Available online www.ijpras.com

International Journal of Pharmaceutical Research & Allied Sciences, 2026, 15(1):22-32
https://doi.org/10.51847/mP4NWbuTcp

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

The Pharmacogenomic Digital Twin for Counterfactual Therapy Selection under Genotype, Phenoconversion, Polypharmacy, Disease Evolution, and Changing Treatment Goals

Elena Fedorova[1]*, Alexey Volkov[1], Irina Morozova[2], Dmitry Kuznetsov[3]

[1]Department of Pharmacogenomic Digital Twins and Counterfactual Therapy, Faculty of Pharmacy, Novosibirsk State University, Novosibirsk, Russia.
[2]Department of Genotype and Phenoconversion, Faculty of Pharmaceutical Sciences, Ural Federal University, Yekaterinburg, Russia.
[3]Department of Polypharmacy and Disease Evolution, Faculty of Pharmacy, Kazan Federal University, Kazan, Russia.

***Email:** elena.fedorova@nsu.ru*

ABSTRACT

Pharmacogenomic decisions are commonly anchored to inherited variation, yet the functional consequences of genotype may change as medications, inflammation, organ function, disease severity, adherence, and treatment priorities evolve. This article proposes a Pharmacogenomic Digital Twin for Counterfactual Therapy Selection as an original, non-empirical architecture for representing these interacting determinants over time. The proposed twin separates relatively stable genomic information from time-varying functional phenotype, medication exposure, disease state, clinical observations, and treatment goals. It combines semantic harmonization, genotype-to-phenotype interpretation, phenoconversion assessment, pharmacokinetic and pharmacodynamic state estimation, counterfactual treatment simulation, constrained benefit–risk comparison, uncertainty characterization, and governed clinical updating. Competing therapeutic objectives are represented explicitly rather than collapsed into a single universal utility, allowing efficacy, safety, treatment burden, urgency, and patient preferences to influence conditional comparisons while preserving non-negotiable safety constraints. New concentrations, biomarkers, adverse effects, medication changes, and disease observations may update the estimated patient state, but population-model modification remains subject to separate change control. The architecture produces conditional treatment trajectories, uncertainty statements, requests for additional information, or abstention rather than autonomous prescriptions. Its original contribution is the integration of pharmacogenomic interpretation with longitudinal state updating and counterfactual comparison within one traceable decision framework. Validation would require component verification, temporal and external validation, subgroup calibration, causal assessment, workflow testing, prospective silent-mode evaluation, and postimplementation monitoring. Principal boundaries include incomplete haplotype inference, phenoconversion uncertainty, unobserved adherence, model misspecification, confounding, limited transportability, changing patient preferences, automation bias, and inequitable data or service access. The proposed architecture therefore establishes a testable research framework, not a validated clinical system.

Key words: *Digital twin, Pharmacogenomics, Precision pharmacotherapy, Genotype, Phenoconversion, Clinical implementation*

Received: 02 October 2025; **Revised:** 18 January 2026; **Accepted:** 21 January 2026

INTRODUCTION

Pharmacogenomics has progressed from explaining inherited variability in drug response toward structured clinical implementation, including prospective evaluation of pre-emptive multigene testing [1-3]. This progress establishes that genomic information can influence treatment decisions, but it does not imply that a genotype result remains an adequate representation of a patient's pharmacological state throughout treatment.

A conventional gene–drug recommendation typically begins with a test result, translates variants into haplotypes or alleles, assigns functional activity, and links the resulting phenotype to a prescribing action. This sequence is clinically useful, yet it can become incomplete when inhibitors, inducers, inflammation, organ dysfunction, adherence, polypharmacy, or disease progression change exposure and response after the original interpretation. Pharmacogenomics consequently informs precision pharmacotherapy without exhausting it [1].

The conceptual gap is not simply insufficient genomic resolution. It is the absence of an integrated representation that distinguishes stable inherited information from changing functional state, connects both to treatment alternatives, and revises the comparison when observations or therapeutic priorities change. Contemporary pharmacogenomics therefore requires attention not only to variant interpretation but also to infrastructure, evidence quality, clinical context, population transferability, and implementation boundaries [2]. Prospective panel implementation demonstrates the feasibility of embedding genotype-guided recommendations in care, but it does not establish a continuously updating model of the individual patient [3].

This Original Digital-Twin Architecture Article proposes the Pharmacogenomic Digital Twin for Counterfactual Therapy Selection, termed PGxDT-CTS. The architecture is an original conceptual synthesis rather than an empirical model, clinical guideline, validated software product, or autonomous prescribing system. Its purpose is to define the patient-state representation, processing layers, counterfactual comparisons, uncertainty controls, validation requirements, and governance boundaries needed to investigate whether dynamic pharmacogenomic decision support can outperform static genotype-only reasoning in defined contexts of use.

