Available online **www.ijpras.com**

International Journal of Pharmaceutical Research & Allied Sciences, 2026, 15(1):123-132
https://doi.org/10.51847/cDo7KLqMC1

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Mechanistic Digital Twins of Drug–Target Dynamics for Counterfactual Perturbation, Adaptive Simulation, and Uncertainty-Aware Molecular Decision-Making

Andrei Popescu[1]*, Mihai Ionescu[1], Elena Stan[2], Cristina Radu[3]

[1]Department of Digital Twins and Drug–Target Dynamics, Faculty of Pharmacy, University of Bucharest, Bucharest, Romania.
[2]Department of Counterfactual Perturbation and Adaptive Simulation, Faculty of Pharmacy, University of Agricultural Sciences Cluj-Napoca, Cluj-Napoca, Romania.
[3]Department of Uncertainty-Aware Molecular Decision-Making, Faculty of Pharmacy, Polytechnic University of Bucharest, Bucharest, Romania.

***Email:** andrei.popescu@usamv.ro*

ABSTRACT

Molecular simulation can characterize conformational ensembles, energy landscapes, ligand interactions, and transition pathways, yet an individual simulation remains a conditional representation rather than a continuously updated molecular counterpart. This article proposes an original mechanistic digital-twin architecture for drug–target dynamics that separates latent molecular state, intervention, physical time, evidence-update time, decision horizon, and observation. The architecture links a persistent entity and provenance record with mechanistic transition models, dynamic conformational-state estimation, counterfactual ligand, mutation, and environmental perturbations, adaptive simulation, and active information acquisition. Uncertainty is propagated across sampling, force fields, parameters, observations, software implementations, learned components, and competing mechanistic hypotheses rather than compressed into a single confidence value. Counterfactual branches are evaluated against structural, thermodynamic, kinetic, and functional consequences, with validation requirements determined by the intended context of use. A governance layer distinguishes exploratory hypothesis generation from bounded molecular decision support and requires escalation, revision, or deferral when provenance, identifiability, validation, or uncertainty conditions are inadequate. The original contribution is the integration of established molecular-simulation methods into a traceable state–intervention–observation–update architecture rather than the proposal of a new simulation algorithm. The framework is intended to generate testable propositions and organize evidence acquisition, but it does not establish predictive validity, causal therapeutic effects, clinical usefulness, regulatory acceptability, universal applicability, or deployment readiness. Its value therefore depends on prospective, context-specific computational and experimental validation.

Key words: *Digital twin, Molecular dynamics, Conformational ensemble, Energy landscape, Sampling, Drug–target recognition*

Received: 21 January 2026; **Revised:** 23 March 2026; **Accepted:** 25 March 2026

INTRODUCTION

Molecular dynamics can connect atomic interactions with conformational change, while machine-learning methods can accelerate representation and simulation. Allosteric behavior further demonstrates that molecular function often emerges from shifting ensembles rather than one static structure [1-3]. Nevertheless, trajectories,

learned representations, and energy landscapes are commonly generated as separate analytical products without a persistent mechanism for evidence updating or counterfactual comparison.

This fragmentation limits interpretation. A calculated affinity may omit kinetic consequences; a sampled conformation may be statistically frequent yet mechanistically irrelevant; and a mutation prediction may depend on an incompletely represented baseline ensemble. Combining more models does not itself resolve these problems because disagreement can arise from sampling, force fields, observation error, structural assumptions, or genuinely competing mechanisms.

This Original Mechanistic Digital-Twin Architecture Article proposes a molecular-scale framework in which the drug–target system retains an explicit computational identity across simulations, interventions, observations, and model revisions. The twin is defined by the traceable relationship among these elements, not by molecular dynamics, machine learning, a structural model, or an experimental dataset acting alone.

The objective is to establish a coherent architecture for state representation, counterfactual perturbation, adaptive information acquisition, uncertainty propagation, validation, and bounded decision use. The contribution is conceptual and non-empirical. It organizes established methods into proposed relationships and decision rules that require system-specific evaluation rather than claiming that a validated molecular twin already exists.

From molecular simulation to continuously updated digital twins

Digital-twin scholarship distinguishes a twin from a standalone model through persistent entity correspondence, connected information, purpose-specific representation, and iterative updating [4-6]. Transferring these principles to molecular science requires caution: a drug–target complex is not continuously observable, and its relevant state may include unmeasured conformations, chemical states, environmental conditions, and transition probabilities.

