Original Article

# The Pharmacotherapy Digital Twin as a Negotiated Model of the Patient rather than a Computational Replica

**Patrick O'Connor[1]*, Grace Murphy[1], Niamh Walsh[2]**

[1]Department of Digital Twins and Pharmacotherapy, School of Pharmacy, University College Cork, Cork, Ireland. [2]Department of Patient Modeling and Clinical Negotiation, Faculty of Pharmacy, University of Galway, Galway, Ireland.

## Abstract

Digital-twin terminology increasingly frames the patient as a computational entity whose changing state can be reproduced, predicted, and tested under simulated interventions. In pharmacotherapy, however, no model can fully reproduce the biological, experiential, behavioral, social, and organizational conditions through which medication effects emerge. Treating a pharmacotherapy digital twin as a replica may consequently conceal model selection, incomplete data, disagreement, uncertainty, and the distribution of decision authority. This article proposes Negotiated Pharmacotherapy Digital-Twin Theory, an original non-empirical account in which the twin is understood as a partial, task-specific, and revisable model assembled through structured interaction among clinical evidence, computational inference, and patient-reported knowledge. The proposed architecture comprises three inspectable model layers, a negotiation workspace, a disagreement register, an intended-use contract, a permission and human-control layer, provenance-preserving updates, and conditional representations of alternative futures. Negotiation does not require consensus or imply that every representation is equally credible. Instead, it makes assumptions, conflicts, uncertainties, and authority boundaries visible before a model is accepted for a specified medication decision. Evaluation would require technical verification, external and temporal validation, calibration and uncertainty assessment, subgroup analysis, human-factors testing, workflow evaluation, patient-concordance assessment, clinical-utility studies, and postdeployment monitoring. The theory does not establish predictive superiority, safer prescribing, improved outcomes, regulatory acceptability, or deployment readiness. Its original contribution is to replace the ontologically strong replica metaphor with a governed model of representation in which the patient remains distinct from—and able to contest—the digital account constructed for pharmacotherapy.

**Keywords:** Digital pharmacy, Artificial intelligence, Pharmacy practice, Medication safety, Human–AI collaboration, Clinical decision support

## Introduction

Digital twins have been presented as computational counterparts through which individual disease trajectories and responses to alternative interventions might be simulated. Within personalized medicine, this ambition has included the construction of computational copies that integrate patient-specific information with mechanistic or data-driven models [1]. The attraction of the replica metaphor is understandable: it suggests continuity between the physical patient and a sufficiently detailed digital representation capable of supporting prospective pharmacotherapy decisions.

Health digital twins are nevertheless defined through differing combinations of longitudinal data, population knowledge, mechanistic models, machine learning, simulation, and bidirectional updating. Their clinical meaning therefore depends not only on computational sophistication but also on implementation conditions, regulatory interpretation, data availability, and the decision

**Address for correspondence:** Patrick O'Connor, Department of Digital Twins and Pharmacotherapy, School of Pharmacy, University College Cork, Cork, Ireland.
patrick.oconnor@ucc.ie
**Received:** 03 November 2025; **Accepted:** 01 February 2026



**How to cite this article:** O'Connor P, Murphy G, Walsh N. The Pharmacotherapy Digital Twin as a Negotiated Model of the Patient rather than a Computational Replica. Arch Pharm Pract. 2026;17(1):31-8. https://doi.org/10.51847/9R2QJ1nnzc

for which the representation is used [2]. A digital object that supports dose estimation, for example, is not necessarily equivalent to one intended to forecast toxicity, compare treatment strategies, or coordinate patient–professional deliberation.

The developing literature confirms considerable heterogeneity in digital-twin definitions, architectures, medical applications, and validation strategies [3]. Reviews of systems explicitly described as patient digital twins similarly show inconsistent terminology and a field that remains predominantly developmental or preclinical [4]. The term can consequently imply a degree of correspondence, continuity, or clinical readiness that the underlying model has not established.