Why pharmacogenomic decisions require dynamic patient models

Patient-specific digital representations and model-informed precision dosing converge on the need to integrate individual data, mechanistic knowledge, longitudinal observations, and decision-specific simulation [4-6]. Digital-twin concepts emphasize a computational representation that changes with the represented individual, while model-informed dosing establishes the importance of linking patient covariates and observations to drug-exposure and response models. Neither concept alone provides the complete pharmacogenomic architecture proposed here.

Genotype is comparatively stable, but its pharmacological expression is conditional. The same diplotype may have different consequences when an interacting drug is initiated, an inflammatory episode suppresses enzyme activity, renal function deteriorates, adherence changes, or a new therapeutic objective alters the acceptable balance between benefit and harm. A dynamic patient model must therefore preserve genotype as one causal determinant without treating it as a permanent summary of functional phenotype.

The proposed patient state contains six coordinated domains: genomic state; functional phenotype; medication and intervention state; disease and physiological state; clinical observations; and treatment goals. A seventh domain records provenance, care context, model version, subgroup applicability, and governance status. Relationships among these domains are partly supported by existing digital-twin and precision-dosing scholarship, whereas their integration into a single pharmacogenomic state-updating system is an original proposal. **Figure 1** presents the original conceptual synthesis for pharmacogenomic digital twin and state-updating architecture, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

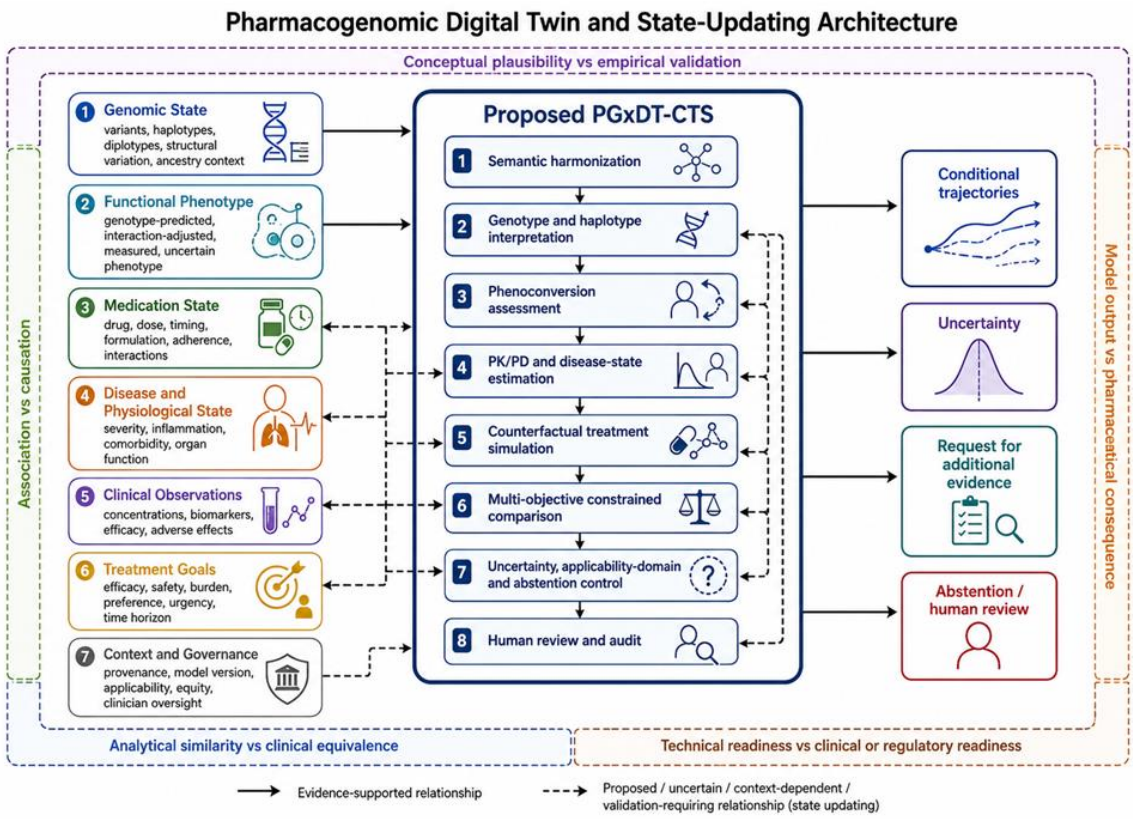

Figure 1. Pharmacogenomic Digital Twin and State-Updating Architecture

The need for updating does not mean that every new observation should immediately alter therapy. A concentration measurement may be mistimed, a medication list may be incomplete, a biomarker may be nonspecific, and a symptom may reflect disease rather than drug exposure. Model-informed precision dosing accordingly requires validated models, informative observations, suitable implementation infrastructure, and interpretation within a defined clinical context [6]. The twin must evaluate whether a new datum is reliable and decision-relevant before changing its state estimate.

