Available online **www.ijpras.com**

International Journal of Pharmaceutical Research & Allied Sciences, 2025, 14(1):35-45
https://doi.org/10.51847/haQ6MUGaAS

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Connecting Femtosecond Motions to Pharmacological Outcomes through a Multiscale Architecture for Residence Time, Allostery, and Functional Response

Sofía Martínez[1]*, Carlos Gómez[1], Lucia Navarro[2]

[1]*Department of Multiscale Molecular Pharmacology, Faculty of Pharmacy, Polytechnic University of Madrid, Madrid, Spain.*
[2]*Department of Residence Time and Allosteric Modeling, Faculty of Pharmacy, University of Barcelona, Barcelona, Spain.*

***Email:** sofia.martinez@upm.es*

ABSTRACT

Molecular pharmacology spans a temporal range that extends from femtosecond-scale bond and electronic motions to clinical-scale patterns of exposure, efficacy, and failure. The central difficulty is not merely computational cost. It is the absence of a disciplined architecture for deciding when information obtained at one scale can legitimately constrain claims at another. This article proposes the Temporal–Mechanistic Translation Architecture, an original conceptual synthesis that connects molecular motion, metastable conformational ensembles, kinetic transitions, residence time, allosteric communication, signalling, and systems-level functional response. The architecture is organized as six linked layers separated by explicit validation gates. Each gate specifies the evidence required to move from one inferential level to the next and prevents thermodynamic, kinetic, structural, functional, and translational quantities from being treated as interchangeable. The framework introduces three governing principles: adjacent-scale validation, bidirectional consistency, and decision-bounded interpretation. It further derives testable propositions concerning state definition, kinetic bottlenecks, residence-time surrogates, allosteric transmission, functional coupling, and exposure-dependent response. Validation is treated as a staged process requiring trajectory stability, kinetic reproducibility, orthogonal experimental constraint, perturbational testing, identifiability analysis, and prospective decision evaluation. The architecture does not claim that atomistic simulation can independently predict pharmacological outcomes, nor does it imply clinical, regulatory, or deployment readiness. Its intended contribution is to organize mechanistic reasoning, expose hidden assumptions, identify failure points, and define the evidence needed before a cross-scale explanation can be treated as decision-relevant.

Key words: *Molecular dynamics, Conformational ensemble, Residence time, Allostery, Binding kinetics, Enhanced sampling*

Received: 08 November 2024; **Revised:** 01 February 2025; **Accepted:** 07 February 2025

INTRODUCTION

Molecular dynamics simulation has become a central instrument for examining protein flexibility, ligand recognition, solvent organization, and conformational change. Its value lies in representing molecular systems as time-dependent ensembles rather than static structures. Yet atomistic trajectories are not direct recordings of pharmacological action. They are model-based realizations generated under a specified force field, solvent model, boundary condition, and sampling protocol, and they therefore require interpretation at the levels of both statistical mechanics and biological function [1].

A second source of conceptual error arises when equilibrium affinity, free energy, association rate, dissociation rate, and residence time are treated as interchangeable descriptions of binding. These quantities answer different questions. Affinity describes an equilibrium relationship; free-energy calculations estimate thermodynamic differences; association and dissociation rates describe kinetic paths; and residence time is the inverse of dissociation rate only within the assumptions of the kinetic model used. A pharmacological explanation that substitutes one of these quantities for another risks collapsing distinct mechanisms into a single, misleading metric [2].

The need for a new conceptual contribution follows from the fact that molecular simulation, binding kinetics, and protein flexibility each describe complementary but incomplete parts of drug action. Molecular dynamics clarifies possible motions, kinetic models organize transitions, and conformational flexibility can modulate both thermodynamic and kinetic barriers, but none alone establishes downstream functional or pharmacological consequence [1-3]. This article therefore proposes the Temporal–Mechanistic Translation Architecture (TMTA), an original, non-empirical synthesis for connecting microscopic dynamics to residence time, allostery, signalling, and functional response. The TMTA is not presented as a validated predictive model. It is a reasoning architecture that specifies permissible inferences, required validation evidence, and decision boundaries across scales.

The temporal-scale mismatch in molecular pharmacology

The temporal-scale mismatch is often described as a gap between short simulations and slow biological events, but duration alone is not the decisive issue. The more fundamental problem is representational: rapid coordinates must be transformed into states and transitions that retain information relevant to slower observables. Markov state models offer one route by partitioning configuration space into metastable states and estimating transition probabilities at a selected lag time. Their scientific use depends on state adequacy, timescale separation, and validation of approximate Markovian behavior rather than on aggregate trajectory duration alone [4].