The problem is particularly important in pharmacotherapy. Medication response emerges from interacting biological scales, changing organ function, comorbidities, co-medications, adherence behavior, treatment burden, patient priorities, care access, and professional action. Complex-systems accounts of precision medicine therefore support multiscale and scenario-based modeling rather than a single reductive representation [5]. Even a technically advanced model remains a selective account produced for a purpose.

This article proposes that a pharmacotherapy digital twin should be understood not as a computational replica but as a negotiated model of the patient. "Negotiated" does not mean that biological facts are determined by opinion or that all perspectives carry equal evidential weight. It denotes a governed process through which heterogeneous representations are compared, challenged, qualified, and provisionally accepted for a bounded decision. The patient remains ontologically and morally distinct from the model. The twin is instead a revisable coordination object whose authority depends on its intended use, evidence, uncertainty, provenance, and relationship to human judgment.

## *Why Pharmacotherapy Models Are Partial and Contestable*

A pharmacotherapy model is partial because it selects which patient properties, treatment mechanisms, outcomes, and time horizons are representable. Selection is unavoidable: a dosing model may represent clearance and exposure while excluding treatment burden; a risk model may include recorded diagnoses but omit unstable housing, medication access, or an unreported adverse effect. Clinical machine learning can also produce unintended consequences when model outputs alter attention, workflow, or professional behavior in ways not captured by technical performance [6].

Partiality becomes clinically consequential when a measurable proxy is treated as equivalent to the underlying need. An algorithm may perform consistently against its chosen target while systematically misrepresenting patients for whom that target is unequally related to illness or care requirements [7]. The question is therefore not only whether a model predicts its target accurately, but whether the target, data, and permitted use provide a defensible representation of the medication decision.

A model is contestable when authorized participants can question its inputs, omissions, assumptions, uncertainty, interpretation, or proposed use. Contestability is not an objection added after prediction. It is a design property required because responsible clinical machine learning depends on safeguards extending beyond accuracy, including attention to data quality, unintended effects, monitoring, and accountability [8]. Clinical AI can produce unintended consequences, encode inequity through a misaligned proxy target, and require safeguards extending beyond predictive performance [6–8].

Contestability is also sociotechnical. The value and safety of a digital model depend on the people expected to interpret it, the organization in which it operates, the clinical task, the perceived value of the technology, and its capacity to adapt without losing control [9]. A pharmacotherapy twin that is technically valid in one setting may therefore remain unusable, misleading, or unsafe in another.

In this theory, partiality is not treated as a temporary deficiency that disappears when more data are collected. Additional measurements may reduce one uncertainty while introducing new temporal, semantic, or causal ambiguity. Contestability likewise does not imply that every output must be disputed. It means that the model cannot authenticate its own completeness or determine the legitimacy of its use.

**Figure 1** presents the original conceptual synthesis for negotiated pharmacotherapy digital-twin architecture, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

