing clinical data with dose models, decision-support interfaces, and iterative monitoring rather than treating dose selection as a one-time calculation [9].

The fourth layer defines the feasible action set. An action may consist of a drug, dose, route, interval, formulation, monitoring plan, supportive intervention, or temporary deferral while additional information is collected. The fifth layer assigns each feasible action a multidimensional counterfactual profile. Efficacy, toxicity, feasibility, burden, and equity remain separate until a transparent process determines whether they can be traded. An uncertainty gate then supports selection, further measurement, intensified monitoring, restriction to a safer option, or abstention.

Germline variation and functional phenotype

Germline variation provides the relatively stable inherited substrate of the CPPDA. In some gene-drug contexts, validated variants have sufficient functional and clinical evidence to constrain an initial dosing strategy. DPYD-guided fluoropyrimidine recommendations illustrate how decreased-function variants can inform a reduced starting dose or avoidance decision within a defined therapeutic context [10]. The architecture treats such guidance as a bounded constraint rather than a universal rule.

A baseline functional phenotype represents expected activity before current phenoconversion factors are incorporated. This distinction matters because genotype-associated exposure differences may be clinically meaningful while leaving substantial unexplained variability. Evidence relating CYP2C19 genotype to escitalopram exposure and therapeutic failure supports a connection between inherited metabolic capacity and treatment outcome, but the retrospective, drug-specific setting does not establish a complete individualized policy [11].

The validity of germline interpretation also depends on population representation. Genome-wide association evidence, allele-function estimates, and polygenic models may not transport uniformly across ancestries because

linkage disequilibrium, allele frequency, genetic architecture, environmental context, and study design differ. Methodological recommendations for ancestrally diverse genomic studies emphasize diversity for both discovery and evaluation of transferability [12]. The architecture therefore attaches an ancestry-aware uncertainty statement to genomic inputs rather than assuming equivalent evidential strength in every population.

The proposed distinction among genotype, baseline functional phenotype, and dynamically observed phenotype creates a boundary between inherited constraint and current biological function. The baseline phenotype should be recorded with the translation method, reference version, laboratory assumptions, ancestry-related limitations, unresolved alleles, and confidence in the gene-drug relationship. This provenance allows later phenoconversion updates without implying that genotype has changed.

Phenoconversion from disease, interactions, and environmental factors

Phenoconversion describes divergence between genotype-predicted function and the phenotype expressed under a particular clinical condition. Enzyme inhibitors may reduce functional activity sufficiently for a genetically predicted normal metabolizer to behave more like an intermediate or poor metabolizer. A clinically oriented CYP2D6 tutorial shows why inhibitor exposure, drug potency, dosing, treatment duration, and substrate dependence must be incorporated when interpreting genotype-guided prescribing [13].

Prospective evidence supports the distinction between inherited and realized phenotype. In genotyped healthy volunteers, drug-drug interactions altered CYP2C19 activity despite unchanged germline background, demonstrating that a genotype-derived category does not fully describe enzyme function during interacting treatment [14]. Such findings support a phenoconversion layer but do not justify transferring the magnitude or consequence of an interaction to every population, substrate, disease state, or regimen.

Disease-related inflammation may suppress metabolic pathways through altered enzyme expression and regulatory signaling. Physiologically based pharmacokinetic modeling has represented combined effects of inflammation and drug interactions on CYP3A and CYP2C19 activity, illustrating how multiple causes may jointly modify exposure [15]. An inferred phenotype can diverge from the genotype-predicted phenotype through inhibitor exposure and inflammation-mediated modulation [4, 13, 15].

The phenoconversion layer should preserve causal provenance and uncertainty. Each update should identify the suspected modifier, timing, affected pathway, supporting evidence, expected duration, and conditions under which the baseline estimate may be restored. The layer is therefore an evidence-organizing mechanism rather than a validated conversion calculator. **Table 1** summarizes its components and boundaries.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in causal precision pharmacotherapy from genotype to dose through phenoconversion.