This distinction is important because rare binding or association events may be reconstructed from many distributed trajectories even when no single trajectory spans the full experimental timescale. Atomic-detail simulations combined with kinetic modeling have demonstrated that encounter complexes, alternative paths, and rate-limiting steps can be organized into a coherent association mechanism. However, such reconstructions remain conditional on the force field, sampling strategy, and state decomposition used to estimate the kinetic network [5].

Adaptive and distributed sampling can further extend kinetic inference into regimes far beyond the length of individual trajectories. The relevant achievement is not that simulation “reaches” seconds in a literal sense, but that many short observations are assembled into a model of slow transitions. The distinction between aggregate sampling and uninterrupted dynamics must therefore remain explicit. Long aggregate time does not remove uncertainty in state definitions, hidden barriers, or kinetic model form [6]. Within the TMTA, this mismatch is handled by the adjacent-scale validation principle: information may move from one scale to the next only when a defined observable links the two and when the uncertainty introduced by coarse-graining is reported rather than silently discarded.

Molecular motions, metastable states, and kinetic transitions

The microscopic layer of the TMTA begins with molecular motion but does not end there. Motion refers to time-dependent coordinate change; a conformational ensemble is a probability distribution over accessible configurations; a metastable state is a region with relatively rapid internal relaxation and slower escape; and a kinetic transition is a probabilistic passage between such states. Deep variational approaches can learn metastable representations and transition structure directly from trajectories, thereby reducing dependence on hand-selected coordinates. Their output nevertheless remains an operational model of kinetics rather than a unique physical partition of the molecule [7].

Slow collective variables provide a bridge between high-dimensional motion and interpretable kinetic organization. Time-lagged autoencoders identify coordinates that retain slowly decorrelating information, which can be useful for exposing conformational modes missed by geometric intuition. Yet a slowly varying coordinate is not automatically causal, mechanistically sufficient, or pharmacologically relevant. It may correlate with a transition while omitting the actual barrier-crossing degree of freedom or a solvent-mediated contribution [8].

Reaction-coordinate construction is therefore treated in the TMTA as a constrained inference problem. A useful coordinate should preserve information about both thermodynamic organization and kinetic progression, while remaining testable against alternative representations. Methods that jointly consider past and future information

illustrate how a latent variable can be optimized to capture slow transitions and free-energy structure [9]. The operational distinctions used throughout the proposed architecture are summarized in **Table 1**.

Table 1. Definitions, components, relationships, and intended uses in Connecting Femtosecond Motions to Pharmacological Outcomes through a Multiscale.

Construct	Operational definition	Scale represented	Relationship within the proposed architecture	Intended use	Prohibited inference	Evidence anchor
Molecular motion	Time-dependent coordinate fluctuations generated under an explicit Hamiltonian and simulation protocol	Femtoseconds to nanoseconds	Supplies raw dynamical observations to the ensemble layer	Identify candidate local rearrangements and solvent-coupled motions	A trajectory is not a direct observation of pharmacological effect	[1]
Conformational ensemble	Probability-weighted set of accessible molecular configurations	Picoseconds to milliseconds	Organizes motions into populations rather than a single structure	Compare ligand-, mutation-, or environment-dependent population shifts	Dominant simulated population does not establish the biologically dominant state	[3]
Metastable state	Operational region with faster internal relaxation than escape	Nanoseconds to milliseconds	Defines coarse-grained nodes for kinetic modeling	Estimate populations, exchange paths, and state lifetimes	State boundaries are not assumed to be unique or observer-independent	[4]
Slow collective variable	Reduced coordinate retaining slowly decorrelating dynamical information	System-dependent	Provides a candidate axis for state discovery or enhanced sampling	Reduce dimensionality and expose slow modes	Slow variation does not establish causal control	[8]
Reaction coordinate	Variable or low-dimensional set intended to track transition progress	Barrier-crossing timescales	Links free-energy organization to transition likelihood	Define sampling bias, committor-like tests, or kinetic analysis	A convenient coordinate is not presumed complete	[9]
Kinetic transition model	Probabilistic representation of movement among states	Microseconds to seconds and beyond	Converts sampled trajectories into rates or transition probabilities	Compare mechanisms and generate rate-linked hypotheses	Aggregate sampling time is not equivalent to continuous physical time	[5]
Residence-time layer	Representation of dissociation kinetics and the persistence of target engagement	Seconds to hours	Receives pathway and barrier information from microscopic layers	Rank kinetic hypotheses or estimate dissociation behavior under calibration	Residence time is not interchangeable with affinity or efficacy	[2]
Allosteric-propagation layer	Representation of how perturbations redistribute structural, dynamic, or population relationships across a protein	Microseconds to functional timescales	Connects local binding events to distal conformational consequences	Generate perturbation-testable coupling hypotheses	Statistical coupling is not treated as causal transmission	[7]
Signalling and functional-response layer	Context-dependent mapping from conformational states and partner interactions to pathway activity	Seconds to days	Connects molecular state occupancy to cellular output	Define pathway-specific tests and functional observables	Structural change alone does not establish signalling	[3]