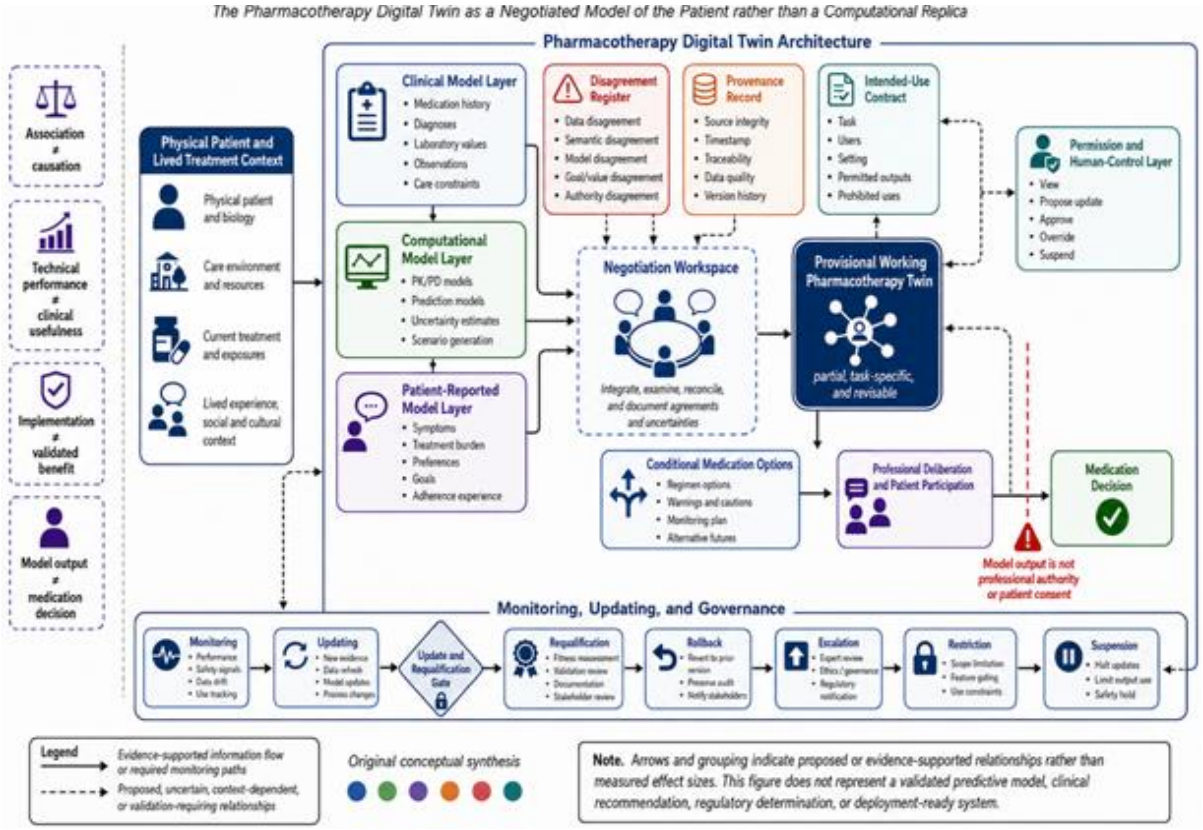


**Figure 1.** Negotiated Pharmacotherapy Digital-Twin Architecture

**Table 1** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article The Pharmacotherapy Digital Twin as a Negotiated Model of the Patient rather than a Computational Replica.

**Table 1.** Definitions, components, relationships, and intended uses for the article The Pharmacotherapy Digital Twin as a Negotiated Model of the Patient rather than a Computational Replica.