Table 1. Definitions, components, relationships, and intended uses in Causal Precision Pharmacotherapy from Genotype to Dose through Phenoconversion.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Germline substrate	Identifies inherited variation relevant to disposition, response, or toxicity	Validated genotype, haplotype, structural variation, allele function	Constrains the baseline functional estimate	More defensible initialization	Analytical validity, gene-drug evidence, ancestry-aware transportability [10]	Misclassification or unresolved alleles	Genotype is an upstream constraint, not an autonomous decision
Baseline functional phenotype	Translates genetic information into expected activity	Genotype, activity score, translation standard	Standardized mapping precedes dynamic updating [4]	Consistent interpretation	Interlaboratory concordance and functional validation	Discordant translation rules	Predicted phenotype is an inference, not direct measurement
Phenoconversion updater	Captures divergence between inherited	Inhibitors, inflammation, disease, organ	Time-varying modifiers update	More relevant current-state estimate	Prospective interaction, exposure, and	Omitted modifiers or substrate mismatch	The update is context-dependent and does not

	prediction and current function	function, environment	phenotype [13]		response evidence		change genotype
Treatment-history ledger	Prevents reliance on the current prescription alone	Previous drugs, doses, duration, response, toxicity, adherence	Historical exposure alters current interpretation	Recognition of path dependence	Longitudinal medication and outcome data	Incomplete records or recall bias	History informs state but does not establish causation
Counterfactual comparison	Distinguishes treatment-effect estimation from prognosis	Options, outcomes, horizon, confounders, effect modifiers	Potential outcomes are compared under causal assumptions [7]	More defensible option comparison	Target-trial specification and causal validation	Confounding, time-zero error, positivity failure	A counterfactual estimate is not an observed individual effect
Uncertainty and abstention	Prevents unsupported precision claims	Data quality, causal uncertainty, transportability, calibration	High uncertainty triggers monitoring, new data, restriction, or abstention	Explicit safety boundary	Context-specific uncertainty validation	False confidence or excessive non-action	No universal abstention threshold is proposed

Treatment history and time-varying patient state

Treatment is a sequence of decisions rather than a single baseline event. Dynamic treatment regimes formalize this principle by mapping accumulated patient information to actions at successive decision points [16]. In causal precision pharmacotherapy, relevant history may include prior drug selection, dose escalation or reduction, interruption, adherence, measured exposure, response, toxicity, interacting medication, and disease progression.

The time-varying state should distinguish information available before a decision from information produced after treatment. Measured concentrations, organ function, microbiological findings, symptoms, adverse effects, and pharmacodynamic markers may update a dose when timing and measurement properties are adequate. Model-informed precision dosing for antibiotics illustrates how repeated concentrations and covariates can revise individualized dosing rather than relying solely on population averages or initial characteristics [17].

Treatment history also matters over longer periods. Longitudinal primary-care data show that patients may encounter multiple pharmacogenomically actionable medicines across the life course, supporting persistent and reusable genomic information [18]. Meaning changes with the prescribed drug, dose, indication, interacting therapy, clinical condition, and available alternative. A treatment-history ledger is therefore proposed to retain provenance and temporal context rather than merely recording whether a genomic result exists.

At each decision time, a prior state and action produce observations that inform the next state, but the observed outcome cannot be attributed to treatment without considering confounding, adherence, measurement error, and concurrent change. The history layer must also avoid improper conditioning on post-treatment variables when estimating effects

Competing efficacy, toxicity, and feasibility objectives

Precision pharmacotherapy is inherently multiobjective. A strategy that improves expected efficacy may increase toxicity, treatment burden, monitoring requirements, cost, or inequitable access. Conversely, the option with the lowest toxicity may be less effective or unavailable. Guidance concerning MT-RNR1 and aminoglycosides illustrates a context in which severe toxicity risk must be considered alongside effective antibacterial therapy and feasible alternatives [19].

Feasibility is not a secondary administrative issue but part of treatment. A theoretically favorable dose may require unavailable measurements, an alternative drug may be inaccessible, or monitoring may impose an unacceptable burden. Economic evaluations of pharmacogenetic testing show that cost-effectiveness depends on the gene-drug pair, testing strategy, pathway, comparator, horizon, and setting [20]. Feasibility must be assessed locally and remain distinct from predicted efficacy or toxicity.

Equity is likewise a substantive objective rather than a retrospective adjustment. Pharmacogenomic practice may reproduce disparities through unequal testing access, underrepresentation, inappropriate use of race, differential coverage, or decision-support systems performing poorly in particular populations. Scholarship on health equality

and pharmacogenomics emphasizes distinguishing genetic ancestry from racial categories and examining how social and healthcare structures shape implementation [21].

