Available online www.ijpras.com

International Journal of Pharmaceutical Research & Allied Sciences, 2026, 15(2):50-60
https://doi.org/10.51847/baB3ZPlMrw

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Counterfactual Dose Twins for Real-Time Precision Pharmacotherapy under Treatment Adaptation, Delayed Outcomes, Safety Constraints, and Incomplete Clinical Observation

Ravi Kumar[1]*, Neha Sharma[1], Aniket Deshmukh[2], Sanjay Kulkarni[3], Arjun Nair[4]

[1]Department of Counterfactual Dose Twins and Real-Time Pharmacotherapy, Faculty of Pharmacy, Indian Agricultural Research Institute, New Delhi, India.
[2]Department of Treatment Adaptation and Delayed Outcomes, Faculty of Pharmacy, IIT Bombay, Mumbai, India.
[3]Department of Safety Constraints and Clinical Observation, Faculty of Pharmaceutical Sciences, Savitribai Phule Pune University, Pune, India.
[4]Department of Precision Pharmacotherapy and Modeling, Faculty of Pharmacy, IIT Kharagpur, Kharagpur, India.

***Email:** ravi.kumar@iari.res.in*

ABSTRACT

Precision pharmacotherapy requires decisions about what dose should be administered next, not merely forecasts of what is likely to occur under previously observed treatment patterns. Current patient models can integrate pharmacokinetic, pharmacodynamic, laboratory, physiological, and treatment data, yet their predictions may remain associational when treatment is repeatedly adapted in response to evolving patient states. Delayed efficacy and toxicity, incomplete clinical observation, treatment-confounder feedback, unsupported dose actions, and model uncertainty further limit direct translation from patient-state prediction to dosing decisions. This Original Causal Digital-Twin Architecture Article proposes the Counterfactual Dose Twin as a patient-specific, time-indexed computational branch that shares the observed clinical history but propagates outcomes under one explicitly defined future dose trajectory. A dose-twin ensemble compares permissible trajectories through linked observation, state-estimation, causal-specification, PK/PD, delayed-outcome, uncertainty, and governance layers. Hard safety constraints, action-support checks, calibrated uncertainty, abstention, and escalation are positioned before clinician authorization rather than appended to an optimized dose output. Real-time updating is conceived as a governed process in which new observations revise the estimated patient state while model versions, data provenance, overrides, missingness, and emerging outcomes remain auditable. Validation would require software verification, causal-recovery experiments, temporal and missing-data stress tests, external retrospective evaluation, prospective silent-mode assessment, human-factors testing, and context-specific interventional evidence. The proposed architecture is an original conceptual synthesis rather than a validated dosing model, treatment recommendation, autonomous prescribing system, regulatory determination, or deployment-ready clinical workflow.

Key words: *Clinical pharmacokinetics, Pharmacodynamics, PK/PD, Dose–exposure–response, Population modelling, Precision dosing*

Received: 16 February 2026; **Revised:** 30 April 2026; **Accepted:** 06 May 2026

INTRODUCTION

Model-informed precision dosing seeks to connect drug, patient, disease, and treatment information with individualized dose selection. Although its scientific foundations are increasingly established, routine clinical

integration remains limited by evidence, workflow, qualification, and governance barriers rather than by pharmacometric capability alone [1].

Effective implementation requires more than an accurate population model or dosing algorithm. Interoperable data, qualified models, usable software, clinical evidence, trained professionals, institutional accountability, and integration with treatment workflows must operate together if model output is to become decision-relevant [2].

Contemporary model-informed dosing commonly combines population PK or PK/PD models, patient covariates, measured concentrations, prior information, Bayesian updating, and dose optimization. However, prospective evidence, transportability, scalable monitoring, and consistent implementation remain incomplete across many therapeutic contexts [3].

A further limitation is conceptual. Most dosing systems estimate a current state or predict an outcome under observed care, whereas clinicians ask what would occur under alternative future dose trajectories. This article proposes a causal digital-twin architecture that separates state estimation from intervention comparison and embeds delayed outcomes, incomplete observation, uncertainty, safety constraints, abstention, and clinical governance within the decision process.

Why descriptive patient models cannot answer dosing counterfactuals

A descriptive model may accurately predict an outcome among patients receiving historically observed treatments without identifying the outcome that would follow if the dose were changed. Prediction under treatment and estimation of an intervention-specific counterfactual are therefore different tasks requiring different assumptions [4, 5].