A digital representation also becomes clinically hazardous when its resemblance to the patient is assumed rather than tested. Digital-twin scholarship provides a useful vision of individualized simulation, but conceptual correspondence does not establish predictive accuracy, causal validity, or clinical utility [4]. The proposed architecture therefore treats uncertainty, applicability boundaries, abstention, and human oversight as central functions rather than secondary additions.

Proposed pharmacogenomic digital-twin architecture

The PGxDT-CTS is defined as a proposed, version-controlled, patient-specific computational representation that integrates relatively stable genomic information with time-varying functional phenotype, medication exposure, disease state, clinical observations, and treatment goals. Its intended use is to compare conditional future treatment trajectories under explicitly represented uncertainty. It is not defined merely as an electronic record, a pharmacogenomic alert, a population model, or a graphical summary.

The first architectural layer performs semantic harmonization and provenance control. It records the assay, tested regions, variant-calling method, phasing assumptions, allele-definition version, phenotype-translation rule, observation time, data source, and missingness. Genomic clinical decision support depends on structured interpretation and delivery at the point of care rather than the passive availability of a test result [7]. The proposed twin extends this principle by preserving the provenance of every state transition and simulated comparison.

The second layer transforms genomic and clinical inputs into a time-indexed functional state. Pre-emptive pharmacogenomic implementation demonstrates that results can be transferred into structured health-system infrastructure and linked to executable prescribing support [8]. In the proposed architecture, however, stored genotype is only the beginning: medication interactions, disease modifiers, measured concentrations, organ function, and previous response may revise the currently estimated phenotype without changing the underlying genomic record.

The third layer links functional state to pharmacokinetic, pharmacodynamic, and disease representations and then simulates candidate therapies, doses, formulations, or monitoring strategies. Each branch begins from the same estimated patient state and returns a multidimensional trajectory rather than a single score. The comparison layer keeps efficacy, toxicity, disease control, treatment burden, uncertainty, and patient priorities distinguishable, while safety constraints may disqualify an otherwise favourable option.

The final layers govern uncertainty, applicability, abstention, and human review. Model credibility must be established for a declared context of use through appropriate verification, validation, and uncertainty assessment [9]. Accordingly, the PGxDT-CTS should decline to rank alternatives when essential data are missing, the patient lies outside the validated domain, component models disagree materially, causal assumptions are implausible, or uncertainty makes differences between trajectories clinically uninterpretable. The architecture is therefore a proposed decision-organizing mechanism, not evidence that any particular implementation is ready for clinical use.

Genotype, haplotype, and functional phenotype

The architecture begins by separating concepts that are often compressed into a single "pharmacogenomic result." A genotype records observed sequence states at specified loci. A haplotype represents variants inferred or measured on the same chromosome, while a diplotype combines the two inherited haplotypes for a gene. Allele function describes the expected activity associated with a haplotype, and phenotype expresses a clinically interpretable functional category. Standardized terminology is necessary for consistent reporting and clinical communication [10].

Translation across these layers is not lossless. Assay coverage may omit relevant variants, phasing may be uncertain, copy-number or hybrid structures may be unresolved, and allele-function assignments may change as knowledge develops. Even when the diplotype is correct, different activity assignments or phenotype thresholds can yield discordant interpretations. Consensus work on CYP2D6 demonstrates why genotype-to-phenotype translation rules must be explicit and standardized [11]. The proposed twin therefore retains the rule version and uncertainty associated with every translation rather than storing phenotype as an unqualified fact.

Haplotype representation must also be versioned and traceable. Curated pharmacogene nomenclature provides a shared basis for connecting variant combinations to named alleles, but the nomenclature itself does not guarantee that a clinical assay detected every relevant structural state or that an allele's function is known in every substrate context [12]. The PGxDT-CTS consequently stores the assay boundary, haplotype-call confidence, structural-variation status, and unresolved alternatives alongside the selected interpretation.

Functional phenotype is represented as time-indexed rather than permanently fixed. The architecture distinguishes genotype-predicted phenotype, interaction-adjusted phenotype, model-estimated latent phenotype, directly measured phenotype, and indeterminate phenotype. These categories are related but not interchangeable. Genomic evidence supplies a relatively stable prior representation; medications, disease, physiology, and observations may modify the estimated current function. This layered treatment preserves the scientific value of genotype while preventing genetic classification from being mistaken for a complete or continuously valid description of drug metabolism or response.