The architecture does not prescribe universal weights for efficacy, toxicity, feasibility, burden, and equity. Relative importance varies with indication, urgency, reversibility of harm, preference, alternatives, and horizon. The response to unresolved conflict may be shared deliberation, additional evidence, staged treatment, intensified monitoring, or abstention rather than automated selection.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for causal precision pharmacotherapy from genotype to dose through phenoconversion.

Table 2. Testable propositions, evidence requirements, and validation criteria for Causal Precision Pharmacotherapy from Genotype to Dose through Phenoconversion.

Architecture layer or process stage	Core function	Information or material flow	Interaction with other components	Validation criterion	Potential failure mode	Human or experimental responsibility	Readiness boundary
Germline interpretation	Establish a reproducible starting point	Laboratory result to allele function and phenotype	Supplies baseline phenotype	Concordance across validated assays and standards [4]	Allele-calling error	Laboratory scientists and guideline developers	No decision use when validity is unresolved
Dynamic phenotype updating	Modify expected function when current causes alter activity	Disease, interaction, environment, physiology to revised phenotype	Updates state used for dosing	Prospective agreement with functional or exposure measures [14]	Missed phenoconversion	Clinical pharmacologists and investigators	Mechanistic plausibility alone is insufficient
Longitudinal-state construction	Represent history at each decision time	Prior actions and outcomes to time-indexed state	Feeds dynamic treatment rules	Temporal completeness and robustness [16]	Truncated history or time-order error	Clinicians, data scientists, causal methodologists	More data do not automatically create a valid state
Feasible-option specification	Define clinically possible actions	Indication, access, formulation, monitoring, preference	Restricts counterfactual comparison	Multidisciplinary confirmation of meaningful alternatives	Omitted comparator	Treating team, pharmacist, patient	Optimality is conditional on the declared option set
Counterfactual estimator	Compare outcomes under alternatives	State, action, outcome, assumptions to effects	Receives options and competing outcomes	Explicit estimand, identification, robustness, external validation [7]	Confounding or extrapolation	Causal specialists and clinical experts	Prediction does not establish treatment-effect validity
Human-governance interface	Preserve accountability and override	Model output and provenance to deliberation	Oversees every architecture layer	Human-factors, workflow, accountability, and audit evaluation	Automation bias or responsibility diffusion	Clinicians, patients, informaticians, governance bodies	Technical functionality does not equal clinical readiness

Counterfactual dose and therapy comparison

Counterfactual comparison asks what would happen under alternative strategies for the same decision context. Because only one strategy is observed for an individual at a given time, unobserved potential outcomes must be estimated from randomized evidence, observational data under causal assumptions, mechanistic models, or combinations of sources. Target-trial emulation provides a framework for specifying eligibility, strategies, assignment, time zero, follow-up, outcome, and causal contrast when observational data are used [22].

The comparison may concern different drugs, doses of the same drug, a treatment-monitoring package, or delay or avoidance. Treatment-effect heterogeneity must not be inferred merely because patients have different baseline

risks. The PATH Statement distinguishes approaches estimating variation in absolute benefit from unsupported subgroup interpretation and emphasizes prespecification, appropriate modeling, and cautious evaluation [23].

Machine-learning methods may estimate conditional average treatment effects or rank options across profiles, but the individual causal effect remains unobservable because both potential outcomes cannot be observed in the same person. Work on individualized treatment-effect estimation identifies confounding, covariate shift, limited overlap, model selection, external validation, and absence of an individual-effect ground truth as central challenges [24].

A defensible comparison should declare the options, decision time, horizon, outcomes, competing events, estimand, assumptions, and target population. Sensitivity analyses should examine alternative state definitions, confounder sets, missing-data assumptions, and model specifications. When options are non-dominated across objectives, the architecture should expose rather than conceal the trade-off.

Uncertainty, missing data, and decision abstention

Uncertainty enters every CPPDA layer. Analytical uncertainty affects genotype calls; interpretive uncertainty affects allele function and phenotype assignment; mechanistic uncertainty affects phenoconversion; measurement uncertainty affects concentrations and clinical state; causal uncertainty affects comparisons; and implementation uncertainty affects whether an output can be delivered safely. Verification, validation, and uncertainty-quantification scholarship for precision-medicine digital twins emphasizes that credibility must be tied to a declared context of use rather than assessed as one abstract property [25].