A dosing counterfactual must specify eligibility, decision time, permissible interventions, treatment assignment, follow-up, outcome, and estimand. Target-trial reasoning helps expose temporal misalignment, immortal-time bias, inconsistent treatment definitions, and inappropriate follow-up, but it does not remove residual confounding or repair unmeasured clinical history [6].

Treatment adaptation intensifies this distinction. Renal function, biomarkers, symptoms, prior concentrations, adverse effects, and disease severity may influence the next dose while also having been affected by earlier treatment. Conditioning on these variables without representing their temporal roles may distort rather than clarify the effect of changing dose.

Incomplete observation creates an additional problem. The measured record is not the patient's complete biological state: concentrations may be sparsely sampled, symptoms may be recorded selectively, and safety outcomes may remain pending. A descriptive forecast can therefore appear precise while depending on untested assumptions about missingness, measurement error, and latent state evolution.

The proposed architecture addresses these limitations by separating observations, latent patient state, treatment history, intervention specification, counterfactual trajectory generation, delayed outcomes, uncertainty, and decision constraints. Predictive calibration remains necessary, but it does not establish the causal effect, clinical acceptability, or safety of changing treatment. **Figure 1** presents the original conceptual synthesis for counterfactual dose-twin architecture, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

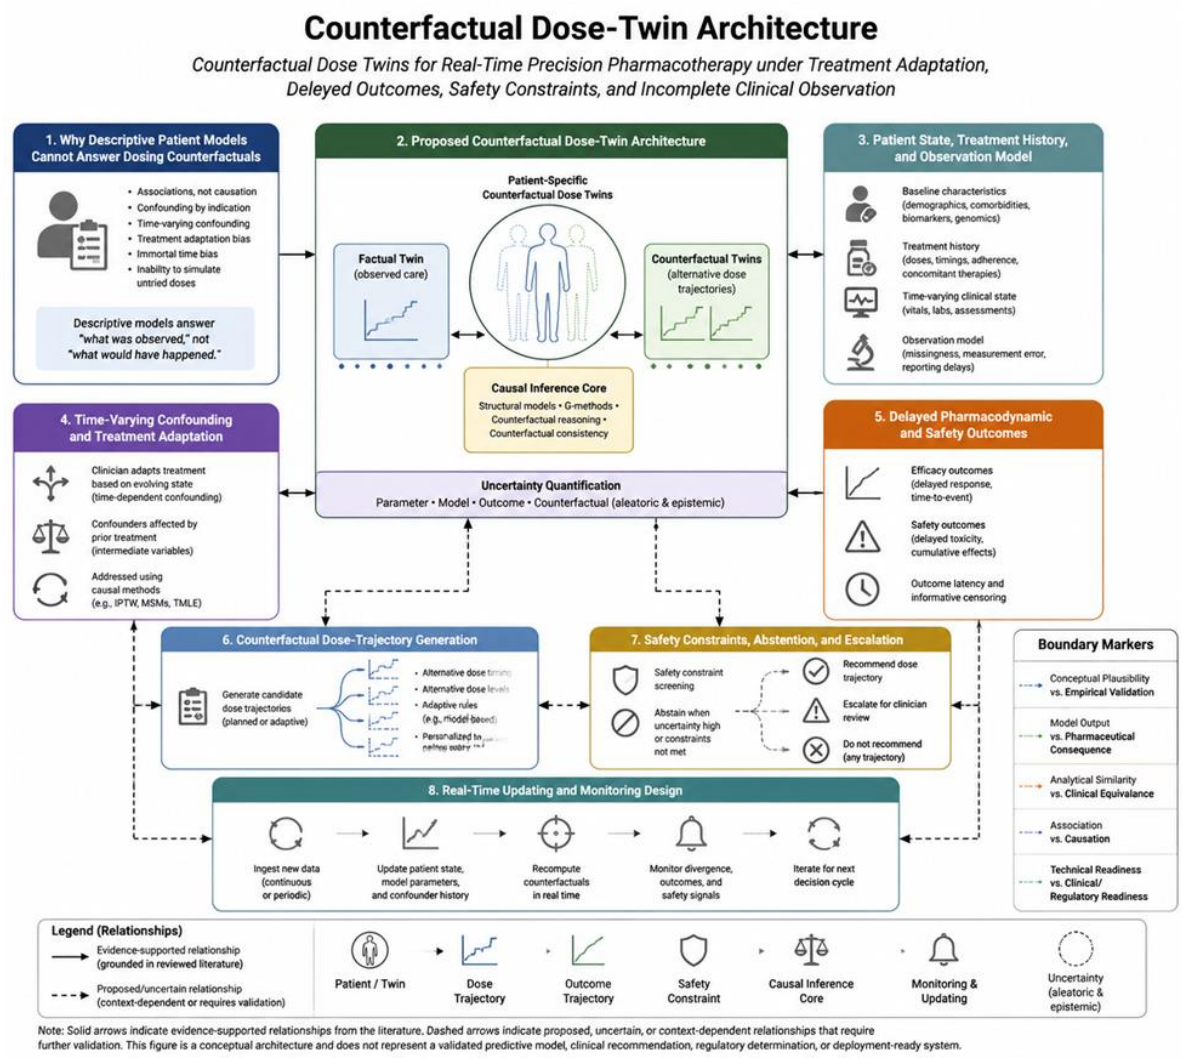

Figure 1. Counterfactual Dose-Twin Architecture

The figure is an original conceptual synthesis developed for "Counterfactual Dose Twins for Real-Time Precision Pharmacotherapy under Treatment Adaptation, Delayed Outcomes, Safety Constraints, and Incomplete Clinical Observation." 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.

Proposed counterfactual dose-twin architecture

Digital-twin scholarship describes individualized computational representations that link patient data with mechanistic or statistical models and can be updated as the represented system changes [7, 8]. The architecture proposed here adds an explicit causal requirement: every simulated future must correspond to a defined dose intervention rather than an unspecified continuation of observed care.

A Counterfactual Dose Twin is defined as a patient-specific, time-indexed computational branch that shares the patient's observed history at a decision time and propagates PK, PD, disease, safety, and observation uncertainty under one permissible future dose trajectory. A set of branches evaluated from the same history forms the proposed dose-twin ensemble.

A pharmacology-specific twin may contain mechanistic PK/PD or quantitative systems pharmacology components, provided their structure, calibration data, and context of use are explicit. Existing disease-specific QSP twin applications demonstrate the feasibility of individualized mechanistic simulation but do not validate a general counterfactual dosing architecture [9].