Phenoconversion and interaction-driven metabolic change

Phenoconversion occurs when non-genetic influences cause current metabolic activity to differ from the phenotype predicted from genotype. Enzyme inhibition, induction, polypharmacy, inflammation, and altered physiology may therefore change drug exposure without changing the inherited diplotype. Concomitant inhibitors and polypharmacy can produce clinically relevant discordance between genotype-predicted and functional metabolic phenotype, requiring explicit phenoconversion assessment [13-15].

The proposed architecture treats phenoconversion as a time-indexed modification rather than a permanent reassignment of genotype. When a relevant inhibitor, inducer, dose change, or interacting combination is detected, the twin re-evaluates functional phenotype and propagates the revised state into exposure and response models. The CYP2D6 tutorial provides an operational precedent for incorporating inhibitor effects into phenotype interpretation [13]. The proposed extension to a broader digital-twin architecture remains conditional on gene-specific evidence.

Direct evidence from CYP2C19-genotyped volunteers shows that inhibitor exposure can alter measured metabolic phenotype despite stable genotype [14]. Such evidence supports dynamic updating but does not justify universal

conversion factors. Effect magnitude may depend on the enzyme, substrate, inhibitor potency, dose, exposure duration, disease context, and measurement method.

Polypharmacy further creates simultaneous drug–drug and drug–drug–gene relationships that may not be captured by pairwise alerts [15]. The twin should therefore preserve the identity and timing of interacting mechanisms, distinguish inferred from measured phenotype, and abstain when competing interactions cannot be resolved. **Table 1** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in the pharmacogenomic digital twin for counterfactual therapy selection under genotype.

Table 1. Definitions, components, relationships, and intended uses in The Pharmacogenomic Digital Twin for Counterfactual Therapy Selection under Genotype.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Genomic state	Stable inherited information may be incompletely represented	Variants, haplotypes, diplotypes, copy number, assay provenance	Forms the relatively stable prior for functional interpretation	Traceable genomic foundation	Analytical validity and allele-definition concordance [10]	Missed variants, phasing error, unresolved structure	Genotype does not equal current phenotype
Haplotype and allele-function layer	Raw variants may not map directly to clinical function	Phasing, allele nomenclature, curated function assignments	Translates variant combinations into versioned functional categories	Interoperable interpretation	Nomenclature and function curation [12]	Version drift or uncertain allele function	Named alleles do not guarantee complete assay coverage
Genotype-predicted phenotype	Clinical interpretation may vary between translation systems	Diplotype, activity assignment, phenotype rule	Produces an initial functional category	Standardized starting state	Gene-specific translation validation [11]	Discordant thresholds or activity assignments	A predicted category is not a direct measurement
Phenoconverted phenotype	Current activity may differ from inherited prediction	Inhibitors, inducers, dose, timing, polypharmacy	Modifies functional state while retaining the underlying genotype	More context-sensitive exposure estimation	Prospective interaction and phenotype evidence [13, 14]	Overgeneralized conversion rules	Conversion remains gene-, drug-, and context-specific
Disease-modified phenotype	Inflammation or organ dysfunction may alter metabolism	Biomarkers, disease severity, organ function	Changes enzyme, transporter, PK, or PD state	Longitudinal clinical relevance	Mechanistic and longitudinal evidence	Nonspecific biomarkers and causal ambiguity	Disease state is an input, not proof of individual progression
Treatment-goal state—proposed construct	The preferred therapy depends on more than biological response	Safety tolerance, efficacy priority, burden, preference, urgency	Reweights or constrains conditional trajectories	Goal-sensitive comparison	Preference elicitation and outcome validation	Unstable or poorly recorded goals	No weighting scheme replaces clinical judgement
Uncertainty and abstention state—proposed construct	Forced rankings may conceal inadequate evidence	Data quality, model disagreement, domain shift, missingness	Blocks or qualifies output when confidence is insufficient	Safer decision support	Context-of-use-specific validation	False certainty or excessive abstention	Abstention is an output, not a treatment recommendation

Disease evolution and longitudinal patient state

Disease can modify pharmacological function independently of genotype. Inflammation may suppress or otherwise alter drug-metabolizing enzymes and transporters, changing exposure during infection, immune activation, or chronic inflammatory disease [16]. The architecture therefore represents disease activity and relevant biomarkers as time-varying modifiers rather than static diagnostic labels.

Multimorbidity and polypharmacy also alter the background against which pharmacogenomic information is interpreted. Prospective hospital evidence links adverse drug reactions with complex medication and morbidity profiles, although such observational relationships do not establish the cause of an individual event [17]. The twin should therefore update safety context when diseases, medications, or organ function change without attributing every change to pharmacogenomics.