					magnitude or phenotype	
Decision-context layer	Explicit statement of the drug-discovery or target-assessment question	Project-dependent	Constrains which outputs are relevant and what validation is sufficient	Support bounded prioritization, not universal prediction	No architecture output is treated as clinical or regulatory proof	Proposed TMTA construct

***Note**:* Evidence anchors support the established concepts in each row. The layer structure, permitted-inference rules, and decision boundaries are original proposed elements of the TMTA. The table contains no validated thresholds, rankings, or effect estimates.

Architecture connecting microscopic dynamics to residence time

The residence-time component of the TMTA is designed as a translation layer rather than a single computational method. Accelerated unbinding approaches can generate ensembles of exit paths and may support calibrated estimation or ranking of residence behavior, provided that pathway diversity, perturbation strength, and comparison to kinetic experiments are explicit [10]. The required inferential bridges between ultrafast molecular motion and pharmacological response are organized in **Figure 1**.

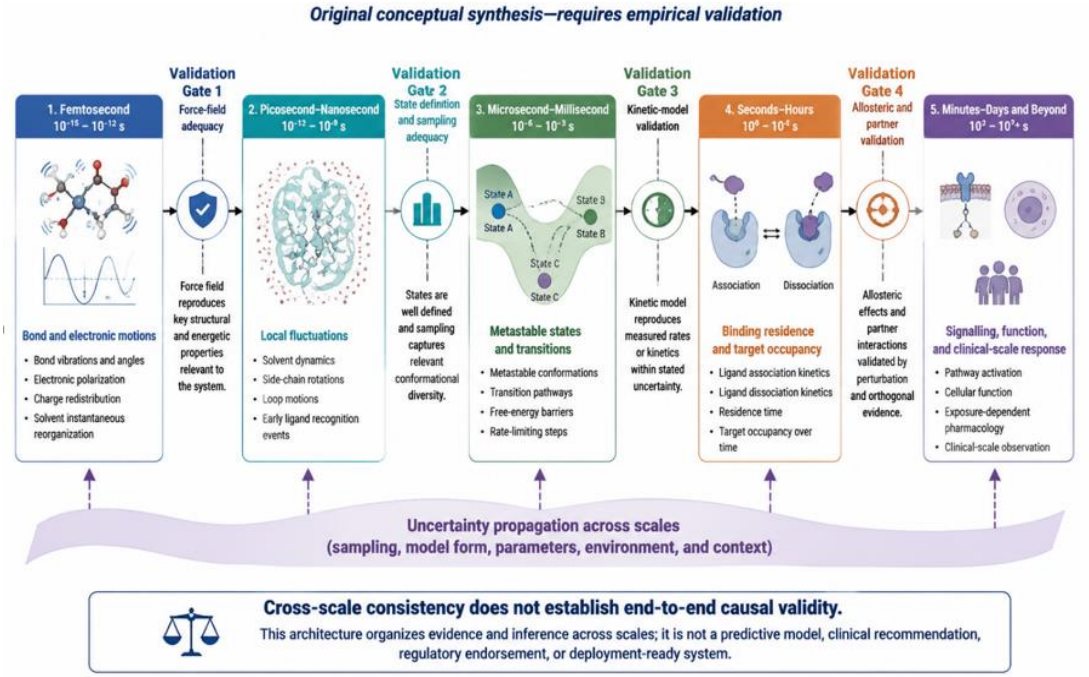

Figure 1. Temporal Bridge from Femtosecond Molecular Motion to Clinical-Scale Drug Response

Relative residence-time estimation illustrates why the layer must distinguish between observable, surrogate, and decision claim. Ensemble steered simulations can produce descriptors that correlate with relative residence behavior, but the work required to force an unbinding event is not itself the physical dissociation time. The TMTA therefore requires every computational residence-time claim to specify whether it concerns absolute rate estimation, rank ordering, pathway classification, or mechanistic explanation [11].

Biased pathway methods add a second route by accelerating escape and comparing the organization of unbinding barriers among ligands. Their utility depends on collective-variable adequacy, reweighting assumptions, and the degree to which the biased path preserves the relevant transition mechanism [12]. The proposed architecture therefore separates three outputs: a pathway hypothesis, a kinetic surrogate, and a decision claim. Progression from one to the next requires independent evidence. A pathway hypothesis may be supported by repeated exit routes; a kinetic surrogate requires calibration against measured kinetics; and a decision claim requires evidence that the kinetic difference matters in the biological and exposure context of interest.