The proposed architecture contains six connected layers: observation and measurement; patient state and treatment history; causal intervention specification; counterfactual trajectory generation; delayed-outcome and uncertainty

propagation; and safety and clinical governance. The layers separate what was measured, what is inferred, what intervention is being compared, and what decision may be authorized.

Evidence-supported pharmacological relationships should remain distinct from newly proposed decision rules. Mechanistic plausibility does not establish causal identification, a simulated trajectory does not establish a treatment effect, and an optimized model objective does not establish clinical acceptability. Each component therefore requires independent verification, validation, and context-specific qualification.

Patient state, treatment history, and observation model

The patient state should distinguish measured information from latent pharmacological and clinical states. Bedside model-informed dosing depends on qualified priors, interpretable historical data, model qualification, and clinically usable individual updating rather than on a population parameter set alone [10].

The observed history includes administered doses, infusion timing, concentrations, laboratory results, physiological measurements, symptoms, concomitant treatments, outcomes, and their timestamps. Therapeutic drug monitoring becomes model-informed dosing when these observations are integrated with population PK/PD structure, exposure targets, covariates, and Bayesian forecasting [11].

The proposed latent state contains PK, PD, disease, and safety components that may not be directly observed. Posterior updating can refine model-based beliefs about these components, but posterior precision is not equivalent to biological truth, causal certainty, or evidence that a new dose will improve outcomes.

A separate treatment–covariate history ledger preserves the temporal order of prior doses, treatment changes, evolving covariates, observations, and outcomes. This shared past anchors every counterfactual branch and prevents future trajectories from being compared under inconsistent histories or selectively reconstructed baselines.

The observation model represents timestamps, provenance, assay quality, measurement error, sampling cadence, missingness, censoring, privacy, security, and workflow constraints. Software evaluations show that integration, validation, output communication, usability, and data governance are material requirements, although feature availability alone does not demonstrate clinical effectiveness [12]. The observation layer therefore defines what the architecture can condition on, not what it can safely infer beyond the recorded information.

Time-varying confounding and treatment adaptation

Repeated dose adaptation creates treatment–confounder feedback when an evolving covariate influences the next dose but has itself been altered by earlier treatment. Conventional adjustment may then block mediated effects or introduce selection bias; longitudinal g-methods are relevant only when their exchangeability, positivity, consistency, and model assumptions are defensible [13].

