



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Residence Time Should Be Treated as an Ensemble Property in Molecular Simulation of Drug Binding

James O'Leary^{1*}, Aisling Dunne¹, Sean O'Brien²

¹Department of Molecular Simulation and Binding Kinetics, School of Pharmacy, Trinity College Dublin, Dublin, Ireland.

²Department of Residence Time and Drug-Target Interactions, School of Pharmacy, University College Cork, Cork, Ireland.

*Email: james.oleary@tcd.ie

ABSTRACT

Drug–target residence time is often discussed as though it were a scalar molecular attribute attached to a ligand–protein pair. Molecular simulation increasingly makes that interpretation difficult to sustain. Long-timescale and enhanced-sampling studies show that dissociation can proceed through multiple bound microstates, kinetically distinct exit routes, transition-state configurations, and protein–solvent reorganizations, while the inferred kinetic observable itself depends on sampling strategy, physical Hamiltonian, state definition, and first-passage estimator. This Perspective argues that simulated residence time should therefore be treated as an ensemble property: a macroscopic kinetic summary emerging from a conditional distribution of microscopic trajectories rather than from one privileged dissociation path. The proposal does not deny the utility of a scalar residence time. Instead, it requires that the scalar be interpreted together with evidence for state, pathway, and estimator convergence. The framework separates route discovery from route weighting, distinguishes sampling uncertainty from Hamiltonian uncertainty, and treats agreement with experiment as a validation gate rather than proof that a simulated microscopic mechanism is unique. This ensemble interpretation provides a stricter basis for comparing kinetic simulations, diagnosing failures, and deciding when a residence-time estimate is sufficiently supported for medicinal-chemistry use.

Keywords: Drug–target residence time, Ligand unbinding, Rare-event sampling, Molecular dynamics, Kinetic ensembles, Enhanced sampling

Received: 21 February 2025; **Revised:** 27 March 2025; **Accepted:** 29 March 2025

INTRODUCTION

Drug binding kinetics has become a distinct design variable in computational medicinal chemistry because equilibrium affinity and dynamical persistence answer different questions. Molecular simulation methods can interrogate association and dissociation processes that are invisible to a static structure or a single binding free energy, but their kinetic predictions require explicit treatment of timescale separation, rare events, and dynamical assumptions [1]. Residence time is conventionally connected to the dissociation rate constant, yet its pharmacological value remains target-, assay-, and exposure-dependent rather than universally beneficial. At the same time, the computational toolbox for estimating binding and unbinding kinetics now spans direct molecular dynamics, enhanced sampling, path-based methods, weighted ensembles, Markovian state approaches, and multiscale reconstruction [2, 3].

A simulated trajectory may reveal one escape event, but experimental residence time is a population-level kinetic observable. Equivalence requires representative initial states, adequate pathway coverage, a suitable Hamiltonian, and an estimator matched to the experimental quantity. Milestoning makes this distinction explicit by reconstructing long-time behavior from ensembles of shorter transitions [4].

This Perspective proposes that a scalar residence time should be treated as the compressed outcome of a conditional trajectory ensemble whose composition must be demonstrated rather than assumed. The following sections examine bound-state heterogeneity, exit-pathway multiplicity, transition ensembles, reorganization, and rare-event sampling before translating them into reporting rules.

Residence time beyond a single dissociation path

Direct rare-event simulations provide a strong existence proof that one ligand–target pair can dissociate through more than one microscopic history. In a weighted-ensemble study of TPPU dissociation from soluble epoxide hydrolase, unbinding occurred through a network of interconnected states and multiple routes rather than through a unique geometric exit [5]. The significance of that result is not that every system must have the same network topology. It is that a single macroscopic dissociation timescale can be consistent with a heterogeneous family of microscopic first-passage trajectories.

Energetic barriers alone are insufficient to define those trajectories. Atomistic analysis of multisecond dissociation showed that dynamical friction, hydration, and protein motion can influence effective crossing times [6]. A pathway is therefore a dynamical object, not merely a line through a free-energy surface.

Accelerated approaches reinforce this view while introducing an important boundary. τ -random acceleration molecular dynamics has recovered useful relative residence-time trends and mechanistic differences in G-protein-coupled receptor systems [7]. Such trajectories can expose recurrent exit motifs, transient interactions, and receptor-dependent detours, but the accelerated clock is not automatically the physical dissociation clock. Mechanistic route evidence and quantitative kinetic evidence must therefore remain analytically separate.