Longitudinal state estimation may incorporate concentrations, laboratory measurements, symptoms, treatment response, and disease markers. Model-informed development of individualized daprodustat regimens illustrates how disease-related biomarkers and response models can inform adaptive dosing in chronic disease [18]. This example supports component plausibility but does not validate the proposed PGxDT-CTS or establish transferability to other drugs.

The architecture must also represent alternative explanations. A rising concentration could reflect inhibition, declining clearance, altered adherence, sampling error, or incorrect dose documentation. State updating should therefore compare plausible mechanisms, retain unresolved uncertainty, and request additional observations when competing explanations would produce different treatment decisions.

Counterfactual therapy and dose simulation

Counterfactual therapy selection asks what may happen under alternatives not simultaneously observed in the same patient. This differs from ordinary prognostic prediction because a variable associated with poor outcome does not necessarily identify which therapy will improve that outcome. Methods for treatment-effect heterogeneity therefore require explicit distinction between baseline risk and differential treatment response [19]. A simulated trajectory becomes an individual causal claim only when assumptions concerning treatment assignment, confounding, consistency, positivity, model specification, and transportability are adequately addressed. Machine-learning predictions of observed outcomes cannot automatically identify outcomes under alternative interventions [20]. PGxDT-CTS outputs must therefore be described as conditional trajectories unless stronger causal evidence exists.

For each candidate therapy or dose, the proposed simulator would project exposure, expected response, toxicity dimensions, disease state, monitoring burden, and uncertainty over a declared time horizon. Bedside model-informed dosing demonstrates the potential value of combining models and patient observations, while also highlighting requirements for appropriate data, software, interpretation, and workflow [21].

The architecture generates testable propositions: dynamic phenotype should outperform an unchanged genotype category when important modifiers arise; stale or missing state variables should reduce ranking reliability; treatment-goal changes may alter the preferred trajectory; informative observations should improve calibration; and the system should abstain during material domain shift or unresolved model disagreement. **Table 2** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for the pharmacogenomic digital twin for counterfactual therapy selection under genotype.

Table 2. Testable propositions, evidence requirements, and validation criteria for The Pharmacogenomic Digital Twin for Counterfactual Therapy Selection under Genotype.

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
Genomic interpretation	Convert assay output into versioned alleles and phenotype prior	Assay → variant → haplotype → function	Supplies the initial state	Concordance with validated interpretation procedures [10]	Incorrect phase or copy number	Laboratory verification and expert adjudication	No use beyond validated assay coverage

Dynamic phenotype engine	Adjust function for interactions and disease	Genotype prior + medication and disease modifiers	Updates PK/PD modules	Better calibration than static phenotype when modifiers are present [13]	Inaccurate conversion magnitude	Prospective phenotype and interaction studies	Gene- and context-specific validation required
State estimator	Infer current patient-specific parameters	Observations + prior state → updated state	Feeds every simulation branch	Temporal calibration and robustness to noisy data	Overreaction to unreliable observations	Sampling-quality control and clinician review	No update from unverified data
Counterfactual simulator	Compare alternative therapies or doses	Common baseline → alternative trajectories	Uses causal, PK/PD, and disease models	External discrimination, calibration, and treatment-effect evaluation [19]	Prognostic association mistaken for treatment effect	Causal-methods review and comparative evaluation	Conditional simulation unless causal assumptions hold
Multi-objective comparison—proposed	Compare efficacy, safety, burden, and preference	Trajectories + goals + constraints	Produces qualified ranking	Stability under plausible goal and uncertainty changes	Hidden trade-offs or unstable ranking	Patient–clinician goal elicitation	No universal utility function
Uncertainty and abstention—proposed	Prevent forced output	Component uncertainty → decision uncertainty	May request data or block ranking	Appropriate coverage and safe abstention behaviour	False reassurance or indiscriminate abstention	Threshold setting for defined use	No ranking outside validated domain
Updating and governance	Control state revision and model change	New observation → patient update; new evidence → controlled model revision	Maintains provenance and auditability	Reproducibility, version traceability, drift detection [9]	Silent model drift	Multidisciplinary change control	Technical verification does not establish clinical utility

Competing therapeutic objectives and safety constraints

Therapy selection rarely optimizes one outcome. A regimen may improve efficacy while increasing toxicity, monitoring burden, interaction risk, or inconvenience. Personalized treatment rules can therefore be formulated as benefit optimization subject to explicit risk constraints rather than unrestricted maximization of one predicted endpoint [22].