Accordingly, "continuously updated" should mean evidence-triggered continuity rather than uninterrupted real-time streaming. New trajectories, structural constraints, affinities, kinetic measurements, mutational observations, or functional data may revise state or parameter beliefs while preserving links to the prior model version. Retraining a model without such provenance would create a replacement model rather than a traceably updated twin.

Figure 1 presents the original conceptual synthesis for mechanistic digital twin of drug–target dynamics, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

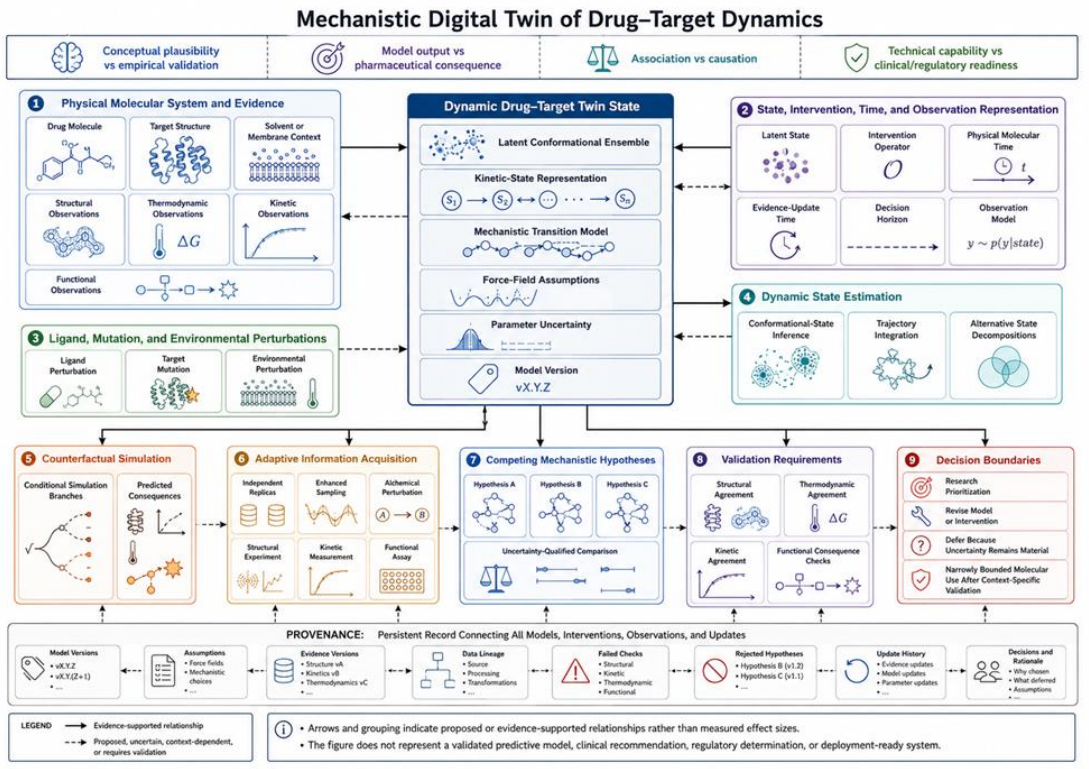

Figure 1. Mechanistic Digital Twin of Drug–Target Dynamics.

The figure is an original conceptual synthesis developed for "Mechanistic Digital Twins of Drug–Target Dynamics for Counterfactual Perturbation, Adaptive Simulation, and Uncertainty-Aware Molecular Decision-Making." 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.

A molecular twin must also retain a declared context of use. A representation intended to compare ligand substitutions may require different states, observations, and validation from one intended to explain mutation-dependent residence time. Updating therefore should not expand the decision claim silently; a change in molecular endpoint or intervention domain requires reconsideration of representation and evidence adequacy.

The proposed transition from simulation to twin is consequently organizational rather than semantic. Calling an existing workflow a digital twin adds no credibility unless entity identity, model versions, observations, interventions, uncertainties, and validation boundaries remain inspectable. This architecture does not establish that complete molecular observability, continuous state recovery, or universally reliable prediction is achievable.

State, intervention, time, and observation representation

The molecular state should be represented as an ensemble or reduced kinetic description that preserves coordinates relevant to the intended mechanism. Collective variables can organize long-timescale behavior, but their adequacy depends on whether they retain the slow transitions and molecular distinctions needed for the stated question [7]. A single coordinate snapshot is therefore an observation of configuration, not the complete twin state.