Pathway multiplicity also depends on how trajectories are represented. Community-detection approaches can partition unbinding trajectories into distinct route families, showing that heterogeneous egress behavior can be recovered at the ensemble level [8]. Yet a route label is still an analytical construct: different coordinates, contact definitions, or clustering schemes can partition the same trajectory set differently. Efficient path-search methods add another layer. SILCS-Pathway can enumerate and cluster plausible ligand dissociation routes from precomputed interaction fields, but route discovery is deliberately separated from subsequent kinetic prediction [9]. A possible path is not necessarily a probable path, and a probable path is not necessarily rate controlling.

Under the proposed ensemble view, the task is not merely to identify an unbinding pathway but to define kinetically relevant pathway support. A dominant route may emerge, but several routes may coexist or exchange importance with receptor state. The object of interest is the weighted first-passage ensemble, not the most visually intuitive trajectory.

Bound-state microstates

Pathway heterogeneity begins before dissociation. A ligand that is nominally “bound” may occupy a family of contact patterns and receptor conformations that share the same crystallographic pose at coarse resolution but differ kinetically. In HSP90, contact-map fingerprints derived from unbinding trajectories separated mechanistically distinct escape behaviors across multiple chemical series, while the same study also documented cases in which relative residence times were underestimated [10]. Those exceptions show that a bound-state fingerprint can be informative without being sufficient.

A related pattern appears in muscarinic receptor simulations. For tiotropium and related antagonists dissociating from the M3 receptor, path refinement and machine-learning analysis identified a structural bottleneck involving receptor extracellular-loop and transmembrane interactions around the transition region [11]. The evidence supports a receptor–ligand coupled bottleneck in that system, not a universal GPCR rule. It nevertheless illustrates why the starting bound ensemble matters: alternative local interaction networks can place otherwise similar ligands at different kinetic distances from the same bottleneck.

For flexible peptide ligands, extended metadynamics of class A GPCR systems resolved multiple bound configurations and alternative sites [12]. A free-energy basin is not automatically a distinct kinetic state, but one bound structure cannot be assumed to represent all dissociation-competent starting conditions.

Receptor context can change those conditions directly. In M3-receptor simulations, mutation-dependent interaction changes altered metastable behavior and the simulated unbinding process of tiotropium [13]. This

undermines the idea that residence time is an immutable property of the ligand alone. More cautiously, it supports the interpretation that the observed dissociation kinetics belong to a ligand–target–state system under a specified physical and experimental context.

The proposed theory therefore treats the bound state as a distribution over kinetically relevant microstates. If changing their populations changes route usage or first-passage statistics, the initial ensemble is part of the kinetic observable.

Exit-pathway multiplicity

Repeated accelerated trajectories provide one practical way to examine the distribution of escape routes. In a large HSP90 ligand series, τ RAMD enabled relative residence-time estimation while linking transient contacts and steric features to escape behavior [14]. The method demonstrates that repeated egress events can contain useful kinetic information, but the acceleration protocol prevents a literal identification of accelerated exit time with physical residence time.

Enhanced sampling also shows why one-dimensional exit coordinates can fail. Temperature-accelerated sliced sampling of avibactam dissociation exposed transverse protein and solvent bottlenecks [15], while funnel metadynamics examined ligand-class-dependent routes for methionine aminopeptidase-II inhibitors [16]. Multiple routes remain informative only after their statistical and physical relevance is tested.

At larger scale, route-resolved information can become predictive. A 2025 SILCS-Kinetics study used free-energy profiles from ligand dissociation pathways together with molecular properties to model experimental off-rates across 329 ligands and 13 proteins [17]. This demonstrates that multi-path descriptors can support scalable kinetic prediction. It does not show that a machine-learning weight is identical to the microscopic probability of a pathway. Predictive sufficiency and mechanistic uniqueness remain different claims.

Simulation can produce routes, energetic profiles, accelerated passage times, or learned descriptors, whereas experiment usually supplies an off-rate or residence time. The following proposed table keeps those evidential layers separate.