Patient preferences may change the relative value of the same clinical outcomes. Preference-sensitive benefit–risk assessment has shown that individualized evaluations can differ when patients assign different importance to benefits and harms [23]. In PGxDT-CTS, treatment goals are therefore represented as a time-indexed state rather than a fixed population utility.

Weights and constraints serve different functions. A preference may alter the importance assigned to treatment burden, whereas a serious contraindication may remain non-negotiable. Methods that jointly assess efficacy and safety reinforce the need to preserve both dimensions rather than combining them prematurely [24]. The proposed architecture consequently screens safety constraints before preference-sensitive ranking.

The output should explain why one trajectory appears preferable, which goals drive the comparison, what uncertainty could reverse it, and which constraints exclude an option. No numerical utility or automated ranking is assumed to be universally valid. Where objectives remain irreconcilable, the appropriate result is structured human deliberation rather than artificial precision.

Continuous updating from clinical observations

Clinical observations give the twin its longitudinal character. Pharmacogenomic point-of-care decision support can alter prescribing behaviour, demonstrating that the timing and presentation of genomic information influence clinical action [25]. However, a changed prescription is not itself evidence of improved outcomes, and updating must extend beyond alert delivery.

Medication initiation, discontinuation, inhibitor exposure, measured concentrations, adverse effects, disease progression, and adherence information may update the patient state. Clinical integration of pharmacogenomics with drug–drug interaction screening illustrates the importance of connecting genomic interpretation to changing medication contexts and follow-up across care settings [26].

Patient-level updating must be distinguished from modification of the underlying population model. A new concentration may update an individual clearance estimate, whereas a new allele-function assignment or revised interaction model requires governed evidence review, validation, versioning, and change control. Recommendations for integrating population pharmacokinetic modelling with precision dosing similarly emphasize data quality, model evaluation, software integration, and continued monitoring [27].

An update should occur only when the observation is reliable, temporally relevant, and capable of changing a decision. Otherwise, the system may retain the previous state, increase uncertainty, or request confirmation. **Table 3** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for failure modes, boundary conditions, governance needs, and implementation priorities for the pharmacogenomic digital twin for counterfactual therapy selection under genotype.

Table 3. Failure modes, boundary conditions, governance needs, and implementation priorities for The Pharmacogenomic Digital Twin for Counterfactual Therapy Selection under Genotype.

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
Static versus updated phenotype	New inhibitor, inducer, inflammation, or organ change	Stability versus responsiveness	Gene- and context-specific evidence [13]	Retain both prior and adjusted states if unresolved	Pharmacology review or additional measurement	More explicit interpretation	That adjustment improves outcome
Immediate versus delayed update	Single noisy observation	Timeliness versus robustness	Verified timing and analytical quality	Increase uncertainty before changing state	Repeat or corroborate observation	Reduced overreaction	That every new datum is informative
Therapy ranking versus abstention	Overlapping trajectories or domain shift	Actionability versus false confidence	Context-of-use validation	Report uncertainty and ranking instability	Request data or defer to human review	Safer boundary control	That abstention identifies the best therapy
Automated alert versus clinical workflow	New PGx or interaction signal	Speed versus alert burden	Demonstrated usability and relevance [25]	Preserve rationale and provenance	Pharmacist or clinician review	Traceable decision support	That alert acceptance improves health
Patient update versus model revision	New individual observation or new scientific evidence	Personalization versus model stability	Separate validation for model changes [27]	Version patient and model states independently	Formal change control	Reproducibility	That local learning is externally valid
Broad deployment versus subgroup restriction	Sparse ancestry or population evidence	Access versus transportability	Subgroup calibration and external validation	Declare unsupported domains	Restrict use and generate evidence	Reduced hidden inequity	That ancestry alone resolves performance gaps
Preference-sensitive ranking versus safety constraint	Changed goals or tolerance	Autonomy versus harm prevention	Valid outcome estimates and documented preferences	Sensitivity analysis across plausible weights	Shared decision-making	Transparent trade-offs	That one utility represents every patient

Validation, uncertainty, and governance

Validation must be organized around intended use. A model used to organize data requires less evidence than one used to rank therapies where an error could cause serious harm. Credibility assessment should therefore scale with model influence, decision consequence, available evidence, verification quality, and uncertainty [28].

The proposed validation pathway begins with data and software verification, followed by validation of genotype interpretation, phenoconversion rules, PK/PD components, disease-state models, and observation-assimilation procedures. It then requires retrospective temporal evaluation, external and subgroup validation, counterfactual

assessment, prospective silent-mode testing, workflow evaluation, and clinical-utility studies. Successful performance in one component cannot compensate automatically for failure elsewhere.