State decomposition also requires kinetic justification. Markov-state approaches can connect metastable states through transition probabilities, yet their conclusions depend on discretization, lag time, sampling coverage, and validation [8]. The proposed architecture consequently stores the state-definition method and competing decompositions rather than presenting one partition as uniquely true.

Data-assisted collective variables may improve state representation by learning coordinates associated with specified transitions or labels [9]. However, learned variables inherit the limitations of their trajectories, supervision targets, and molecular domain. Predictive separation of classes does not establish that a variable is causally sufficient, physically complete, or transferable to a new ligand, mutation, or environment.

The architecture separates state, intervention, time, and observation. An intervention specifies a ligand transformation, mutation, protonation change, environmental alteration, or restraint. Physical time denotes simulated molecular evolution; update time records the sequence of incoming evidence; and the decision horizon defines the future consequence being queried. Observations are linked to latent state through an explicit observation model rather than treated as direct access to molecular reality.

Proposed drug–target digital-twin architecture

The proposed architecture integrates mechanistic, multiscale, learned, and observational components because no single representation spans all relevant molecular scales or evidence types [10]. Its central unit is a versioned twin state connected to the physical molecular entity, an intervention registry, an observation model, simulation engines, an uncertainty ledger, and a context-of-use declaration.

Physics-informed learning may embed conservation relationships, differential constraints, symmetries, or mechanistic priors within learned modules [11]. Such constraints can narrow implausible behavior, but they do not independently demonstrate empirical validity. The architecture therefore records whether a relationship is physically imposed, statistically learned, experimentally calibrated, or merely proposed.

Learned coarse-grained force fields illustrate how reduced representations may bridge molecular scales while introducing assumptions about resolution, training support, and transferability [12]. Accordingly, each surrogate or reduced model is linked to its parent representation, training domain, parameter version, and reconstruction limits. Outputs that cross these limits must trigger uncertainty escalation rather than silent extrapolation.

Ten connected components constitute the proposed twin: entity identity and provenance; latent molecular state; mechanistic transition model; intervention operator; observation model; conformational-state estimator; counterfactual simulation engine; adaptive information-acquisition controller; uncertainty and competing-hypothesis ledger; and validation and governance gate. No component alone qualifies as the twin, and replacing one component requires an auditable version transition.

Information flows are conditional. Observations update state or parameter beliefs only after compatibility assessment; interventions generate parallel branches from a shared baseline; uncertainty guides additional simulation or experimentation; and validation determines whether outputs remain exploratory or may support a

narrowly bounded molecular decision. This integration is an original conceptual synthesis that requires prospective computational and experimental testing and does not represent a validated architecture.

Dynamic conformational-state estimation

Dynamic state estimation treats the conformational ensemble as a latent, time-dependent quantity inferred from trajectories and observations. VAMPnets demonstrate that kinetic state assignments and transition structure can be learned jointly, although the resulting states remain dependent on training coverage and model selection [13]. Enhanced sampling can be coupled to learned latent variables so that new trajectories refine the coordinates used to explore the landscape. RAVE illustrates this iterative relationship between dimensionality reduction, reweighting, and sampling, but it does not guarantee that the learned coordinate captures every relevant slow process [14].

Generative approaches can propose configurations from an estimated equilibrium distribution and may reduce barriers to exploring separated regions of configuration space [15]. Equilibrium plausibility, however, does not establish kinetic fidelity, mechanistic relevance, or transferability to a perturbed chemical system.

The proposed twin therefore stores alternative state decompositions, posterior state uncertainty, sampling history, and the relationship between latent states and measurable observations. State estimates may be revised as evidence accumulates, but revisions must preserve provenance and must not be interpreted as recovery of a uniquely true molecular state.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in mechanistic digital twins of drug–target dynamics for counterfactual perturbation.