Table 1. Reconciling Exit-Route Evidence Before Assigning Ensemble Weight

Exit-route evidence pattern	Context	Observation and Interpretation and Boundary and Validation
Scientific anchor	Molecular or state context	Simulation observable What the pattern can justify What it cannot justify Discriminating test Residence-time ensemble interpretation
One route repeatedly dominates across independently initialized trajectories	Bound microstates remain similar and route topology is stable	Simulation observable Recurrent contact/path sequence and concentrated first-passage distribution What the pattern can justify A dominant route hypothesis What it cannot justify Proof that minor routes are absent Discriminating test before ensemble weighting Increase starting-state diversity and test route stability to additional sampling Residence-time ensemble interpretation Dominant contribution is plausible if route occupancy remains stable
Several geometrically distinct routes recur	Multiple exit corridors or transient interaction networks are sampled	Simulation observable Path clusters or route-specific free-energy profiles What the pattern can justify Genuine pathway multiplicity What it cannot justify Equal kinetic importance of every route Discriminating test before ensemble weighting Estimate route-specific first-passage support and sensitivity to representation Residence-time ensemble interpretation Residence time should be integrated over supported route families

Route identity changes with receptor or ligand microstate	Alternative starting conformations alter access to exits	Simulation observable State-conditioned path frequencies What the pattern can justify State-dependent pathway selection What it cannot justify A ligand-intrinsic route independent of receptor state Discriminating test before ensemble weighting Reweight or resample starting microstates and compare route populations Residence-time ensemble interpretation Initial-state distribution becomes part of the kinetic ensemble
A path is found by geometric or energetic search but rarely appears dynamically	Candidate corridor exists without repeated event support	Simulation observable Low-cost path or minimum-cost trajectory What the pattern can justify Physical plausibility of an exit corridor What it cannot justify Kinetic relevance Discriminating test before ensemble weighting Test with dynamical trajectories initiated from multiple nearby states Residence-time ensemble interpretation Candidate path remains unweighted until dynamic support exists
A route predicts experimental ranking despite ambiguous microscopic interpretation	Aggregate path descriptors correlate with measured off-rates	Simulation observable Relative kinetic score or learned prediction What the pattern can justify Decision-useful predictive signal within the tested domain What it cannot justify Unique microscopic mechanism Discriminating test before ensemble weighting External-series validation plus perturbations that distinguish competing paths Residence-time ensemble interpretation Predictive value can be retained while mechanism remains unresolved
Accelerated trajectories give stable route ordering but distorted absolute times	Forced or biased dynamics reproducibly separate compounds	Simulation observable Accelerated escape-time distribution What the pattern can justify Relative kinetic ordering What it cannot justify Direct physical residence time Discriminating test before ensemble weighting Compare against experimental kinetics or an absolute-rate method Residence-time ensemble interpretation Route evidence is usable, but physical time requires a separate calibration layer

Transition-state ensembles

If multiple bound states and exit routes exist, one transition-state structure may be inadequate. Free-energy calculations combined with machine learning show that dissociation-region features can contain kinetic information [18], but feature importance does not prove causal transition-state physics or uniquely define a dividing surface.

State-learning methods provide a complementary perspective. Enhanced sampling guided by a state predictive information bottleneck has been used to estimate ligand residence times across several benchmark complexes [19]. The methodological implication is that useful kinetic coordinates may need to encode information about metastable-state evolution rather than one intuitive geometric distance. Yet a learned coordinate remains representation dependent; its success does not establish that the learned variable is the unique mechanistic reaction coordinate.

The proposed interpretation therefore uses a transition-state ensemble when the evidence demands it. Structurally distinct configurations may share kinetic function or contribute through route-specific bottlenecks. A representative structure can remain useful, but should not be mistaken for the full statistical object.

Protein and solvent reorganization

Ligand Gaussian accelerated molecular dynamics enhances binding and unbinding while permitting reweighting of thermodynamic and kinetic quantities from recurrent events in benchmark systems [20]. Repeated transitions also reveal how starting configurations, intermediates, and reorganizations contribute to first-passage distributions.

Protein conformational gating can become especially important when the binding pocket is not continuously accessible. LiGaMD2 introduced additional boosting of pocket-residue degrees of freedom and recovered repeated binding and dissociation in closed-pocket T4 lysozyme systems, with calculated kinetic and thermodynamic quantities compared against available reference data [21]. In such systems, pocket opening is not external “noise” around ligand motion; it is part of the sequence of states that makes dissociation possible.

LiGaMD3 extended this logic to systems requiring broader ligand and protein conformational sampling [22]. Together with evidence for state-dependent friction and hydration [6], these studies support keeping barrier, solvent, friction, pocket-gating, and ligand-reorganization effects analytically distinct.

The figure below therefore treats protein and solvent reorganization as a structured support field that conditions pathway access and kinetic inference rather than as a decorative background around a ligand trajectory.




Figure 1. Reorganization Support Determines Which Dissociation Histories Become Kinetically Accessible

Sampling rare dissociation events

Rare events are difficult not only because they are slow, but because computational strategies alter the mapping between generated trajectories and physical time. Targeted molecular dynamics produced a useful residence-time ranking for CDK8 inhibitors [23], but that score should not be interpreted as an absolute residence time.