The causal target should be a sustained or dynamic dosing strategy rather than a decontextualized exposure label. Its population, eligibility criteria, decision times, adherence definition, censoring rules, follow-up horizon, and outcome must be specified because a per-protocol effect is conditional on the strategy actually being evaluated [14].

Naive conditioning on treatment-affected renal function, biomarkers, symptoms, or toxicity can produce biased comparisons between dose trajectories. The causal layer must therefore distinguish confounders, mediators, colliders, and observation variables according to their temporal roles rather than their names in the clinical record [15].

The proposed twin records these assumptions and identifies candidate estimators but cannot create exchangeability, recover unmeasured clinical reasoning, or overcome unsupported treatment actions. Counterfactual interpretation remains conditional on longitudinal data quality, intervention consistency, adequate action support, and sensitivity analyses for plausible violations.

Delayed pharmacodynamic and safety outcomes

Pharmacodynamic response and clinical events may evolve on different timescales. Joint longitudinal and time-to-event models can connect changing biomarkers or responses with delayed events, although the fitted association structure does not by itself establish a causal biological mechanism [16].

Pending outcomes must remain visible during dose comparison. Methodology for delayed-toxicity dose finding demonstrates why partial follow-up should contribute time-dependent information rather than being classified prematurely as absence of toxicity, although this principle requires separate validation in routine pharmacotherapy [17].

Cumulative exposure, treatment duration, and evolving patient factors may contribute to delayed safety risk. Exposure–time-to-event modelling provides an applicable structural example, but disease- and product-specific findings cannot define a universal exposure–toxicity relationship [18].

The proposed architecture therefore carries unresolved efficacy and safety states across decision times, distinguishes observed absence from unobserved or pending outcome, and compares trajectories over a common horizon. **Table 1** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in counterfactual dose twins for real-time precision pharmacotherapy under treatment.

Table 1. Definitions, components, relationships, and intended uses in Counterfactual Dose Twins for Real-Time Precision Pharmacotherapy under Treatment.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Observed clinical state	Records are incomplete and time-dependent	Doses, concentrations, laboratories, physiology, symptoms, timestamps	Defines the information available at each decision time [10]	Reconstructs what was known and when	Provenance, timestamp, and measurement-quality verification	Missing or selectively measured variables	Observed state is not the complete biological state
Latent PK state	Individual disposition is partly unobserved	Population prior, covariates, dose and concentration history	Bayesian state estimation propagates exposure under candidate doses [11]	Patient-specific exposure forecasts	External calibration and concentration-prediction checks	Structural misspecification or weak sampling	Posterior precision is not biological or causal certainty
Latent PD and safety state	Benefit and harm may emerge after exposure	Biomarkers, outcomes, disease state, cumulative exposure	Links longitudinal response with pending events [16]	Represents delayed benefit and risk	Longitudinal-event validation	Surrogate failure or incorrect lag structure	A fitted association is not a validated clinical endpoint
Treatment-history ledger	Adaptation changes later covariates and decisions	Ordered treatment, covariate, monitoring, and outcome history	Provides the common past for every branch [13]	Preserves temporal ordering	Complete administration and monitoring records	Unmeasured clinical reasoning	The ledger does not remove confounding
Observation model	Missingness may be informative	Sampling process, assay error, censoring, quality flags	Models how observations arise rather than treating them as complete state [12]	Exposes measurement uncertainty	Missing-data and sampling-cadence stress tests	Non-identifiable observation processes	Unmeasured information cannot be recovered without assumptions
Causal intervention specification	Forecasts do not automatically answer treatment questions	Decision time, eligible actions, horizon, outcomes, estimand	Converts prediction into an explicit dose-strategy comparison [6]	Aligns branches with a defined counterfactual question	Target-trial specification and causal diagnostics	Positivity or exchangeability violations	Specification does not guarantee identification
Counterfactual Dose Twin branch	One future dose path must be distinguished from alternatives	Shared history and one permissible future trajectory	Propagates state and outcomes under one proposed	Creates a comparable patient-specific future	Mechanistic, causal, and external validation	Simulator error or unsupported action	Proposed construct; not a biological clone or treatment effect