Table 1. Definitions, components, relationships, and intended uses in Mechanistic Digital Twins of Drug–Target Dynamics for Counterfactual Perturbation

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Latent conformational state	Static structures omit ensemble behavior	Trajectories, structural constraints, kinetic descriptors	**Proposed:** infer an uncertainty-qualified ensemble state	Mechanistically interpretable state tracking	Kinetic-state validation and observation compatibility [13]	State aliasing or missing slow modes	An inferred state is not molecular ground truth
Adaptive state representation	Fixed coordinates may miss emerging transitions	Existing and newly acquired trajectories	**Proposed:** revise collective variables during updating	Better targeting of unresolved landscape regions	Reweighting and sampling diagnostics [14]	Feedback bias or unstable coordinates	Adaptive refinement does not guarantee convergence
Generative state proposals	Rare states may remain inaccessible	Learned equilibrium distribution and physical constraints	**Proposed:** propose candidate configurations for testing	Expanded hypothesis generation	Distributional and energetic validation [15]	Nonphysical or kinetically misleading samples	Generated states are candidates, not validated mechanisms
Intervention representation	Perturbations are often mixed with observed outcomes	Ligand, mutation, protonation, and environmental definitions	**Proposed:** encode interventions separately from state	Traceable counterfactual comparison	Common baseline and transformation validity	Intervention–baseline incompatibility	A simulated contrast is conditional, not a clinical causal effect
Observation model	Experiments do not directly reveal complete latent state	Structural, affinity, kinetic, and functional data	**Proposed:** map latent states to expected observations	Evidence assimilation with explicit measurement assumptions	Measurement-model calibration	Observation error or non-identifiability	Agreement with one observation does not validate the full twin

Provenance and decision boundary	Updates can obscure model identity and evidentiary scope	Model versions, assumptions, evidence history, context of use	**Proposed:** preserve an auditable update chain and use gate	Reproducibility and bounded interpretation	Version consistency and context-specific validation	Silent model replacement or scope expansion	Technical plausibility is not decision readiness

Counterfactual ligand, mutation, and environmental perturbations

Ligand counterfactuals compare outcomes under explicitly defined chemical transformations while holding the relevant baseline representation constant. Nonequilibrium alchemical calculations demonstrate that relative binding free energies can be estimated across ligand series, but their validity remains conditional on chemical overlap, sampling, and system preparation [16].

Mutation counterfactuals alter the target rather than the ligand. Alchemical calculations have been used to estimate mutation-dependent changes in inhibitor binding, showing how mechanistic hypotheses about resistance may be examined computationally [17]. Such estimates remain molecular and system-specific and cannot by themselves establish clinical resistance.

Environmental counterfactuals include changes in protonation, solvent, ionic composition, membrane context, cofactors, or conformational restraints. Protein flexibility can modulate both binding thermodynamics and kinetics, demonstrating that an intervention cannot always be interpreted independently of the ensemble from which it is applied [18].

The proposed engine therefore branches interventions from a shared, uncertainty-qualified baseline and records which assumptions remain invariant. Counterfactual admissibility requires a defined intervention, a sufficiently represented baseline state, compatible parameterization, and an outcome that the molecular model can legitimately address.

Adaptive simulation and active information acquisition

Adaptive simulation changes computational allocation in response to what has already been learned. Reinforcement-based sampling can reward exploration of underrepresented conformational regions, providing a mechanism for directing new trajectories toward unresolved portions of the landscape [19].

Data-driven collective variables offer a complementary strategy in which accumulating trajectory information is used to refine the coordinates that guide enhanced sampling [20]. The resulting efficiency depends on whether the learned coordinate remains relevant to the molecular question and does not conceal alternative slow processes. Deep learning of slow modes may improve access to rare events by focusing simulation on dynamically informative coordinates [21]. Nevertheless, accelerated discovery of transitions does not demonstrate that all relevant states have been found or that the learned representation is transferable.

The proposed controller selects the next simulation or experiment according to expected reduction of decision-relevant uncertainty. It may request replicas, enhanced sampling, an alchemical perturbation, or a discriminating experiment, but the selection remains advisory and must preserve human and experimental oversight.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for mechanistic digital twins of drug–target dynamics for counterfactual perturbation.