Uncertainty should remain decomposable. The twin should distinguish measurement error, genotype ambiguity, parameter uncertainty, model-form uncertainty, causal uncertainty, transportability uncertainty, and uncertainty arising from changing treatment goals. Transparent reporting of data, modelling choices, intended use, evaluation procedures, limitations, and performance is necessary for independent assessment [29]. Reporting compliance, however, does not establish validity.

Governance requires a model registry, version control, provenance, access controls, audit logs, override documentation, incident review, drift surveillance, and predefined criteria for restriction or withdrawal. Changes to patient state may occur routinely, but changes to model structure or scientific knowledge require formal review. No single retrospective metric or technical checklist can establish clinical decision readiness.

Limitations and implementation priorities

The architecture inherits limitations from both pharmacogenomics and predictive modelling. Genetic effects may be estimated unevenly across populations, and uncritical transfer of genetic prediction can worsen disparities [30]. Although polygenic scores and pharmacogenomic variants are not equivalent, the broader lesson is directly relevant: ancestry-labelled inputs cannot substitute for representative data, subgroup calibration, environmental context, or equitable clinical access.

Implementation also depends on laboratory capacity, informatics, education, workflow design, reimbursement, evidence interpretation, and sustained multidisciplinary responsibility [31].

Early implementation should proceed through restricted contexts of use. Priorities are interoperable genomic representation, prospective phenoconversion studies, longitudinal multimodal datasets, causal treatment-comparison methods, external and subgroup validation, usable uncertainty communication, and silent-mode evaluation before decision influence. Clinical implementation research must test whether the architecture improves decisions and outcomes rather than merely increasing information volume.

Important unknowns remain. The optimal balance between mechanistic and data-driven models is context-dependent; some treatment effects may be weakly identifiable; rare haplotypes may lack evidence; patient goals may be incompletely elicited; and model updates may introduce new inequities. These limitations define the research programme and justify abstention, restricted applicability, and continued human responsibility.

CONCLUSION

The PGxDT-CTS is proposed as an original architecture for integrating inherited pharmacogenomic information with phenoconversion, polypharmacy, disease evolution, clinical observations, and changing treatment goals. Its central contribution is to separate stable genomic state from dynamic functional state and to compare conditional treatment trajectories through explicit efficacy, safety, burden, uncertainty, and governance layers. The framework does not establish that a digital twin can identify an individual's optimal therapy, infer causal effects from routine prediction, or replace clinical judgement. Its scientific value depends on falsifiable propositions, context-specific validation, transparent uncertainty, subgroup transportability, governed updating, and the ability to abstain when evidence is inadequate. It should therefore be treated as a research architecture for evaluating dynamic pharmacogenomic decision support rather than as a validated, universally applicable, or deployment-ready clinical system.

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. Pirmohamed M. Pharmacogenomics: Current status and future perspectives. Nat Rev Genet. 2023;24(6):350-62. doi:10.1038/s41576-022-00572-8
3. 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
4. Björnsson B, Borrebaeck C, Elander N, Gasslander T, Gawel DR, Gustafsson M, et al. Digital twins to personalize medicine. Genome Med. 2020;12(1):4. doi:10.1186/s13073-019-0701-3
5. Corral-Acero J, Margara F, Marciniak M, Rodero C, Loncaric F, Feng Y, et al. The 'Digital Twin' to enable the vision of precision cardiology. Eur Heart J. 2020;41(48):4556-64. doi:10.1093/eurheartj/ehaa159
6. Darwich AS, Polasek TM, Aronson JK, Ogungbenro K, Wright DFB, Achour B, et al. Model-informed precision dosing: Background, requirements, validation, implementation, and forward trajectory of individualizing drug therapy. Annu Rev Pharmacol Toxicol. 2021;61:225-45. doi:10.1146/annurev-pharmtox-033020-113257
7. Freimuth RR, Formea CM, Hoffman JM, Matey E, Peterson JF, Boyce RD. Implementing genomic clinical decision support for drug-based precision medicine. CPT Pharmacometrics Syst Pharmacol. 2017;6(3):153-5. doi:10.1002/psp4.12173
8. Tamraz B, Shin J, Khanna R, Van Ziffle J, Knowles S, Stregowski S, et al. Clinical implementation of preemptive pharmacogenomics testing for personalized medicine at an academic medical center. J Am Med Inform Assoc. 2025;32(3):566-71. doi:10.1093/jamia/ocae293
9. Musuamba FT, Skottheim Rusten I, Lesage R, Russo G, Bursi R, Emili L, et al. Scientific and regulatory evaluation of mechanistic in silico drug and disease models in drug development: Building model credibility. CPT Pharmacometrics Syst Pharmacol. 2021;10(8):804-25. doi:10.1002/psp4.12669
10. 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 (CPIC). Genet Med. 2017;19(2):215-23. doi:10.1038/gim.2016.87
11. 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
12. Gaedigk A, Ingelman-Sundberg M, Miller NA, Leeder JS, Whirl-Carrillo M, Klein TE. The Pharmacogene Variation (PharmVar) Consortium: Incorporation of the Human Cytochrome P450 (CYP) Allele Nomenclature Database. Clin Pharmacol Ther. 2018;103(3):399-401. doi:10.1002/cpt.910
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. Stingl JC, Viviani R. Pharmacogenetic guided drug therapy – how to deal with phenoconversion in polypharmacy? Expert Opin Drug Metab Toxicol. 2025;21(4):399-407. doi:10.1080/17425255.2025.2451440
16. Stanke-Labesque F, Gautier-Veyret E, Chhun S, Guilhaumou R; French Society of Pharmacology and Therapeutics. Inflammation is a major regulator of drug metabolizing enzymes and transporters: Consequences for the personalization of drug treatment. Pharmacol Ther. 2020;215:107627. doi:10.1016/j.pharmthera.2020.107627
17. Osanlou R, Walker L, Hughes DA, Burnside G, Pirmohamed M. Adverse drug reactions, multimorbidity and polypharmacy: A prospective analysis of 1 month of medical admissions. BMJ Open. 2022;12(7). doi:10.1136/bmjopen-2021-055551
18. Mahar KM, Yang S, Johnson B, Visser SAG, van den Berg P, van Noort M, et al. Model-informed drug development for daprodustat supports the design of individualized dosing regimens in chronic kidney disease patients with anemia. Clin Pharmacol Ther. 2026;119(2):322-35. doi:10.1002/cpt.70094