			intervention [9]				
Dose-twin ensemble	Alternatives require a common comparison basis	Multiple branches with the same past and horizon	Contrasts permissible dose trajectories	Supports constrained comparison	Candidate-set completeness and policy-support analysis	Relevant actions omitted	Proposed construct; not an exhaustive clinical choice set
Uncertainty envelope	Point forecasts conceal decision risk	Parameter, structural, measurement, and distributional uncertainty	Accompanies every trajectory and can activate review [19]	Makes uncertainty decision-relevant	Calibration and distribution-shift testing	Confident error or miscalibration	Uncertainty estimation is not proof of safety
Constraint, abstention, and escalation controller	Unsupported outputs may reach care	Clinical constraints, action support, uncertainty, governance rules	Rejects, withholds, or escalates candidate actions [20]	Limits automated overreach	Software verification and human-factors evaluation	Constraint bypass or automation bias	Proposed controller; not an autonomous prescriber

Counterfactual dose-trajectory generation

Sequential decision methods are relevant because each dose changes the future state on which subsequent decisions depend. Reinforcement learning can formalize states, actions, transitions, and objectives, but pharmacological structure and clinical constraints are required to limit data demands and unsafe exploration [21]. A PK/PD or disease model can provide a simulated environment for comparing sequential actions and testing policies before clinical evaluation. Model-enhanced reinforcement learning demonstrates this computational possibility, but success within a theoretical simulator remains conditional on simulator fidelity and does not establish clinical effectiveness [22].

Adaptive dosing also requires an explicit objective. A policy may balance exposure targets, response, toxicity, treatment burden, and uncertainty, but the relative value assigned to these outcomes can materially alter the preferred trajectory [23].

The proposed trajectory generator therefore enumerates only permissible actions, propagates each branch across a common horizon, and records objective components separately. A learned policy becomes counterfactual only when its intervention, estimand, causal assumptions, action support, and uncertainty are explicit.

Safety constraints, abstention, and escalation

Candidate trajectories require hard clinical constraints, action-support checks, conservative off-policy evaluation, and safeguards against reward misspecification. Healthcare reinforcement learning remains vulnerable to confounding, limited coverage of alternative actions, unsafe exploration, and evaluation against an unreliable observational policy [20].

Decision-relevant uncertainty must be quantified and communicated in a form that can trigger second review rather than being hidden behind a point recommendation [19, 24]. Aleatoric, epistemic, structural, measurement, and distributional uncertainty should remain distinguishable because they imply different corrective actions.

The proposed controller returns one of three states: eligible for clinician comparison, abstain or withhold, or escalate for additional measurement or specialist review. Constraint violation, inadequate action support, unstable state estimation, unresolved delayed toxicity, or excessive uncertainty prevents an affirmative recommendation.

Abstention reduces overreach but cannot remove residual clinical risk or guarantee that escalation will be timely. **Table 2** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for counterfactual dose twins for real-time precision pharmacotherapy under treatment.

Table 2. Testable propositions, evidence requirements, and validation criteria for Counterfactual Dose Twins for Real-Time Precision Pharmacotherapy under Treatment.

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

Causal specification	Define dose-strategy estimand	History and candidate actions enter a longitudinal causal model	Constrains trajectory interpretation	Recovery of known effects in simulated data [13]	Bias under feedback or assumption violation	Design causal-recovery and sensitivity experiments	Simulation recovery does not establish clinical validity
Observation and state layer	Estimate current PK, PD, disease, and safety state	Measurements update latent state	Supplies the shared branch origin	Calibration under altered sampling and missingness [12]	Unstable estimates or informative missingness	Conduct external and missing-data stress testing	Calibration is context-specific
Delayed-outcome layer	Preserve pending benefit and harm	Longitudinal responses inform delayed-event projections	Feeds safety and comparison gates	Time-dependent calibration with partial follow-up [16]	Premature reassurance	Validate landmarked longitudinal-event predictions	Association does not establish causation
Trajectory generator	Create permissible alternative dose paths	State and actions enter mechanistic or simulated propagation	Produces branches for comparison	Reproduction of expected dynamics in controlled simulations [20]	Simulator exploitation or unsupported extrapolation	Verify equations, software, and candidate-action support	Simulation performance is not clinical benefit
Objective and comparison layer	Compare competing outcomes	Branch outcomes enter an explicit objective	Operates only after constraint screening	Stability across plausible objective specifications [21]	Reward misspecification	Prespecify outcomes with clinical and patient input	No universal utility function is claimed
Safety gate	Reject prohibited or unsupported branches	Constraints and support diagnostics filter trajectories	Precedes clinician review	No hard-constraint bypass in adversarial tests [22]	Unsafe action reaches comparison	Test fail-closed behaviour independently	Passing software tests does not prove treatment safety
Uncertainty and abstention controller	Route uncertain decisions away from automation	Uncertainty status produces compare, abstain, or escalate state	Interacts with observation and governance layers	Prespecified calibration and selective-risk performance [23]	Confident error or excessive abstention	Define escalation pathways and review thresholds prospectively	Thresholds require context-specific validation
Governed updating and human use	Preserve traceability across repeated decisions	New observations, actions, overrides, and outcomes enter an audit trail	Reopens state estimation without silently changing authorization	Version stability, auditability, and human-factors acceptability	Drift, automation bias, or alert fatigue	Evaluate prospectively in silent and supervised modes	Proposed process; later clinical evidence remains necessary