| Component or scientific dimension | Problem addressed | Inputs or determinants | Proposed mechanism or relationship | Expected contribution | Evidence required | Failure or uncertainty risk | Boundary statement |
|---|---|---|---|---|---|---|---|
| **Physical patient** | Replica language may collapse the person into the model. | Biology, experience, relationships, environment, treatment context. | The patient remains external to and only partially represented by every digital layer. | Preserves conceptual and moral distinction. | Evidence that represented variables are relevant to the intended task. | Reification of the model as the patient. | No model constitutes the complete patient [1–5]. |
| **Clinical layer** | Records contain incomplete and professionally interpreted information. | Diagnoses, medicines, laboratory findings, observations, care constraints. | Provides one inspectable representation rather than a complete ground truth. | Supports clinical interpretation and medication review. | Data-quality, timeliness, provenance, and task-relevance assessment. | Missing, stale, miscoded, or context-poor data. | Clinical documentation does not exhaust lived treatment reality. |
| **Computational layer** | Model output may be mistaken for direct knowledge of the patient. | Mechanistic models, statistical relationships, population priors, uncertainty estimates. | Converts selected inputs into estimates or conditional scenarios. | Supports structured inference unavailable from unaided review. | Verification, validation, calibration, transportability, and uncertainty evidence. | Proxy error, drift, overconfidence, or invalid extrapolation [6–8]. | Prediction is not causation, authority, or guaranteed benefit. |
| **Patient-reported layer** | Symptoms, burdens, goals, and medication-use | Patient reports, preferences, functioning, | Adds an independently inspectable | Makes patient-relevant omissions | Valid constructs, suitable timing, | Under-reporting, measurement burden, exclusion, | Patient reporting is indispensable but not infallible. |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | experience may be absent from records. | adherence experience, treatment burden. | experiential representation. | and conflicts visible. | accessibility, and interpretation. | or coercive collection. | |
| **Negotiation workspace (proposed)** | Heterogeneous representations may conflict without an explicit resolution process. | Layer outputs, assumptions, uncertainty, stakeholder questions. | Enables comparison, challenge, qualification, and provisional acceptance. | Converts hidden model conflict into visible deliberation. | Human-factors, comprehension, concordance, and decision-quality evaluation. | Tokenistic participation or forced consensus. | Negotiation does not make all claims equally credible. |
| **Disagreement register (proposed)** | Fusion may conceal unresolved differences. | Data, semantic, model, goal, value, authority, and temporal conflicts. | Records the type, consequence, owner, and disposition of disagreement. | Preserves uncertainty and accountability. | Reliability and utility of disagreement classification. | Alert burden or accumulation without resolution. | Some disagreements may remain unresolved without blocking every use. |
| **Working twin (proposed)** | "Twin" may imply a permanent and comprehensive replica. | Qualified layer content, accepted assumptions, unresolved warnings. | Creates a temporary representation for one intended use. | Supports bounded coordination. | Comparative evaluation against simpler models and usual decision processes. | Scope expansion or false permanence. | Acceptance for one task does not authorize another. |
| **Intended-use and permission structure (proposed)** | Model capability may be confused with permission to act. | Task, population, setting, users, outputs, exclusions, authority. | Constrains viewing, updating, approval, override, suspension, and action. | Clarifies responsibility and prevents unauthorized use. | Governance, workflow, auditability, and implementation evidence [8, 9]. | Informal workarounds or responsibility gaps. | Technical availability does not create clinical authority. |

## *Proposed Negotiated Digital-Twin Theory*

Negotiated Pharmacotherapy Digital-Twin Theory defines the twin as a provisional working representation created for a specified medication decision through governed comparison of partial models. The term "negotiated" draws on shared decision-making as a staged process in which options, preferences, and decisions are made discussable rather than reduced to a final consent event [10]. The theory extends that deliberative principle to the constitution of the model itself: participants may question what is represented, how it was inferred, which uncertainty matters, and what actions the representation may support.

The theory comprises seven connected constructs. First, layered representation preserves distinctions among clinical evidence, computational inference, and patient-reported knowledge. Second, declared incompleteness requires each layer to state material omissions and assumptions. Third, a negotiation workspace makes those layers jointly inspectable. Fourth, a disagreement register retains unresolved data, semantic, model, goal, value, authority, or temporal conflicts. Fifth, an intended-use contract specifies the medication task, population, setting, users, output, prohibited uses, and conditions requiring reassessment. Sixth, a permission layer defines who may view, modify, approve, override, suspend, or act. Seventh, a conditional-futures structure represents more than one plausible trajectory when the evidence does not justify a single determinate forecast.

Negotiation requires intelligibility, but intelligibility is role-dependent. Patients, pharmacists, prescribers, model developers, and governance bodies may need different explanations because they ask different questions and bear different consequences [11]. A patient may need to understand treatment burden and uncertainty; a pharmacist may need data provenance, dosing assumptions, and contraindication logic; a developer may need evidence of drift or model failure. The theory therefore rejects a single generic explanation as sufficient.