Even methods designed for kinetic reweighting retain statistical assumptions. Short-time infrequent metadynamics was developed in response to limitations of conventional infrequent-metadynamics inference, including sensitivity to collective-variable quality, bias-deposition behavior, and assumptions about first-passage statistics [24]. This is a different failure mode from insufficient event count. More trajectories cannot repair an estimator whose probabilistic assumptions are inappropriate for the observed process.

A second uncertainty axis lies in the physical Hamiltonian. QMrebind showed that local force-field reparameterization can materially alter predicted kinetics, while multiscale kinase calculations exposed sensitivity to protonation, charge assignment, and difficult chemical parameterization [25, 26]. Sampling uncertainty asks whether the relevant ensemble was explored; Hamiltonian uncertainty asks whether that ensemble belongs to an adequate physical model.

For practical use, rare-event failures should therefore be diagnosed by mechanism rather than placed under a single label of “simulation error.” The proposed triage below distinguishes misleading patterns that can look similar numerically but demand different corrective tests.

Table 2. Failure-Mode Triage for Rare-Event Residence-Time Calculations

Likely uncertainty source Competing diagnosis Residual consequence		Observable signature Corrective test	
Misleading pattern Residence-time estimate shifts when additional independent runs are added			
Likely uncertainty source Insufficient event or pathway sampling Competing diagnosis Hamiltonian error producing apparently unstable routes Residual consequence after mitigation Even converged sampling may remain wrong if the physical model is inaccurate		Observable signature Broad or unstable first-passage distribution; route frequencies continue to change Corrective test Extend independent sampling and test state/path occupancy convergence	
Misleading pattern Stable accelerated exit times disagree with experiment			
Likely uncertainty source Mapping from driven time to physical kinetics is invalid Competing diagnosis Incorrect protonation, charges, or force field Residual consequence after mitigation Relative ranking may remain useful even when absolute interpretation is rejected		Observable signature Reproducible ranking but poor absolute timescale Corrective test Compare with an independently calibrated absolute-rate method and perturb the physical model	
Misleading pattern Free-energy barrier appears converged but inferred rate remains unstable			
Likely uncertainty source Kinetic estimator or friction treatment is inadequate Competing diagnosis Missing slow protein or solvent coordinate Residual consequence after mitigation A converged energetic landscape does not guarantee converged kinetics		Observable signature Similar barrier profile with changing rate estimate Corrective test Test alternative state definitions, friction treatment, or first-passage model	
Misleading pattern Different collective variables produce different pathways			
Likely uncertainty source Representation incompleteness Competing diagnosis Genuine multi-route kinetics Residual consequence after mitigation Mechanistic uniqueness should remain unresolved if alternatives persist		Observable signature Path topology or transition ordering depends on coordinate choice Corrective test Use state/pathway analyses not tied to one coordinate and test recurrence	
Misleading pattern One chemical subset is systematically mispredicted			
Likely uncertainty source Hamiltonian or chemical-state mismatch Competing diagnosis Sampling deficit specific to that subset Residual consequence after mitigation Domain restrictions should be reported even if aggregate performance is good		Observable signature Errors cluster by scaffold, protonation state, halogenation, or interaction motif Corrective test Parameter sensitivity, protonation/tautomer checks, and targeted resampling	
Misleading pattern Experimental rank ordering is reproduced but route interpretation varies across methods			
Likely uncertainty source Predictive–mechanistic discordance Competing diagnosis Methods are observing different valid route components Residual consequence after mitigation Prediction can remain decision-useful while mechanism is explicitly uncertain		Observable signature Similar kinetic ordering with different microscopic narratives Corrective test Apply perturbations designed to discriminate route hypotheses	

The same logic can be represented as a trajectory-to-inference architecture with feedback. The purpose is to prevent a successful scalar output from obscuring failure at an earlier level.




Figure 2. Rare-Event Trajectories Require Recursive Validation Before They Become Residence-Time Evidence

Relating simulated kinetics to experimental residence time

The value of a simulated residence time ultimately depends on what experimental quantity it is intended to reproduce. Computational studies do not all generate the same observable. Some produce accelerated escape times, others relative kinetic scores, first-passage distributions, reconstructed k_{off} values, or absolute residence-time estimates. Treating these outputs as interchangeable obscures both their strengths and their limitations.