Table 2. Testable propositions, evidence requirements, and validation criteria for Mechanistic Digital Twins of Drug–Target Dynamics for Counterfactual Perturbation

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
State-sufficiency test	Determine whether represented states preserve intervention-relevant dynamics	Trajectories → latent-state model	Constrains intervention and observation models	Stable kinetic interpretation under resampling [13]	Missing slow coordinate	Review state definitions and diagnostics	Research use until state adequacy is demonstrated

Update-consistency test	Ensure new evidence changes beliefs traceably	Evidence → state or parameter update	Connected to provenance ledger	Reproducible update with preserved version history [14]	Unstable or contradictory updates	Examine evidence compatibility	No bounded use after unexplained model replacement
Counterfactual-admissibility test	Verify that a perturbation has a defensible baseline	Baseline state + intervention → comparison branches	Precedes simulation engine	Transformation and parameter compatibility [16]	Invalid chemical or state mapping	Approve intervention definition	No causal interpretation without admissibility
Adaptive-value test	Select informative next actions	Uncertainty → simulation or experiment request	Uses state estimator and hypothesis ledger	Reduced prespecified uncertainty without biased inference [19]	Exploration unrelated to decision question	Approve resource allocation and experiments	Adaptive operation remains supervised
Rare-event test	Assess whether sampling captures relevant transitions	Learned slow modes → enhanced trajectories	Updates state representation	Replicated transition evidence [21]	Apparent convergence from biased coordinates	Require independent checks	No claim of landscape completeness
Validation-triangulation test	Compare predicted consequences across evidence layers	Predictions → structural, kinetic, functional checks	Feeds governance gate	Coherent, context-specific agreement	Success on one convenient endpoint	Define prospective acceptance criteria	Single-endpoint agreement is insufficient
Decision-boundary test	Restrict use to supported claims	Validation and uncertainty → governance outcome	Final architecture gate	Evidence proportional to intended consequence	Scope expansion beyond validation	Authorize, defer, or reject intended use	No clinical or regulatory inference from molecular evidence alone

Uncertainty propagation and competing mechanistic hypotheses

Uncertainty should be decomposed across stochastic sampling, implementation reproducibility, and learned-model extrapolation rather than reported as one confidence value [22-24]. These sources may interact but require different diagnostic and corrective actions.

Independent replicas are necessary because a single trajectory can support a false mechanistic conclusion through chance sampling or incomplete exploration [22]. Replicate agreement reduces one source of uncertainty but does not correct shared force-field or structural misspecification.

Cross-software comparison can reveal implementation-dependent discrepancies, while learned molecular potentials require explicit attention to out-of-domain behavior [23, 24]. Neither reproducibility nor a narrow uncertainty interval proves that the represented mechanism is correct.

The proposed hypothesis ledger therefore preserves alternative state models, transition pathways, binding mechanisms, and perturbation explanations. Hypotheses are updated separately until discriminating evidence is available; averaging incompatible mechanisms into one prediction would conceal structural uncertainty.

Validation against structural, kinetic, and functional consequences

Validation must correspond to the intended use rather than to the most convenient benchmark. Large-scale assessments of binding free-energy calculations show why performance across realistic projects is more informative than success in an isolated, favorable system [25].

When a claim concerns binding or conformational kinetics, validation should examine transition pathways, state populations, or rate-related behavior. Atomistic simulation combined with Markov modeling has demonstrated how association kinetics may be connected to detailed molecular pathways [26].

When a proposed mechanism includes biological function, structural agreement alone is inadequate. Molecular analysis of GPCR-mediated arrestin activation illustrates how conformational transitions can be related to experimentally observed functional consequences [27].

The architecture therefore requires layered validation: numerical and reproducibility checks; structural or thermodynamic agreement; kinetic consequences where relevant; and functional evidence when function is claimed. Molecular validation does not establish therapeutic benefit, patient outcome, or manufacturing performance.

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for failure modes, boundary conditions, governance needs, and implementation priorities for mechanistic digital twins of drug–target dynamics for counterfactual perturbation.

Table 3. Failure modes, boundary conditions, governance needs, and implementation priorities for Mechanistic Digital Twins of Drug–Target Dynamics for Counterfactual Perturbation