The proposed working twin is neither an autonomous agent nor a consensus document. Human and computational capabilities may complement one another, but computational capacity does not replace relational judgment, contextual interpretation, or professional responsibility [12]. Negotiation can end in acceptance, provisional acceptance with warnings, preservation of parallel representations, request for additional evidence, restriction of use, or refusal to generate an actionable output. Unresolved disagreement is not necessarily system failure; concealing a material disagreement is.

Multimodal artificial intelligence may expand the range of information that can be represented and reveal relationships not apparent within a single data source [13]. Yet additional modality does not dissolve the theory's boundaries. A more complex representation may remain poorly calibrated, causally ambiguous, inaccessible to patients, or incompatible with workflow. The twin's authority must therefore derive from evidence and governance appropriate to its intended use, not from computational scale or realism.

Several propositions follow. Layer-explicit systems should make source conflict more detectable than fully fused representations. Preserving unresolved disagreement should reduce false certainty. Intended-use contracts should reduce unauthorized scope expansion. Conditional future branches

should encourage more calibrated interpretation than single deterministic forecasts. Permissioned updating and provenance records should improve auditability. These propositions are testable but unvalidated. The theory would require comparison with simpler decision-support systems, ordinary medication review, and existing model-informed dosing approaches before any claim of added clinical value could be justified.

## Clinical, Computational, and Patient-Reported Model Layers

The clinical layer represents information ordinarily interpreted within pharmacotherapy practice: medication history, diagnoses, organ function, laboratory measurements, co-medications, allergies, prior responses, professional observations, treatment setting, and operational constraints. Model-informed precision dosing demonstrates how population knowledge can be combined with patient-specific characteristics, but it also shows that clinical translation is limited by incomplete data, model suitability, implementation barriers, and uncertainty about how outputs should enter practice [14]. The clinical layer is therefore not a ground-truth layer. It is a professionally mediated account whose content depends on documentation, timing, measurement, and context.

The computational layer contains mechanistic, pharmacometric, statistical, or machine-learning representations used to estimate latent states, exposure, response, risk, or possible regimen consequences. Therapeutic drug monitoring and model-informed dosing illustrate how observed concentrations and treatment information can update individualized estimates during therapy [15]. Such updating can be clinically informative without turning the model into a replica. Estimated clearance, exposure, or response remains conditional on model structure, measurement quality, population priors, and the applicability of the model to the patient and decision.

The patient-reported layer represents symptoms, functioning, treatment burden, preferences, goals, adherence experience, and outcomes that may not be recoverable from clinical records or computational outputs. Patient-reported outcomes require explicit definitions, suitable instruments, prespecified timing, and plans for incomplete information if they are to contribute interpretable evidence [16]. In the negotiated architecture, patient reporting is not reduced to satisfaction or used merely to confirm a professional interpretation. It constitutes an independently inspectable model layer.

These layers may converge. A reported adverse effect may correspond with a laboratory change and a model-estimated exposure increase. They may also diverge. Patient-generated information can serve different needs for patients and professionals, producing disagreement about relevance, interpretation, workload, privacy, or the action that should follow [17]. A patient may report an intolerable burden while a computational model forecasts favorable disease control; a clinician may question the reliability of a home measurement that the patient considers meaningful.

The theory proposes that disagreement first be classified rather than immediately resolved. Data disagreement concerns incompatible observations. Semantic disagreement concerns differing meanings assigned to the same variable. Model disagreement occurs when plausible computational approaches generate materially different outputs. Goal and value disagreements concern desired outcomes or acceptable burdens. Authority disagreement concerns who may decide, update, or override. Temporal disagreement occurs when historical and current representations conflict.

No universal weighting formula is proposed. Clinical measurements should not automatically override lived experience, and patient reports should not automatically invalidate a technically justified estimate. The appropriate disposition depends on the decision, consequences of error, quality of evidence, urgency, and authority of the participants. Possible dispositions include verification, additional measurement, provisional acceptance, preservation of parallel representations, escalation, restriction, or suspension. The three-layer structure is therefore an inspectable architecture for research and evaluation, not evidence that heterogeneous representations can always be reconciled or that negotiation necessarily improves pharmacotherapy outcomes.