Missingness is causal and operational rather than merely statistical. A missing concentration may reflect an unavailable assay, uncollected sample, early discharge, treatment discontinuation, clinician judgment, or patient refusal. These causes may carry information about prognosis, treatment, access, and feasibility. PROBAST distinguishes model performance from the risk that study design, participants, predictors, outcomes, or analysis produce misleading conclusions [26].

Calibration is essential when estimated risks or outcomes compare therapeutic options. A model may discriminate higher- and lower-risk patients while systematically overestimating or underestimating absolute outcomes, distorting benefit-harm comparison. Calibration has therefore been characterized as a central weakness of predictive analytics [27].

Decision abstention is not necessarily therapeutic inaction. It may mean confirming a genotype, measuring a concentration, resolving the medication list, treating an acute condition before reassessment, choosing a conservative option, intensifying monitoring, seeking specialist review, or declining to issue a model-based recommendation. Abstention rules must be proportionate to consequences and evaluated for both safety and inequity. **Table 3** summarizes the principal escalation pathways.

Table 3. Failure modes, boundary conditions, governance needs, and implementation priorities for Causal Precision Pharmacotherapy from Genotype to Dose through Phenoconversion.

Decision scenario or design trade-off	Trigger or context	Competing objective	Evidence threshold	Uncertainty treatment	Corrective or escalation action	Expected implication	What cannot be concluded
Unresolved allele	Incomplete or discordant genotype interpretation	Rapid decision versus analytical validity	Validated assay and interpretable translation [4]	Preserve genotype state as unresolved	Confirm testing or avoid genotype-dependent action	Prevents false phenotype assignment	The most likely allele cannot be treated as confirmed
Suspected phenoconversion	Inhibitor, inflammation, organ dysfunction, changing exposure	Immediate treatment versus current functional accuracy	Context-specific mechanistic and clinical evidence [13]	Represent baseline and updated phenotype separately	Review interactions, repeat measurement, modify monitoring	Clarifies genotype-phenotype divergence	Genotype-predicted phenotype cannot represent current activity automatically
Misaligned decision time	Treatment begins before baseline state is measured	Data availability versus causal validity	Common time zero and explicit	Treat post-treatment measurements	Redesign analysis or emulate a target trial	Reduces time-related bias	A misaligned association cannot establish effect

			assignment [22]	as treatment affected			
Poor calibration or drift	Changing population, practice, assay, or disease context	Stability versus adaptation	Context-specific calibration and temporal validation [27]	Report direction and consequence of error	Recalibrate or suspend use	Protects absolute benefit-harm comparison	Stable discrimination does not prove reliable probabilities
Conflicting patient goals	Benefit, toxicity, burden, fertility, cost, quality of life	Multiple legitimate objectives	Explicit preference elicitation	Preserve incomparable dimensions	Shared deliberation or staged decision	Makes value judgments visible	The architecture cannot infer universal preference weights
Excessive abstention	Uncertainty rules disproportionately exclude groups	Safety versus access and equity	Prospective evaluation of abstention and downstream outcomes	Audit abstention by population and context	Revise rules, improve data, provide alternative pathway	Balances harm prevention with access	More abstention does not automatically mean greater safety

Validation and clinical governance

Validation must proceed layer by layer before the CPPDA is evaluated as an integrated architecture. Analytical validity concerns genotype and measurements; functional validity concerns genotype-to-phenotype translation and phenoconversion; model validity concerns exposure, response, and counterfactual estimators; implementation validity concerns whether outputs reach appropriate users at the correct time. Multisite evidence from the eMERGE Network shows that pharmacogenomic decision-support designs and process outcomes vary across organizations [28].

Human-facing representation requires separate validation. Pharmacogenomic alerts must communicate the implicated drug, relevant result, clinical consequence, recommended action, provenance, and limiting conditions. Work developing guidance for healthcare professionals demonstrates the need for clear and consistent representation of pharmacogenomic alert information [29].

Observed clinician behavior is another distinct endpoint. Clinician response to pharmacogenetic alerts varies across gene-drug pairs and clinical contexts, indicating that alert exposure does not guarantee acceptance, appropriate action, or improved outcomes [30]. Counterfactual outputs require prespecified estimands, principled treatment-effect heterogeneity analysis, and external validation before decision use [22-24].