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
Affinity ranking versus mechanistic explanation	Ligand comparison	Predictive ordering versus causal interpretation	Prospective project-relevant validation [25]	Separate statistical error from mechanism uncertainty	Restrict claim or acquire discriminating evidence	Bounded molecular prioritization	Therapeutic superiority
Structural fit versus kinetic validity	Plausible bound structure	Geometric agreement versus transition realism	Pathway or rate evidence [26]	Retain kinetic alternatives	Extend sampling or obtain kinetic measurements	Better-supported mechanism	Correct residence time from structure alone
Conformational change versus functional consequence	Allosteric or signaling claim	Structural plausibility versus biological effect	Compatible functional evidence [27]	Propagate observation-model uncertainty	Test downstream consequence	Stronger mechanistic coherence	Clinical efficacy
Precision versus reproducibility	Narrow predicted interval	Apparent certainty versus implementation robustness	Replica and workflow consistency	Preserve between-run and between-engine variance	Repeat, compare, or revise workflow	More credible computational evidence	Absence of model bias
Adaptation versus traceability	Repeated model updating	Learning efficiency versus provenance	Complete version and evidence history	Record update-specific uncertainty	Freeze, audit, or roll back model	Inspectable evidence evolution	That later versions are automatically better
Research utility versus decision use	Output appears actionable	Speed versus evidentiary adequacy	Context-specific validation and governance review	Escalate unresolved structural uncertainty	Limit to research, revise, or defer	Prevention of unsupported use	Clinical, regulatory, or deployment readiness

Decision boundaries and governance

Model credibility should be proportional to the intended context of use and the consequences of an incorrect decision [28]. A twin used to prioritize simulations therefore requires a different evidentiary standard from one proposed to guide compound selection or interpret resistance.

Credibility assessment should connect model purpose, verification, validation, uncertainty, documentation, and evidentiary adequacy rather than treating them as isolated checkboxes [29]. Governance begins before simulation by defining the permissible claim and the evidence needed to support it.

The proposed decision gate has four outcomes: research prioritization, model or intervention revision, deferral because material uncertainty remains, and narrowly bounded molecular use after context-specific validation. These outcomes are not maturity scores and do not imply regulatory qualification.

Human responsibility remains necessary for defining questions, approving intervention assumptions, assessing experimental compatibility, interpreting unresolved hypotheses, and preventing scope expansion. Automated generation of molecular outputs does not transfer scientific or decision accountability to the model.

Limitations and implementation priorities

Force fields remain conditional representations of molecular interactions. Improvements designed for folded and intrinsically disordered proteins demonstrate that performance can depend on the conformational regime and evaluation target [30].

Transferability across folded, disordered, mutated, ligand-bound, and environmentally altered states must therefore be tested rather than assumed. Force-field development spanning distinct protein-state regimes illustrates both the importance and difficulty of this requirement [31].

Additional limitations include incomplete sampling, hidden chemical states, uncertain protonation, state aliasing, observation-model error, adaptive feedback bias, mechanistic non-identifiability, and loss of provenance during updating. No uncertainty method can guarantee that all omitted mechanisms have been recognized.

Figure 2 presents the original conceptual synthesis for adaptive counterfactual-simulation and evidence-updating cycle, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