## Updating, Disagreement, and Representation of Alternative Futures

A negotiated twin must distinguish routine data incorporation from changes that alter model meaning. Verification, validation, uncertainty quantification, and context-of-use assessment are therefore required when patient information, parameters, model structure, or intended use changes [18]. Updating should not be presumed beneficial: previously adequate calibration can deteriorate as outcomes and predictor relationships change [19].

The process must also separate individual patient-state change from population, institutional, or data-generating shift [20]. New evidence should be classified as a data correction, patient-state update, parameter update, structural revision, or intended-use modification. Minor updates within qualified boundaries may be accepted through an authorized pathway; material changes should trigger reassessment or requalification.

When layers conflict, the proposed disagreement register records the source, type, consequence, and disposition of the conflict. The system may request verification, preserve parallel representations, restrict use, or escalate review rather than forcing immediate fusion. Alternative futures should likewise be represented as conditional scenarios with explicit assumptions, uncertainty, time horizon, and causal status. Predictive associations must not be presented as treatment

effects unless the required causal assumptions and evidence are defensible [21]. Parallel futures are therefore aids to deliberation, not forecasts of what will necessarily happen.

### *Evaluation, Governance, and Intended-Use Boundaries*

Evaluation must address the complete human–model–workflow system. Early live assessment should examine actual interaction, error recovery, workload, adaptations, and use within clinical practice [22]. Comparative studies should transparently report model inputs, outputs, human involvement, failure handling, and analysis [23], while protocols should prospectively specify intended users, integration, input management, output use, and foreseeable errors [24]. Evaluation should document early live clinical use and prospectively specify and transparently report the distinctive behavior of an AI-supported intervention [22–24].

The proposed evaluation pathway progresses from component verification to retrospective, external, temporal, subgroup, human-factors, workflow, clinical-utility, and postdeployment assessment. Governance should address fairness, universality, traceability, usability, robustness, and explainability throughout the lifecycle [25]. Passing one stage would justify only the next bounded stage, not unrestricted medication use. Material changes in model structure, population, workflow, or intended purpose should trigger requalification, restriction, rollback, or suspension.

### *Implications for Precision Pharmacotherapy*

Model-informed dosing becomes clinically meaningful only when pharmacometric outputs are translated into usable and reviewable bedside decisions [26]. The negotiated twin extends that requirement by incorporating patient experience, visible disagreement, conditional futures, and explicit authority boundaries.

**Figure 2** presents the original conceptual synthesis for updating, disagreement, and alternative-future representation, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