Clinical governance should define ownership of genomic interpretation, phenotype updating, medication reconciliation, model maintenance, alert design, override review, adverse-event surveillance, equity auditing, and retirement of obsolete components. It should specify when specialist confirmation is required and how conflicting patient and professional judgments are resolved. Prospective evaluation of utility, safety, workflow, distributive consequences, accountability, and unintended behavior remains necessary.

Limitations and implementation priorities

The CPPDA is limited by the evidence and observability available at each decision point. Important determinants may be incompletely measured, including adherence, diet, nonprescription exposure, disease fluctuation, social constraints, and preferences. Institutional implementation processes also vary. Multisite evaluation of pharmacogenetic testing for antidepressant therapy identified differences in workflows, resources, stakeholder engagement, and implementation determinants, showing that local adaptation is unavoidable [31].

A second limitation concerns the pharmacogenomic digital-twin metaphor. A computational patient representation is necessarily selective: it models particular subsystems, variables, scales, and horizons rather than reproducing the complete biological and social person. A complex-systems perspective emphasizes challenges involving multiscale coupling, changing states, partial observability, and model boundaries [32].

Figure 2 presents the original conceptual synthesis for counterfactual comparison of precision-treatment options, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

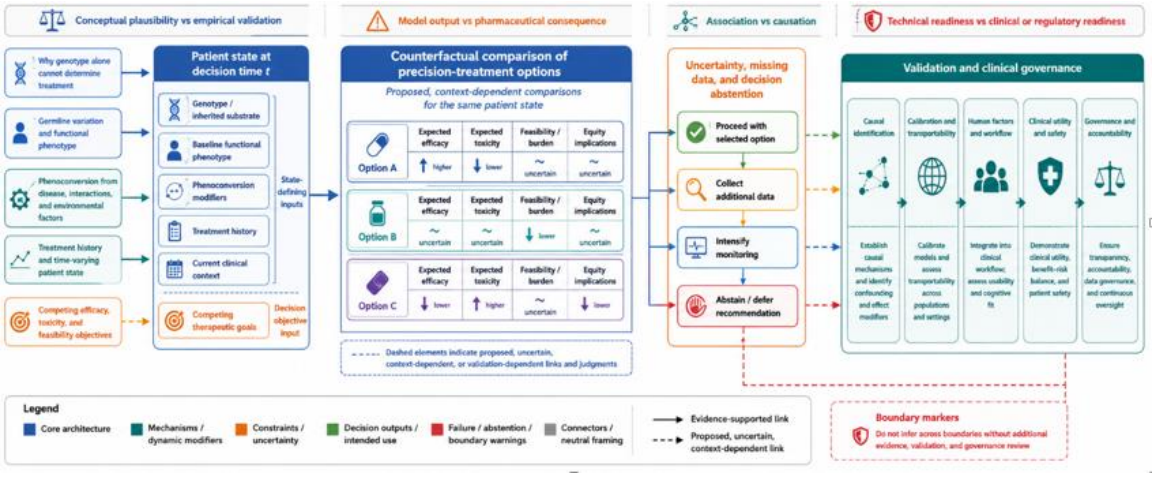

Figure 2. Counterfactual Comparison of Precision-Treatment Options

Implementation priorities follow from non-substitutability. Early studies should validate genotype interpretation and dynamic phenotype updating for specific gene-drug-disease contexts before testing combined policies. Subsequent work should compare estimators, examine temporal and ancestry transportability, assess calibration, and determine whether abstention reduces harmful overconfidence without worsening inequity. Human-factors studies should evaluate comprehension, automation bias, override behavior, responsibility, and presentation of multidimensional trade-offs.

The architecture should be revised when its layers do not add decision-relevant information, when dynamic updating performs worse than a simpler rule, when counterfactual estimates fail external validation, when uncertainty cannot be calibrated, or when implementation increases burden or disparity without corresponding benefit. The CPPDA is not intended to require maximal data or computational complexity. Its purpose is to identify the minimum causal, temporal, therapeutic, and governance information necessary for a defensible decision and to make unresolved insufficiency visible.