Propagation of allosteric information across protein structures

Allostery is incorporated into the TMTA as an ensemble-level phenomenon rather than a mandatory physical pathway connecting two sites. Contemporary accounts encompass population shifts, altered fluctuations, network

reorganization, structural relays, and changes in energetic coupling. These mechanisms are not mutually exclusive, and different representations may be appropriate for different proteins or perturbations [13]. The architecture therefore avoids defining allostery solely as a visible conformational displacement.

Trajectory-derived interaction networks can identify long-range couplings that are difficult to infer from static structure. Neural relational approaches, for example, can learn residue relationships from simulation data and generate hypotheses about dynamic communication across a protein [14]. In the TMTA, such edges are classified as candidate couplings until they survive perturbational testing. A strong statistical dependency may reflect a common driver, an omitted solvent variable, or geometric constraint rather than directional information transfer.

Experimental observations of ligand-specific flexibility changes support the possibility that distal regulation can occur with limited change in the average structure. In the lac repressor, altered conformational flexibility has been linked to long-range allosteric behavior, illustrating how dynamic redistribution can mediate function without a single rigid pathway [15]. The proposed allosteric layer therefore contains three competing explanatory modes: population redistribution, network-mediated coupling, and state-dependent mechanical propagation. These modes can coexist, but the manuscript does not treat any as universally dominant. Their relative credibility depends on whether mutations, ligand substitutions, spectroscopic observables, or other perturbations discriminate among them.

Linking conformational change to signalling and functional response

The next inferential step concerns how conformational change becomes signalling. GPCR research shows why this mapping cannot be reduced to a binary active-versus-inactive switch. Receptors occupy ensembles whose populations depend on ligand, membrane environment, intracellular partner, and post-translational state. Distinct conformational features may preferentially support different signalling partners or pathway outputs [16]. The TMTA therefore treats functional response as a conditional mapping from state occupancy to partner engagement, not as an intrinsic label attached to one structure.

Mechanistic studies of arrestin activation illustrate how specific receptor contacts can stabilize functionally important partner conformations. Molecular simulation can help identify candidate interaction sequences, while biochemical and structural experiments test whether those contacts are required for activation [17]. Within the architecture, such evidence supports an interaction-to-competence claim: a molecular configuration can be linked to the ability to engage or activate a signalling partner. It does not yet establish whole-cell response magnitude, pathway duration, or organism-level effect.

Combinatorial phosphorylation patterns further demonstrate that functional translation depends on molecular context. Distinct receptor modification patterns can tune arrestin conformation and alter downstream signalling behavior [18]. This supports a layered view in which binding state, modification state, partner state, and cellular context jointly shape response. A conformational change is therefore treated as functionally meaningful only when an orthogonal observable establishes a link to partner engagement or pathway activity. The TMTA expressly prohibits direct inference from a simulated state population to a clinical outcome.

Multiscale integration of simulation, kinetics, and systems pharmacology

Residence time enters pharmacology through several pathways rather than a single linear mechanism. It can influence the persistence of target occupancy, interact with pharmacokinetic exposure, and matter differently depending on target turnover, rebinding, compartmentalization, and vulnerability. Experimental work examining in vivo occupancy shows that residence time can affect target engagement through multiple routes, but it does not independently determine efficacy [19]. The TMTA therefore rejects residence time as a universal optimization objective.

Quantitative systems pharmacology provides a natural upper layer for integrating exposure, target engagement, pathway response, and disease context. Its relevance lies not in replacing molecular detail but in determining whether a microscopic difference can propagate under realistic concentrations, temporal patterns, feedback loops, and biological constraints [20]. In the proposed architecture, molecular parameters enter a systems model only when their uncertainty, identifiability, and intended decision role are specified. The resulting system-level output is treated as a conditional projection, not as an automatic consequence of the molecular model.

Cross-scale integration also requires interoperable definitions of parameters, observables, noise models, and estimation problems. Formal specifications for parameter estimation help separate biological assumptions from implementation choices and make model–data relationships auditable [21]. The TMTA adopts this principle through a translation contract for every layer interface. Each contract records the incoming variable, its units and

uncertainty, the transformation applied, the observable used for validation, and the decision that the output is permitted to inform. This contract is a proposed governance device rather than a validated technical standard.

Validation criteria across temporal and biological scales

Validation in the TMTA is staged and scale-specific. Independent replicas are required because stochastic trajectories can support apparently stable conclusions that disappear when variability is properly sampled [22]. The complete proposed architecture and its validation gates are presented in **Figure 2**.