Real-time updating and monitoring design

Continued learning can revise population and individual model components as clinical data accumulate, but updates require diagnostic evaluation, version control, and protection against degradation [25]. The proposed architecture separates patient-state updating from changes to the governing model.

Real-time bedside support also requires automated data acquisition, timestamp reconciliation, model execution, interpretable output, and workflow integration. Existing dosing-support development demonstrates these feasibility components without establishing general clinical effectiveness or validating the proposed architecture [26].

Each decision cycle should preserve the input snapshot, model version, candidate trajectories, rejected branches, uncertainty state, clinician action, override rationale, administered treatment, and subsequent observations. Monitoring should detect calibration drift, altered case mix, new assay behaviour, missing-data changes, and unsupported actions.

A new observation may update the patient posterior immediately when the validated model is unchanged. By contrast, parameter, structure, objective, constraint, or interface changes should enter controlled shadow evaluation and independent review before affecting recommendations.

Validation and clinical governance

Clinical evaluation must examine intended use, eligible users, human–system interaction, workflow fit, errors, overrides, safety events, and prospective performance in the actual decision environment [27]. Reporting guidance supports transparent evaluation but does not itself validate a system.

Computational credibility requires distinct verification, validation, uncertainty quantification, and context-of-use assessment [28]. Equation correctness, software correctness, biological adequacy, causal recovery, transportability, and clinical usefulness are separate claims and should not be collapsed into one accuracy statistic.

A staged programme would proceed through software verification, simulated causal-recovery tests, retrospective temporal and missingness stress tests, multisite external evaluation, prospective silent-mode assessment, human-factors studies, and controlled interventional evaluation. Failure at one stage should restrict the intended use rather than be averaged away.

Clinical governance should define model ownership, authorized users, change control, audit review, incident response, override handling, monitoring frequency, and retirement criteria. Passing a computational stage does not authorize prescribing, eliminate clinician responsibility, or establish regulatory acceptance.

Limitations and implementation priorities

Claims of patient benefit require transparent problem formulation, representative data, reproducibility, prospective evaluation, workflow analysis, and ethical accountability [29]. These requirements remain context-specific and cannot be satisfied solely through technical documentation or retrospective performance.

Post hoc explanation methods may support review but cannot substitute for causal validity, calibration, robustness, safety evaluation, or accountability [30]. An apparently intuitive explanation may still describe an unreliable association or obscure dependence on unsupported data.

Figure 2 presents the original conceptual synthesis for real-time updating and constrained comparison of dose trajectories, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