CONCLUSION

Causal precision pharmacotherapy requires more than connecting genotype to a drug or dose. The CPPDA separates inherited variation from functional phenotype, updates phenotype through disease, interactions, environment, and organ function, incorporates treatment history and time-varying clinical state, defines a feasible option set, preserves competing efficacy, toxicity, feasibility, burden, and equity objectives, and compares options through explicit counterfactual questions. An uncertainty and abstention layer prevents incomplete data or unsupported inference from becoming false precision, while staged validation and governance preserve human accountability. The architecture is an original conceptual synthesis intended to generate testable propositions and clarify boundaries. It does not establish a validated model, clinical recommendation, regulatory pathway, universal treatment policy, or deployment-ready pharmacogenomic digital twin.

The proposed architecture also establishes a practical research sequence: verify inherited inputs, test dynamic phenotype updates, define feasible alternatives, specify causal contrasts, assess calibration and transportability, evaluate patient-centered trade-offs, and audit real-world use. Progress at one layer cannot compensate for failure at another. Consequently, future studies should report not only whether a model produces an answer, but whether the answer is causally interpretable, temporally current, clinically actionable, equitably available, transparent to users, and appropriately withheld when current evidence is insufficient.

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. 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
3. Manolio TA, Rowley R, Williams MS, Roden D, Ginsburg GS, Bult C, et al. Opportunities, resources, and techniques for implementing genomics in clinical care. Lancet. 2019;394(10197):511-20. doi:10.1016/S0140-6736(19)31140-7
4. 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
5. Darwich AS, Ogungbenro K, Vinks AA, Powell JR, Reny JL, Marsousi N, et al. Why has model-informed precision dosing not yet become common clinical reality? Lessons from the past and a roadmap for the future. Clin Pharmacol Ther. 2017;101(5):646-56. doi:10.1002/cpt.659
6. 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
7. Feuerriegel S, Frauen D, Melnychuk V, Schweisthal J, Hess K, Curth A, et al. Causal machine learning for predicting treatment outcomes. Nat Med. 2024;30(4):958-68. doi:10.1038/s41591-024-02902-1
8. Venkatesh KP, Raza MM, Kvedar JC. Health digital twins as tools for precision medicine: Considerations for computation, implementation, and regulation. NPJ Digit Med. 2022;5(1):150. doi:10.1038/s41746-022-00694-7
9. Kluwe F, Michelet R, Mueller-Schoell A, Maier C, Klopp-Schulze L, van Dyk M, et al. Perspectives on model-informed precision dosing in the digital health era: Challenges, opportunities, and recommendations. Clin Pharmacol Ther. 2021;109(1):29-36. doi:10.1002/cpt.2049
10. Amstutz U, Henricks LM, Offer SM, Barbarino J, Schellens JHM, Swen JJ, et al. Clinical Pharmacogenetics Implementation Consortium (CPIC) guideline for dihydropyrimidine dehydrogenase genotype and fluoropyrimidine dosing: 2017 update. Clin Pharmacol Ther. 2018;103(2):210-6. doi:10.1002/cpt.911
11. Jukić MM, Haslemo T, Molden E, Ingelman-Sundberg M. Impact of CYP2C19 genotype on escitalopram exposure and therapeutic failure: A retrospective study based on 2,087 patients. Am J Psychiatry. 2018;175(5):463-70. doi:10.1176/appi.ajp.2017.17050550
12. Peterson RE, Kuchenbaecker K, Walters RK, Chen CY, Popejoy AB, Periyasamy S, et al. Genome-wide association studies in ancestrally diverse populations: Opportunities, methods, pitfalls, and recommendations. Cell. 2019;179(3):589-603. doi:10.1016/j.cell.2019.08.051
13. Cicali EJ, Elchynski AL, Cook KJ, Houder JT, Thomas CD, Smith DM, et al. How to integrate CYP2D6 phenoconversion into clinical pharmacogenetics: A tutorial. Clin Pharmacol Ther. 2021;110(3):677-87. doi:10.1002/cpt.2354
14. 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
15. Lenoir C, Niederer A, Rollason V, Desmeules JA, Daali Y, Samer CF. Prediction of cytochromes P450 3A and 2C19 modulation by both inflammation and drug interactions using physiologically based pharmacokinetics. CPT Pharmacometrics Syst Pharmacol. 2022;11(1):30-43. doi:10.1002/psp4.12730
16. Zhang Y, Laber EB, Davidian M, Tsiatis AA. Interpretable dynamic treatment regimes. J Am Stat Assoc. 2018;113(524):1541-9. doi:10.1080/01621459.2017.1345743
17. Wicha SG, Märtson AG, Nielsen EI, Koch BCP, Friberg LE, Alffenaar JWC, et al. From therapeutic drug monitoring to model-informed precision dosing for antibiotics. Clin Pharmacol Ther. 2021;109(4):928-41. doi:10.1002/cpt.2202
18. Kimpton JE, Carey IM, Threapleton CJD, Robinson A, Harris T, Cook DG, et al. Longitudinal exposure of English primary care patients to pharmacogenomic drugs: An analysis to inform design of pre-emptive pharmacogenomic testing. Br J Clin Pharmacol. 2019;85(12):2734-46. doi:10.1111/bcp.14100