19. 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
20. Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. Nat Mach Intell. 2020;2(7):369-75. doi:10.1038/s42256-020-0197-y
21. Keizer RJ, ter Heine R, Frymoyer A, Lesko LJ, Mangat R, Goswami S. Model-informed precision dosing at the bedside: Scientific challenges and opportunities. CPT Pharmacometrics Syst Pharmacol. 2018;7(12):785-7. doi:10.1002/psp4.12353
22. Wang Y, Fu H, Zeng D. Learning optimal personalized treatment rules in consideration of benefit and risk: With an application to treating type 2 diabetes patients with insulin therapies. J Am Stat Assoc. 2018;113(521):1-13. doi:10.1080/01621459.2017.1303386
23. Tervonen T, Prawitz T, Chua GN, Hyacinthe J, Pinto CA. Net clinical benefit of antiplatelet therapy was affected by patient preferences: A personalized benefit-risk assessment. J Clin Epidemiol. 2022;144:84-92. doi:10.1016/j.jclinepi.2021.11.036
24. Yeung WY, Reigner B, Beyer U, Diack C, Sabanés Bové D, Palermo G, et al. Bayesian adaptive dose-escalation designs for simultaneously estimating the optimal and maximum safe dose based on safety and efficacy. Pharm Stat. 2017;16(6):396-413. doi:10.1002/pst.1818
25. O'Donnell PH, Wadhwa N, Danahey K, Borden BA, Lee SM, Hall JP, et al. Pharmacogenomics-based point-of-care clinical decision support significantly alters drug prescribing. Clin Pharmacol Ther. 2017;102(5):859-69. doi:10.1002/cpt.709
26. Tremmel R, Schreeck F, Jaeger S, Schricker S, Schaeffeler E, Igel S, et al. Clinical implementation of pharmacogenomics and drug–drug interaction screening in a German academic teaching hospital and outpatient follow-up. Clin Pharmacol Ther. 2026;119(1):241-54. doi:10.1002/cpt.70083
27. Berezowska M, Hayden IS, Brandon AM, Zats A, Patel M, Barnett S, et al. Recommended approaches for integration of population pharmacokinetic modelling with precision dosing in clinical practice. Br J Clin Pharmacol. 2025;91(4):1064-79. doi:10.1111/bcp.16335
28. Kuemmel C, Yang Y, Zhang X, Florian J, Zhu H, Tegenge M, et al. Consideration of a credibility assessment framework in model-informed drug development: Potential application to physiologically-based pharmacokinetic modeling and simulation. CPT Pharmacometrics Syst Pharmacol. 2020;9(1):21-8. doi:10.1002/psp4.12479
29. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385. doi:10.1136/bmj-2023-078378
30. 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
31. Chang WC, Tanoshima R, Ross CJD, Carleton BC. Challenges and opportunities in implementing pharmacogenetic testing in clinical settings. Annu Rev Pharmacol Toxicol. 2021;61:65-84. doi:10.1146/annurev-pharmtox-030920-025745