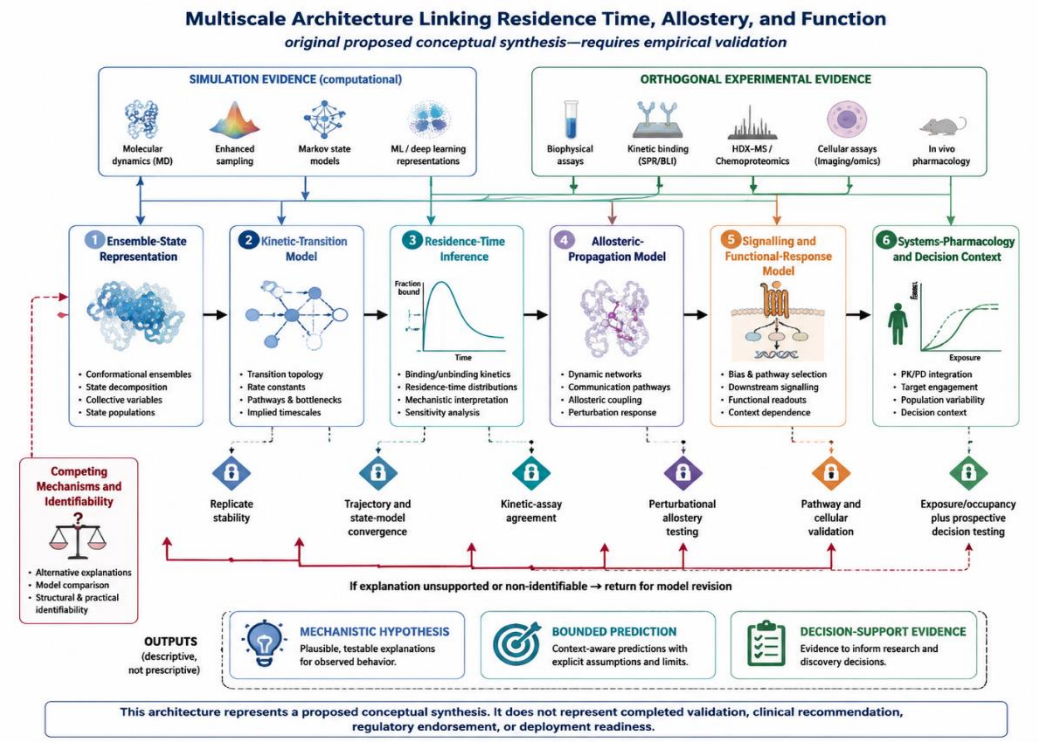

Figure 2. Multiscale Architecture Linking Residence Time, Allostery, and Function

Implementation reproducibility forms a second validation gate. Agreement across simulation engines or workflows is meaningful only when force fields, protonation states, initial conditions, algorithms, restraint definitions, and analysis conventions are harmonized. Cross-software studies of free-energy calculations show that nominally similar methods can diverge when these details are not controlled [23]. Reproducibility is therefore necessary but not sufficient: two implementations may agree while sharing the same incorrect physical assumption. The propositions and validation requirements derived from the architecture are specified in **Table 2.**

Table 2. Testable propositions, evidence requirements, and validation criteria for Connecting Femtosecond Motions to Pharmacological Outcomes through a Multiscale.

Proposition	Proposed relationship	Observable prediction	Required computational evidence	Required experimental or cross-scale evidence	Failure criterion	Decision boundary	Evidence anchor
P1: State adequacy	A useful metastable-state model preserves the slow processes relevant to the intended kinetic claim	Implied timescales and state populations remain stable across defensible lag times and state resolutions	Replicate trajectories, state sensitivity analysis, transition-model diagnostics	Orthogonal conformational or kinetic observables sensitive to the proposed states	Kinetic conclusions change materially under plausible state definitions	State models may organize hypotheses but cannot support rate claims without stability	[4]
P2: Reaction-coordinate sufficiency	A reaction coordinate adequate for	Alternative coordinate sets yield	Committor-related tests, reweighting	Perturbations expected to alter the proposed barrier	Coordinate-dependent mechanism	Enhanced sampling supports	[9]