19. McDermott JH, Wolf J, Hoshitsuki K, Huddart R, Caudle KE, Whirl-Carrillo M, et al. Clinical Pharmacogenetics Implementation Consortium guideline for the use of aminoglycosides based on MT-RNR1 genotype. Clin Pharmacol Ther. 2022;111(2):366-72. doi:10.1002/cpt.2309
20. Morris SA, Alsaidi AT, Verbyla A, Cruz A, Macfarlane C, Bauer J, et al. Cost effectiveness of pharmacogenetic testing for drugs with Clinical Pharmacogenetics Implementation Consortium (CPIC) guidelines: A systematic review. Clin Pharmacol Ther. 2022;112(6):1318-28. doi:10.1002/cpt.2754
21. Magavern EF, Gurdasani D, Ng FL, Lee SSJ. Health equality, race and pharmacogenomics. Br J Clin Pharmacol. 2022;88(1):27-33. doi:10.1111/bcp.14983
22. Hernán MA, Wang W, Leaf DE. Target trial emulation: A framework for causal inference from observational data. JAMA. 2022;328(24):2446-7. doi:10.1001/jama.2022.21383
23. Kent DM, Paulus JK, van Klaveren D, D'Agostino R, Goodman S, Hayward R, et al. The Predictive Approaches to Treatment effect Heterogeneity (PATH) statement. Ann Intern Med. 2020;172(1):35-45. doi:10.7326/M18-3667
24. Curth A, Peck RW, McKinney E, Weatherall J, van der Schaar M. Using machine learning to individualize treatment effect estimation: Challenges and opportunities. Clin Pharmacol Ther. 2024;115(4):710-9. doi:10.1002/cpt.3159
25. Sel K, Hawkins-Daarud A, Chaudhuri A, Osman D, Bahai A, Paydarfar D, et al. Survey and perspective on verification, validation, and uncertainty quantification of digital twins for precision medicine. NPJ Digit Med. 2025;8(1):40. doi:10.1038/s41746-025-01447-y
26. Wolff RF, Moons KGM, Riley RD, Whiting PF, Westwood M, Collins GS, et al. PROBAST: A tool to assess the risk of bias and applicability of prediction model studies. Ann Intern Med. 2019;170(1):51-8. doi:10.7326/M18-1376
27. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW; Topic Group 'Evaluating diagnostic tests and prediction models' of the STRATOS initiative. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
28. Herr TM, Peterson JF, Rasmussen LV, Caraballo PJ, Peissig PL, Starren JB. Pharmacogenomic clinical decision support design and multi-site process outcomes analysis in the eMERGE Network. J Am Med Inform Assoc. 2019;26(2):143-8. doi:10.1093/jamia/ocy156
29. Carter JPL, Critchlow J, Jackson S, Sanghvi S, Feger H, Chaudhry A, et al. Pharmacogenomic alerts: Developing guidance for use by healthcare professionals. Br J Clin Pharmacol. 2022;88(7):3201-10. doi:10.1111/bcp.15234
30. Lemke LK, Cicali EJ, Williams R, Nguyen KA, Cavallari LH, Wiisanen K. Clinician response to pharmacogenetic clinical decision support alerts. Clin Pharmacol Ther. 2023;114(6):1350-7. doi:10.1002/cpt.3051
31. Tuteja S, Salloum RG, Elchynski AL, Smith DM, Rowe E, Blake KV, et al. Multisite evaluation of institutional processes and implementation determinants for pharmacogenetic testing to guide antidepressant therapy. Clin Transl Sci. 2022;15(2):371-83. doi:10.1111/cts.13154
32. 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