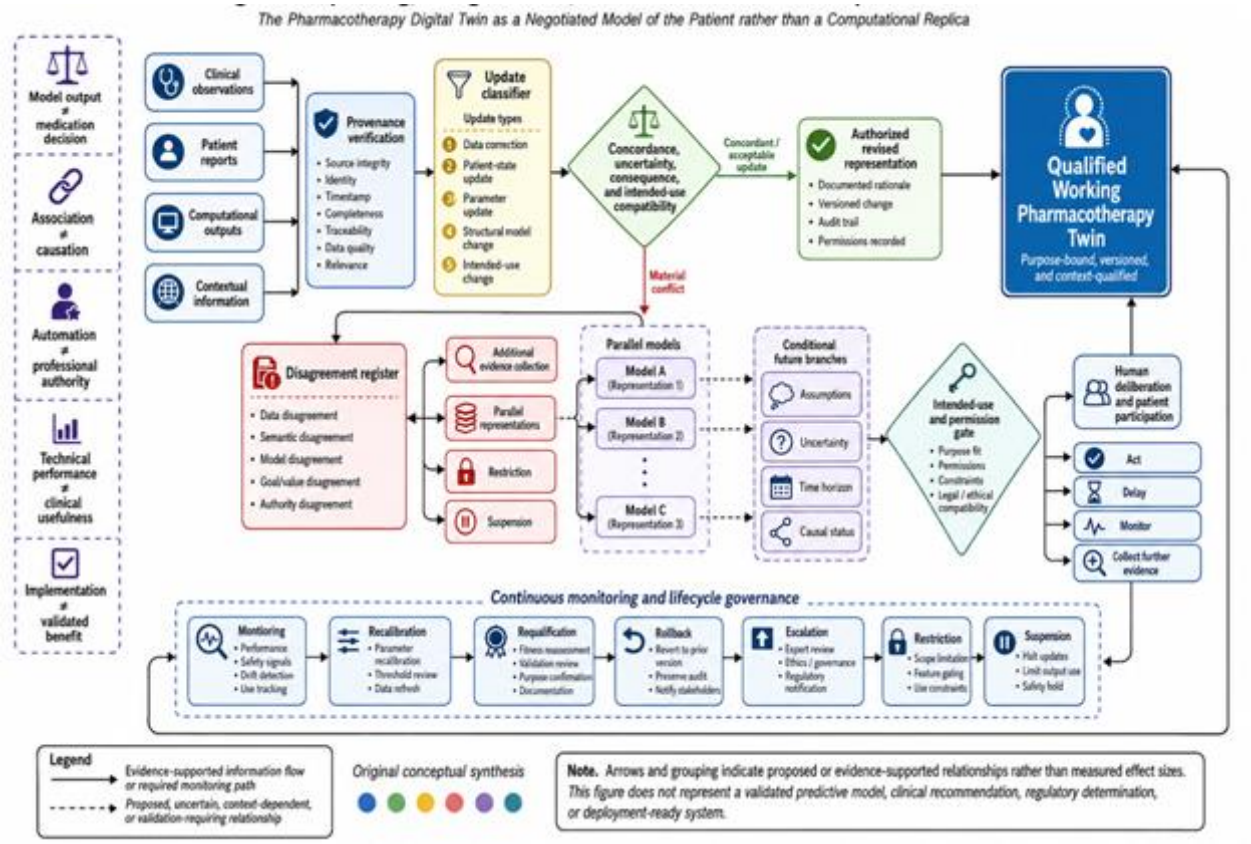


**Figure 2.** Updating, Disagreement, and Alternative-Future Representation

Precision dosing already seeks to connect population evidence with individual characteristics [27]. The negotiated theory adds that individualization must also remain conditional on model validity, workflow integration, practical validation, and evidence of clinical utility [28]. It further rejects reducing precision to genomic variation alone: treatment history, environment, behavior, access, and patient experience may also shape medication response and feasibility [29]. The theory does not authorize automated dose changes or identify a universally preferred pharmacotherapy method.

### *Limitations*

The theory remains conceptual. Explanations may create an appearance of understanding without establishing validity, causality, robustness, or fairness [30]. Implementation would require prospective evidence, workflow redesign, resources, professional acceptance, and organizational change [31]. Computational representation also raises unresolved concerns involving autonomy, privacy, ownership, identity, and responsibility [32]. Moreover, satisfactory aggregate performance can conceal systematic failure in underserved groups [33]. The theory's constructs, transition gates, and disagreement classifications require empirical testing across medications, populations, pharmacy settings, infrastructures, and jurisdictions.

## Conclusion

A pharmacotherapy digital twin should not be treated as a computational duplicate whose increasing detail progressively eliminates uncertainty. It is better understood as a partial, purpose-bound, and revisable representation constructed through interaction among clinical evidence, computational inference, patient-reported knowledge, and human authority.

Negotiated Pharmacotherapy Digital-Twin Theory proposes inspectable model layers, a negotiation workspace, a disagreement register, conditional alternative futures, permissioned updating, provenance, and an intended-use contract. These constructs may support research on how pharmacotherapy models represent patients, communicate uncertainty, distribute responsibility, and remain contestable.