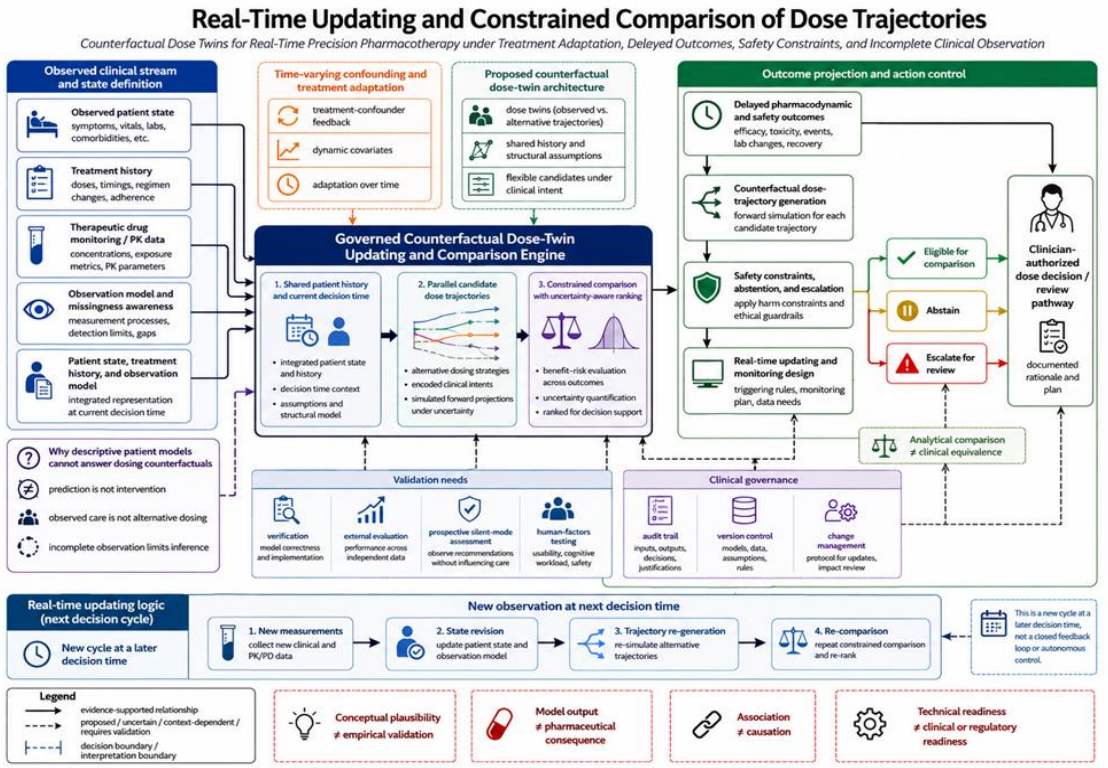

Figure 2. Real-Time Updating and Constrained Comparison of Dose Trajectories

The figure is an original conceptual synthesis developed for "Counterfactual Dose Twins for Real-Time Precision Pharmacotherapy under Treatment Adaptation, Delayed Outcomes, Safety Constraints, and Incomplete Clinical Observation." 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.

Principal limitations include unmeasured treatment rationale, imperfect adherence records, sparse sampling, uncertain structural transportability, delayed or competing outcomes, changing clinical practice, and objective misspecification. Priorities are therefore prospective data provenance, causal benchmark environments, multisite validation, calibrated abstention, human-factors evaluation, and governance mechanisms that fail closed when evidence is insufficient.

CONCLUSION

The Counterfactual Dose Twin is proposed as a causal, pharmacological, uncertainty-aware architecture for comparing permissible future dose trajectories from a shared patient history. Its contribution is the separation of observation, latent state, treatment adaptation, causal intervention specification, delayed outcomes, trajectory generation, safety constraints, abstention, updating, and clinical authorization. The architecture does not establish that a simulated trajectory is a causal treatment effect, that an optimized objective is clinically acceptable, or that a recommendation is safe. Its value must be tested through staged computational, causal, external, prospective, human-factors, and interventional validation within narrowly defined contexts of use.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. 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
2. Polasek TM, Rostami-Hodjegan A, Yim DS, Jamei M, Lee H, Kimko H, et al. What does it take to make model-informed precision dosing common practice? Report from the 1st Asian Symposium on Precision Dosing. AAPS J. 2019;21(2):17. doi:10.1208/s12248-018-0286-6
3. Minichmayr IK, Dreesen E, Centanni M, Wang Z, Hoffert Y, Friberg LE, et al. Model-informed precision dosing: State of the art and future perspectives. Adv Drug Deliv Rev. 2024;215:115421. doi:10.1016/j.addr.2024.115421
4. van Geloven N, Swanson SA, Ramspek CL, Luijken K, van Diepen M, Morris TP, et al. Prediction meets causal inference: The role of treatment in clinical prediction models. Eur J Epidemiol. 2020;35(7):619-30. doi:10.1007/s10654-020-00636-1
5. 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
6. Dickerman BA, García-Albéniz X, Logan RW, Denaxas S, Hernán MA. Avoidable flaws in observational analyses: An application to statins and cancer. Nat Med. 2019;25(10):1601-6. doi:10.1038/s41591-019-0597-x
7. 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
8. 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