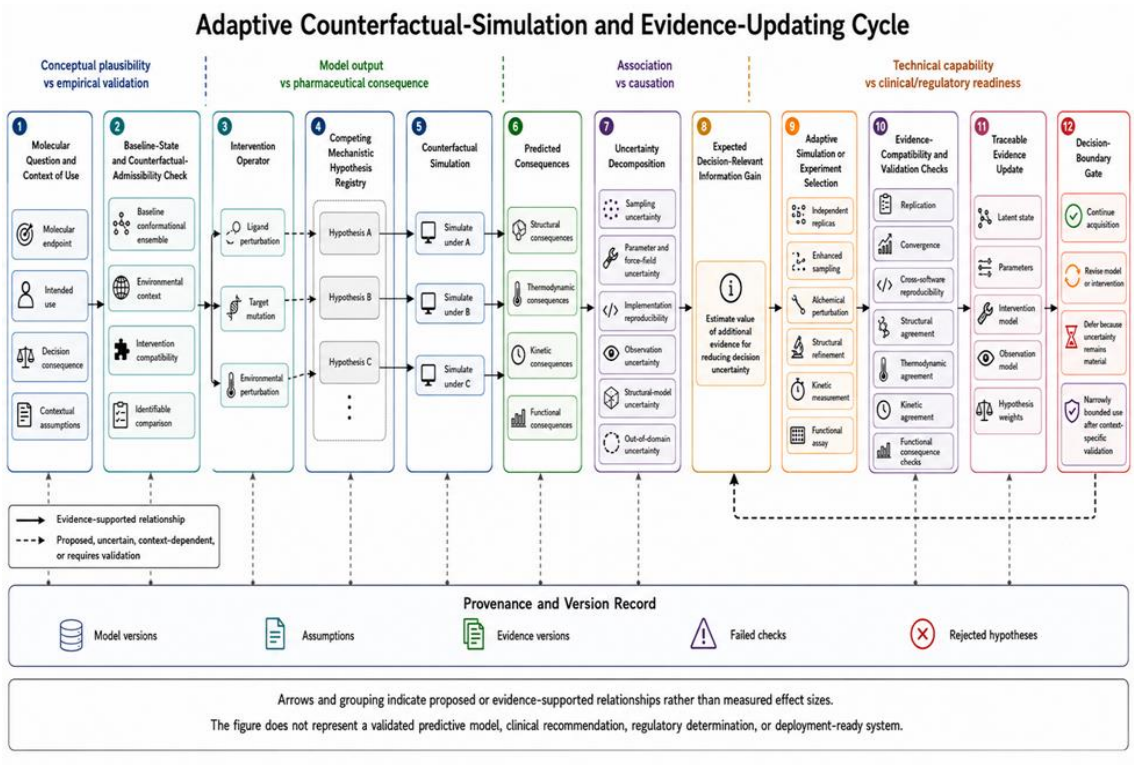

Figure 2. Adaptive Counterfactual-Simulation and Evidence-Updating Cycle.

The figure is an original conceptual synthesis developed for "Mechanistic Digital Twins of Drug–Target Dynamics for Counterfactual Perturbation, Adaptive Simulation, and Uncertainty-Aware Molecular Decision-Making." 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.

Implementation should begin with bounded molecular questions for which interventions, observations, and validation endpoints can be prespecified. Early priorities are versioned provenance, independent simulation replicas, uncertainty decomposition, and prospective comparison with discriminating experiments. Expansion to consequential decisions should occur only after the relevant context has been separately validated.

CONCLUSION

This article proposes a mechanistic drug–target digital twin as a persistent, traceable relationship among molecular state, interventions, time, observations, simulation models, evidence updates, uncertainty, and decision boundaries. Its principal contribution is an architecture for organizing dynamic conformational-state estimation, counterfactual perturbation, adaptive information acquisition, competing mechanistic hypotheses, and layered validation without treating any component as sufficient alone. The framework may support testable molecular

questions and disciplined evidence acquisition, but its outputs remain conditional on representation, force fields, sampling, observations, and context of use. It does not establish predictive validity, clinical benefit, regulatory acceptability, universal applicability, or deployment readiness.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Hollingsworth SA, Dror RO. Molecular dynamics simulation for all. Neuron. 2018;99(6):1129-43. doi:10.1016/j.neuron.2018.08.011
2. Noé F, Tkatchenko A, Müller KR, Clementi C. Machine learning for molecular simulation. Annu Rev Phys Chem. 2020;71:361-90. doi:10.1146/annurev-physchem-042018-052331
3. Wodak SJ, Paci E, Dokholyan NV, Berezovsky IN, Horovitz A, Li J, et al. Allostery in its many disguises: From theory to applications. Structure. 2019;27(4):566-78. doi:10.1016/j.str.2019.01.003
4. Jones D, Snider C, Nassehi A, Yon J, Hicks B. Characterising the Digital Twin: A systematic literature review. CIRP J Manuf Sci Technol. 2020;29(Pt A):36-52. doi:10.1016/j.cirpj.2020.02.002
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. Laubenbacher R, Mehrad B, Shmulevich I, Trayanova N. Digital twins in medicine. Nat Comput Sci. 2024;4(3):184-91. doi:10.1038/s43588-024-00607-6
7. Noé F, Clementi C. Collective variables for the study of long-time kinetics from molecular trajectories: Theory and methods. Curr Opin Struct Biol. 2017;43:141-7. doi:10.1016/j.sbi.2017.02.006
8. Husic BE, Pande VS. Markov state models: From an art to a science. J Am Chem Soc. 2018;140(7):2386-96. doi:10.1021/jacs.7b12191
9. Sultan MM, Pande VS. Automated design of collective variables using supervised machine learning. J Chem Phys. 2018;149(9):094106. doi:10.1063/1.5029972
10. Alber M, Buganza Tepole A, Cannon WR, De S, Dura-Bernal S, Garikipati K, et al. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. NPJ Digit Med. 2019;2:115. doi:10.1038/s41746-019-0193-y
11. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. Nat Rev Phys. 2021;3(6):422-40. doi:10.1038/s42254-021-00314-5
12. Wang J, Olsson S, Wehmeyer C, Pérez A, Charron NE, de Fabritiis G, et al. Machine learning of coarse-grained molecular dynamics force fields. ACS Cent Sci. 2019;5(5):755-67. doi:10.1021/acscentsci.8b00913
13. Mardt A, Pasquali L, Wu H, Noé F. VAMPnets for deep learning of molecular kinetics. Nat Commun. 2018;9(1):5. doi:10.1038/s41467-017-02388-1
14. Ribeiro JML, Bravo P, Wang Y, Tiwary P. Reweighted autoencoded variational Bayes for enhanced sampling (RAVE). J Chem Phys. 2018;149(7):072301. doi:10.1063/1.5025487
15. Noé F, Olsson S, Köhler J, Wu H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. Science. 2019;365(6457). doi:10.1126/science.aaw1147
16. Gapsys V, Pérez-Benito L, Aldeghi M, Seeliger D, van Vlijmen H, Tresadern G, et al. Large scale relative protein ligand binding affinities using non-equilibrium alchemy. Chem Sci. 2020;11(4):1140-52. doi:10.1039/C9SC03754C
17. Hauser K, Negron C, Albanese SK, Ray S, Steinbrecher T, Abel R, et al. Predicting resistance of clinical Abl mutations to targeted kinase inhibitors using alchemical free-energy calculations. Commun Biol. 2018;1:70. doi:10.1038/s42003-018-0075-x
18. Amaral M, Kokh DB, Bomke J, Wegener A, Buchstaller HP, Eggenweiler HM, et al. Protein conformational flexibility modulates kinetics and thermodynamics of drug binding. Nat Commun. 2017;8(1):2276. doi:10.1038/s41467-017-02258-w