The contribution is theoretical rather than prescriptive. It does not establish improved prediction, safer medication decisions, equitable benefit, regulatory acceptance, or deployment readiness. Its central boundary is that the patient always exceeds the model: no digital representation can independently determine what the patient is, what matters to them, or what medication decision should follow.

ACKNOWLEDGMENTS: None
CONFLICT OF INTEREST: None
FINANCIAL SUPPORT: None
ETHICS STATEMENT: None

## References

1. 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
2. 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
3. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, et al. Digital twins for health: A scoping review. NPJ Digit Med. 2024;7(1):77. doi:10.1038/s41746-024-01073-0
4. Drummond D, Gonsard A. Definitions and characteristics of patient digital twins being developed for clinical use: Scoping review. J Med Internet Res. 2024;26. doi:10.2196/58504
5. 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
6. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA. 2017;318(6):517-8. doi:10.1001/jama.2017.7797
7. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-53. doi:10.1126/science.aax2342
8. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: A roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. doi:10.1038/s41591-019-0548-6
9. Greenhalgh T, Wherton J, Papoutsi C, Lynch J, Hughes G, A'Court C, et al. Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. J Med Internet Res. 2017;19(11). doi:10.2196/jmir.8775
10. Elwyn G, Durand MA, Song J, Aarts J, Barr PJ, Berger Z, et al. A three-talk model for shared decision making: Multistage consultation process. BMJ. 2017;359. doi:10.1136/bmj.j4891
11. Amann J, Blasimme A, Vayena E, Frey D, Madai VI. Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. BMC Med Inform Decis Mak. 2020;20(1):310. doi:10.1186/s12911-020-01332-6
12. Topol EJ. High-performance medicine: The convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7
13. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. Nat Med. 2022;28(1):31-8. doi:10.1038/s41591-021-01614-0
14. 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
15. 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
16. Calvert M, Kyte D, Mercieca-Bebber R, Slade A, Chan AW, King MT, et al. Guidelines for inclusion of patient-reported outcomes in clinical trial protocols: The SPIRIT-PRO extension. JAMA. 2018;319(5):483-94. doi:10.1001/jama.2017.21903
17. Reading MJ, Merrill JA. Converging and diverging needs between patients and providers who are collecting and using patient-generated health data: An integrative review. J Am Med Inform Assoc. 2018;25(6):759-71. doi:10.1093/jamia/ocy006
18. 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
19. Davis SE, Lasko TA, Chen G, Siew ED, Matheny ME. Calibration drift in regression and machine learning models for acute kidney injury. J Am Med Inform Assoc. 2017;24(6):1052-61. doi:10.1093/jamia/ocx030
20. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041
21. 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
22. 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
23. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Nat Med. 2020;26(9):1364-74. doi:10.1038/s41591-020-1034-x
24. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ; SPIRIT-AI and CONSORT-AI Working Group. Guidelines for clinical trial

protocols for interventions involving artificial intelligence: The SPIRIT-AI extension. Nat Med. 2020;26(9):1351-63. doi:10.1038/s41591-020-1037-7
25. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP, et al. FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388. doi:10.1136/bmj-2024-081554
26. 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
27. Gonzalez D, Rao GG, Bailey SC, Brouwer KLR, Cao Y, Crona DJ, et al. Precision dosing: Public health need, proposed framework, and anticipated impact. Clin Transl Sci. 2017;10(6):443-54. doi:10.1111/cts.12490
28. 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
29. Peck RW. Precision medicine is not just genomics: The right dose for every patient. Annu Rev Pharmacol Toxicol. 2018;58:105-22. doi:10.1146/annurev-pharmtox-010617-052446
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
31. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17(1):195. doi:10.1186/s12916-019-1426-2
32. Huang PH, Kim KH, Schermer M. Ethical issues of digital twins for personalized health care service: Preliminary mapping study. J Med Internet Res. 2022;24(1). doi:10.2196/33081
33. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-82. doi:10.1038/s41591-021-01595-0