	mechanistic interpretation captures transition progress, not merely slow correlation	consistent barrier ordering and transition-path statistics	checks, hidden-barrier analysis	change measured kinetics	with no discriminating evidence	mechanism only within validated coordinate scope	
P3: Residence-surrogate validity	A computational unbinding surrogate is useful only after calibration to the kinetic quantity claimed	Surrogate preserves prospective ligand ordering or calibrated rate behavior	Independent unbinding ensembles and uncertainty estimates	Direct association/dissociation measurements in the same target context	Retrospective fit without prospective transfer or pathway instability	Surrogate output is not treated as physical residence time by default	[10]
P4: Allosteric-coupling causality	A candidate coupling becomes mechanistic only when targeted perturbation changes the predicted distal state or function	Mutations or ligand changes selectively alter the proposed coupling and downstream observable	Competing network models and perturbation simulations	Structural, spectroscopic, biochemical, or functional perturbation tests	Correlation persists but perturbation fails to alter the predicted outcome	Statistical coupling remains a hypothesis	[14]
P5: Ensemble-to-function mapping	Functional response depends on state populations, partner context, and regulatory state	Changes in state occupancy correspond to partner-specific activation patterns	Ensemble comparison under ligand, partner, and modification conditions	Partner-binding, pathway, or cellular assays that discriminate outputs	Similar state shifts produce inconsistent functional effects across contexts	No direct structural-to-clinical inference	[16]
P6: Cross-scale consistency	Parameters transferred between layers retain meaning, units, uncertainty, and context	System-level predictions remain consistent under propagated molecular uncertainty	Explicit translation contracts and uncertainty propagation	Target-occupancy, exposure, pathway, and functional measurements	Output depends on hidden unit conversion, parameter reinterpretation, or untracked calibration	Cross-scale use is limited to the decision for which the contract was qualified	[20]
P7: Reproducibility adequacy	A mechanistic claim should survive stochastic and implementation variation	Main conclusion persists across replicas and harmonized implementations	Independent replicas, workflow provenance, cross-engine or cross-laboratory checks	Repetition under matched assay or measurement conditions	Conclusion is driven by one seed, one workflow, or one analysis choice	Reproducibility does not equal mechanistic truth	[22]
P8: Identifiability and discrimination	Competing mechanisms require data that distinguish their predictions	Alternative models make separable predictions under a designed perturbation	Profile, posterior, or sensitivity analyses across candidate mechanisms	Prospective discriminating experiments	Multiple mechanisms fit all available observables equally well	No mechanism is preferred solely because it is more detailed	Proposed TMTA criterion; evidence discussed in Section 12

Note: The propositions, validation ladder, failure criteria, and decision boundaries are original proposed elements of the TMTA. Evidence anchors support the methodological basis for each proposition. No universal numerical cutoffs, validation scores, or regulatory thresholds are asserted.

The strongest validation claim requires convergence of three distinct forms of evidence: replicate stability, implementation reproducibility, and orthogonal experimental constraint [22-24]. Augmented kinetic models demonstrate how experimental observations can be incorporated into simulation-based state models rather than used only as retrospective decoration [24]. The TMTA extends this logic by requiring that validation data test an aspect of the model not exhausted by calibration. Agreement with fitted observables is model consistency; successful prediction of a withheld perturbation is stronger mechanistic evidence; and successful prospective use in a defined decision is a separate level of readiness.

Uncertainty, non-identifiability, and competing mechanistic explanations

A multiscale model can fit observed data while leaving its internal mechanism unresolved. Structural non-identifiability arises when distinct parameter combinations produce identical outputs even with ideal data, whereas practical non-identifiability arises when available data are too noisy or uninformative to determine parameters precisely [25]. These problems are especially important when molecular parameters are passed into larger systems models, because apparent precision at one layer can conceal ambiguity inherited from another.

Representation learning introduces a related uncertainty. Time-lagged latent variables may identify useful slow modes, but their performance can deteriorate in nonlinear, degenerate, or poorly sampled systems [26]. The TMTA therefore requires comparison of alternative state spaces and reaction coordinates before a learned representation is interpreted mechanistically. If multiple latent descriptions reproduce the same kinetics but imply different physical explanations, the correct conclusion is non-identifiability, not confidence.

Force-field uncertainty creates a further class of competing explanation. A change in simulated population or transition barrier may reflect actual molecular physics, inadequate sampling, or limitations in the potential energy model. Work aimed at balancing folded and disordered protein states demonstrates how force-field design can reshape conformational ensembles [27]. The architecture therefore separates sampling uncertainty, representation uncertainty, parameter uncertainty, and model-form uncertainty. These categories are propagated independently wherever possible, because increasing sampling cannot correct a systematically biased Hamiltonian, and improving a force field cannot compensate for a non-identifiable systems model.