9. Kaddi C, Tao M, Bergeler S, George K, Geerts H, van der Graaf PH, et al. Quantitative systems pharmacology-based digital twins approach supplements clinical trial data for enzyme replacement therapies in Pompe disease. Clin Pharmacol Ther. 2025;117(2):579-88. doi:10.1002/cpt.3498
10. 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
11. 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
12. Kantasiripitak W, Van Daele R, Gijsen M, Ferrante M, Spriet I, Dreesen E. Software tools for model-informed precision dosing: How well do they satisfy the needs? Front Pharmacol. 2020;11:620. doi:10.3389/fphar.2020.00620
13. Naimi AI, Cole SR, Kennedy EH. An introduction to g methods. Int J Epidemiol. 2017;46(2):756-62. doi:10.1093/ije/dyw323
14. Hernán MA, Robins JM. Per-protocol analyses of pragmatic trials. N Engl J Med. 2017;377(14):1391-8. doi:10.1056/NEJMsm1605385
15. Mansournia MA, Etminan M, Danaei G, Kaufman JS, Collins GS. Handling time varying confounding in observational research. BMJ. 2017;359. doi:10.1136/bmj.j4587
16. Zhudenkov K, Gavrilov S, Sofronova A, Stepanov O, Kudryashova N, Helmlinger G, et al. A workflow for the joint modeling of longitudinal and event data in the development of therapeutics: Tools, statistical methods, and diagnostics. CPT Pharmacometrics Syst Pharmacol. 2022;11(4):425-37. doi:10.1002/psp4.12763
17. Lin R, Yuan Y. Time-to-event model-assisted designs for dose-finding trials with delayed toxicity. Biostatistics. 2020;21(4):807-24. doi:10.1093/biostatistics/kxz007
18. Kågedal M, Samineni D, Gillespie WR, Lu D, Fine BM, Girish S, et al. Time-to-event modeling of peripheral neuropathy: Platform analysis of eight valine-citrulline-monomethylauristatin E antibody-drug conjugates. CPT Pharmacometrics Syst Pharmacol. 2019;8(8):606-15. doi:10.1002/psp4.12442
19. Begoli E, Bhattacharya T, Kusnezov D. The need for uncertainty quantification in machine-assisted medical decision making. Nat Mach Intell. 2019;1(1):20-3. doi:10.1038/s42256-018-0004-1
20. Gottesman O, Johansson F, Komorowski M, Faisal A, Sontag D, Doshi-Velez F, et al. Guidelines for reinforcement learning in healthcare. Nat Med. 2019;25(1):16-8. doi:10.1038/s41591-018-0310-5
21. Ribba B, Dudal S, Lavé T, Peck RW. Model-informed artificial intelligence: Reinforcement learning for precision dosing. Clin Pharmacol Ther. 2020;107(4):853-7. doi:10.1002/cpt.1777
22. Ribba B, Bräm DS, Baverel PG, Peck RW. Model enhanced reinforcement learning to enable precision dosing: A theoretical case study with dosing of propofol. CPT Pharmacometrics Syst Pharmacol. 2022;11(11):1497-510. doi:10.1002/psp4.12858
23. Tosca EM, De Carlo A, Ronchi D, Magni P. Model-informed reinforcement learning for enabling precision dosing via adaptive dosing. Clin Pharmacol Ther. 2024;116(3):619-36. doi:10.1002/cpt.3356
24. Kompa B, Snoek J, Beam AL. Second opinion needed: Communicating uncertainty in medical machine learning. NPJ Digit Med. 2021;4(1):4. doi:10.1038/s41746-020-00367-3
25. Maier C, de Wiljes J, Hartung N, Kloft C, Huisinga W. A continued learning approach for model-informed precision dosing: Updating models in clinical practice. CPT Pharmacometrics Syst Pharmacol. 2022;11(2):185-98. doi:10.1002/psp4.12745
26. Roggeveen LF, Guo T, Driessen RH, Fleuren LM, Thoral P, van der Voort PHJ, et al. Right dose, right now: Development of AutoKinetics for real time model informed precision antibiotic dosing decision support at the bedside of critically ill patients. Front Pharmacol. 2020;11:646. doi:10.3389/fphar.2020.00646
27. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
28. Viceconti M, Pappalardo F, Rodriguez B, Horner M, Bischoff J, Musuamba Tshinanu F. In silico trials: Verification, validation and uncertainty quantification of predictive models used in the regulatory evaluation of biomedical products. Methods. 2021;185:120-7. doi:10.1016/j.ymeth.2020.01.011

29. Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368. doi:10.1136/bmj.l6927
30. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9