19. Shamsi Z, Cheng KJ, Shukla D. Reinforcement learning based adaptive sampling: REAPing rewards by exploring protein conformational landscapes. J Phys Chem B. 2018;122(35):8386-95. doi:10.1021/acs.jpcb.8b06521
20. Bonati L, Rizzi V, Parrinello M. Data-driven collective variables for enhanced sampling. J Phys Chem Lett. 2020;11(8):2998-3004. doi:10.1021/acs.jpclett.0c00535
21. Bonati L, Piccini G, Parrinello M. Deep learning the slow modes for rare events sampling. Proc Natl Acad Sci U S A. 2021;118(44). doi:10.1073/pnas.2113533118
22. Knapp B, Ospina-Forero L, Deane CM. Avoiding false positive conclusions in molecular simulation: The importance of replicas. J Chem Theory Comput. 2018;14(12):6127-38. doi:10.1021/acs.jctc.8b00391
23. Loeffler HH, Bosisio S, Duarte Ramos Matos G, Suh D, Roux B, Mobley DL, et al. Reproducibility of free energy calculations across different molecular simulation software packages. J Chem Theory Comput. 2018;14(11):5567-82. doi:10.1021/acs.jctc.8b00544
24. Wen M, Tadmor EB. Uncertainty quantification in molecular simulations with dropout neural network potentials. NPJ Comput Mater. 2020;6:124. doi:10.1038/s41524-020-00390-8
25. Schindler CEM, Baumann H, Blum A, Böse D, Buchstaller HP, Burgdorf L, et al. Large-scale assessment of binding free energy calculations in active drug discovery projects. J Chem Inf Model. 2020;60(11):5457-74. doi:10.1021/acs.jcim.0c00900
26. Plattner N, Doerr S, De Fabritiis G, Noé F. Complete protein-protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling. Nat Chem. 2017;9(10):1005-11. doi:10.1038/nchem.2785
27. Latorraca NR, Wang JK, Bauer B, Townshend RJL, Hollingsworth SA, Olivieri JE, et al. Molecular mechanism of GPCR-mediated arrestin activation. Nature. 2018;557(7705):452-6. doi:10.1038/s41586-018-0077-3
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. 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
30. Huang J, Rauscher S, Nawrocki G, Ran T, Feig M, de Groot BL, et al. CHARMM36m: An improved force field for folded and intrinsically disordered proteins. Nat Methods. 2017;14(1):71-3. doi:10.1038/nmeth.4067
31. Robustelli P, Piana S, Shaw DE. Developing a molecular dynamics force field for both folded and disordered protein states. Proc Natl Acad Sci U S A. 2018;115(21). doi:10.1073/pnas.1800690115