Application scenarios in lead optimization and target assessment

The TMTA is intended for bounded use in lead optimization when a target-specific kinetic hypothesis exists. Industry experience with residence-time programs indicates that kinetic optimization is most useful when connected to target biology, assay design, exposure, and a clearly defined decision objective [28]. The architecture therefore begins each application with a decision statement such as: “Does modification of this exit-path interaction produce a kinetically persistent and biologically relevant increase in target engagement under the intended exposure profile?” It does not begin with the generic instruction to maximize residence time.

A second use case concerns target assessment and kinetic selectivity. Broad surveys of kinase binding kinetics show that compounds with similar equilibrium potency can display substantially different kinetic profiles across related targets [29]. This creates a plausible role for simulation in generating hypotheses about why a compound dissociates slowly from one target and rapidly from another. The decision boundary remains strict: computational differences must be tested with direct kinetic assays, and kinetic selectivity must be evaluated in relevant cellular and exposure contexts before it is linked to therapeutic selectivity.

A third use case is negative decision support. The TMTA can justify stopping a mechanistic line of inquiry when the required bridge cannot be validated. Examples include a residence-time surrogate that fails prospectively, an allosteric network that does not survive perturbation, a systems model whose key parameters remain non-identifiable, or a functional response that is not robust across cellular context. In this sense, the architecture is not only a route for connecting scales; it is also a formal method for preventing unsupported connection.

Limitations and research priorities

The first limitation is that enhanced sampling does not eliminate the dependence of free-energy and kinetic interpretation on collective variables, reweighting assumptions, and convergence diagnostics [30]. Hidden barriers may remain unsampled, and apparent convergence along a selected coordinate may coexist with unresolved orthogonal motion. The TMTA therefore cannot guarantee that a simulated landscape is complete. It can only require explicit diagnostics, alternative coordinates, and boundary statements that restrict the inference.

Future progress depends jointly on better exploration of complex landscapes and more transferable physical representations, but neither advance alone guarantees pharmacological prediction [30, 31]. Emerging learned

potentials may improve energetic accuracy and extend the feasible scale of atomistic simulation, yet their transferability across proteins, ligands, protonation states, and chemical environments must be established independently [31]. Kinetic fidelity is an especially important research priority because small energetic errors can produce large changes in rare-event rates.

A broader limitation is that the architecture remains conceptual. Its modules, translation contracts, validation gates, and propositions require empirical testing across diverse targets and decision contexts. Priority studies should prospectively compare alternative state models, kinetic surrogates, and allosteric explanations; propagate uncertainty into target-occupancy and systems models; and assess whether architecture-guided decisions outperform simpler baselines. Even successful evaluation would not imply universal validity. Qualification would remain specific to the target class, assay system, exposure regime, and decision question examined.

CONCLUSION

The central challenge in connecting molecular dynamics to pharmacological outcome is not the existence of many timescales alone. It is the temptation to treat evidence from one scale as sufficient for another. The proposed Temporal–Mechanistic Translation Architecture addresses this problem by separating molecular motion, conformational ensembles, kinetic transitions, residence time, allosteric propagation, signalling, and systems-level response into linked but non-interchangeable layers.

Three principles govern the architecture: every cross-scale step requires adjacent-scale validation, every mechanistic explanation must remain consistent in both upward and downward directions, and every output must be bounded by a specific decision context. These principles convert a broad multiscale ambition into a set of falsifiable propositions, explicit translation contracts, and defined stopping rules.

The architecture is an original conceptual synthesis rather than a validated predictive system. Its value lies in clarifying what would have to be true before femtosecond-scale motions could support claims about residence time, allostery, function, or pharmacological response. By making uncertainty, non-identifiability, and competing explanations visible, it provides a disciplined basis for future computational–experimental programs while preventing plausible molecular narratives from being mistaken for established pharmacology.

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. Decherchi S, Cavalli A. Thermodynamics and kinetics of drug-target binding by molecular simulation. Chem Rev. 2020;120(23):12788-833. doi:10.1021/acs.chemrev.0c00534
3. 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;8:2276. doi:10.1038/s41467-017-02258-w
4. 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
5. 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
6. Paul F, Wehmeyer C, Abualrous ET, Wu H, Crabtree MD, Schöneberg J, et al. Protein-peptide association kinetics beyond the seconds timescale from atomistic simulations. Nat Commun. 2017;8:1095. doi:10.1038/s41467-017-01163-6
7. Mardt A, Pasquali L, Wu H, Noé F. VAMPnets for deep learning of molecular kinetics. Nat Commun. 2018;9:5. doi:10.1038/s41467-017-02388-1

8. Wehmeyer C, Noé F. Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. J Chem Phys. 2018;148(24):241703. doi:10.1063/1.5011399
9. Wang Y, Ribeiro JML, Tiwary P. Past-future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics. Nat Commun. 2019;10:3573. doi:10.1038/s41467-019-11405-4
10. Kokh DB, Amaral M, Bomke J, Grädler U, Musil D, Buchstaller HP, et al. Estimation of drug-target residence times by τ-random acceleration molecular dynamics simulations. J Chem Theory Comput. 2018;14(7):3859-69. doi:10.1021/acs.jctc.8b00230
11. Potterton A, Husseini FS, Southey MWY, Bodkin MJ, Heifetz A, Coveney PV, et al. Ensemble-based steered molecular dynamics predicts relative residence time of A2A receptor binders. J Chem Theory Comput. 2019;15(5):3316-30. doi:10.1021/acs.jctc.8b01270
12. Callegari D, Lodola A, Pala D, Rivara S, Mor M, Rizzi A, et al. Metadynamics simulations distinguish short- and long-residence-time inhibitors of cyclin-dependent kinase 8. J Chem Inf Model. 2017;57(2):159-69. doi:10.1021/acs.jcim.6b00679
13. 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
14. Zhu J, Wang J, Han W, Xu D. Neural relational inference to learn long-range allosteric interactions in proteins from molecular dynamics simulations. Nat Commun. 2022;13:1661. doi:10.1038/s41467-022-29331-3
15. Glasgow A, Hobbs HT, Perry ZR, Wells ML, Marqusee S, Kortemme T. Ligand-specific changes in conformational flexibility mediate long-range allostery in the lac repressor. Nat Commun. 2023;14:1179. doi:10.1038/s41467-023-36798-1
16. Latorraca NR, Venkatakrishnan AJ, Dror RO. GPCR dynamics: Structures in motion. Chem Rev. 2017;117(1):139-55. doi:10.1021/acs.chemrev.6b00177
17. 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
18. Latorraca NR, Masureel M, Hollingsworth SA, Heydenreich FM, Suomivuori CM, Brinton C, et al. How GPCR phosphorylation patterns orchestrate arrestin-mediated signaling. Cell. 2020;183(7):1813-25.e18. doi:10.1016/j.cell.2020.11.014
19. Lee KSS, Yang J, Niu J, Ng CJ, Wagner KM, Dong H, et al. Drug-target residence time affects in vivo target occupancy through multiple pathways. ACS Cent Sci. 2019;5(9):1614-24. doi:10.1021/acscentsci.9b00770
20. Bradshaw EL, Spilker ME, Zang R, Bansal L, He H, Jones RDO, et al. Applications of quantitative systems pharmacology in model-informed drug discovery: Perspective on impact and opportunities. CPT Pharmacometrics Syst Pharmacol. 2019;8(11):777-91. doi:10.1002/psp4.12463
21. Schmiester L, Schälte Y, Bergmann FT, Camba T, Dudkin E, Egert J, et al. PEtab—Interoperable specification of parameter estimation problems in systems biology. PLoS Comput Biol. 2021;17(1). doi:10.1371/journal.pcbi.1008646
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. Olsson S, Wu H, Paul F, Clementi C, Noé F. Combining experimental and simulation data of molecular processes via augmented Markov models. Proc Natl Acad Sci U S A. 2017;114(31):8265-70. doi:10.1073/pnas.1704803114
25. Wieland FG, Hauber AL, Rosenblatt M, Tönsing C, Timmer J. On structural and practical identifiability. Curr Opin Syst Biol. 2021;25:60-9. doi:10.1016/j.coisb.2021.03.005
26. Chen W, Sidky H, Ferguson AL. Capabilities and limitations of time-lagged autoencoders for slow mode discovery in dynamical systems. J Chem Phys. 2019;151(6):064123. doi:10.1063/1.5112048
27. 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
28. Schuetz DA, de Witte WEA, Wong YC, Knasmueller B, Richter L, Kokh DB, et al. Kinetics for drug discovery: An industry-driven effort to target drug residence time. Drug Discov Today. 2017;22(6):896-911. doi:10.1016/j.drudis.2017.02.002

29. Georgi V, Schiele F, Berger BT, Steffen A, Marin Zapata PA, Briem H, et al. Binding kinetics survey of the drugged kinome. J Am Chem Soc. 2018;140(46):15774-82. doi:10.1021/jacs.8b08048
30. Bussi G, Laio A. Using metadynamics to explore complex free-energy landscapes. Nat Rev Phys. 2020;2(4):200-12. doi:10.1038/s42254-020-0153-0
31. Wang T, He X, Li M, Li Y, Bi R, Wang Y, et al. Ab initio characterization of protein molecular dynamics with AI2BMD. Nature. 2024;635(8040):1019-27. doi:10.1038/s41586-024-08127-z