The distinction becomes especially important when simulation enters medicinal-chemistry decision making. ModBind illustrates how a rapid simulation-based predictor can be evaluated retrospectively against experimentally characterized compounds and then used prospectively to prioritize additional synthesis within a defined discovery program [27]. The relevant lesson is not that one computational score has universal transferability. It is that kinetic predictions become decision-relevant only after their relationship to measured kinetics has been established within an appropriate target, chemical series, and assay context.

More physics-explicit calculations are also approaching absolute timescale prediction. A two-stage workflow combining random-acceleration pathway exploration with infrequent metadynamics was recently evaluated across 29 protein–ligand complexes from five targets and reported an RMSE of 1.22 and an R^2 of 0.80 for log-transformed residence times [28]. This constitutes substantial evidence that absolute residence-time estimation can be feasible across multiple systems. It does not establish universal accuracy, nor does numerical agreement alone demonstrate that the simulated route distribution is the unique physical mechanism.

The distinction between kinetic agreement and mechanistic agreement is central to the ensemble interpretation proposed here. Two computational workflows could yield comparable k_{off} values while assigning different weights to bound microstates or exit routes. Conversely, a method could reproduce experimentally plausible routes yet generate a biased absolute timescale. Experimental concordance should therefore be treated as a validation gate with at least three questions: is the computational observable commensurate with the experimental one; is agreement maintained across the intended chemical domain; and is

the microscopic interpretation stable to reasonable changes in sampling, state definition, and physical-model specification?

This leads to a proposed observable-alignment rule. Before simulation and experiment are declared concordant, the reported computational quantity should be identified explicitly as a relative rank, method-specific score, $\text{koffk}_{\text{off}}$, residence time, or another first-passage statistic. Experimental assay conditions and the computational initial-state definition should also be stated where they affect interpretation. Correlation between non-equivalent observables can still be useful for prioritization, but it should not be presented as direct validation of absolute kinetics.

An ensemble theory of residence time

The preceding evidence supports a theoretical reconstruction of simulated residence time as an ensemble property. The proposal begins from a simple distinction: the reported residence time is a scalar, whereas the molecular process generating it may be distributed across many microscopic histories. Direct weighted-ensemble simulations already demonstrate networked dissociation pathways in individual ligand–target systems [5]. Atomistic kinetic analysis additionally shows that effective dissociation times can depend on hydration, friction, and dynamical reorganization rather than only on a single energetic barrier [6].

Accordingly, the proposed ensemble consists of several conditionally coupled levels. The first is the bound-state ensemble, comprising kinetically distinguishable ligand–protein microstates. The second is the exit-route ensemble, comprising dynamically supported pathways accessible from those starting states. The third is the transition ensemble, comprising route-dependent configurations associated with kinetically important bottlenecks rather than necessarily one unique transition structure. The fourth is the reorganization ensemble, comprising protein, ligand, hydration, and local interaction-network changes that alter route accessibility or crossing dynamics.

These levels should not be interpreted as a universal fixed decomposition. Their relevance depends on the molecular system and on the simulation representation. Mutation-dependent changes in receptor microstates, for example, can alter simulated escape behavior [13], while state-learning approaches show that the coordinates needed for kinetic prediction may themselves be representation dependent [19]. The ensemble is therefore conditional on the preparation protocol, protonation and chemical state, Hamiltonian, initial-state distribution, pathway definition, and kinetic estimator.

A useful scalar residence time can still emerge from this complexity. The proposed theory does not replace $\text{koffk}_{\text{off}}$ or residence time with an unwieldy catalogue of trajectories. Instead, it changes what constitutes adequate support for that scalar. A residence-time estimate should be considered well supported when the macroscopic statistic is stable to reasonable expansion of the underlying state and pathway ensemble and when remaining uncertainty has been localized rather than hidden.

This distinction also prevents an important category error. Ensemble variability is not necessarily computational failure. Genuine microscopic heterogeneity may persist after convergence. The objective is therefore not to force every system into one dominant route, but to determine whether the calculated kinetic observable is stable despite route-level heterogeneity and whether the inferred route weights are sufficiently supported for the intended scientific claim.

Computational tests and reporting standards

An ensemble interpretation becomes useful only if it changes computational practice. The first proposed test is event and pathway convergence. Increasing simulation length alone is insufficient if newly sampled trajectories continue to reveal additional bound states or route families. Convergence should therefore be assessed against the quantities actually entering the kinetic estimate: event statistics, state occupancies, pathway distributions, or milestone transitions, depending on method [4].

Second, representation sensitivity should be distinguished from sampling convergence. If pathway assignments change substantially with reasonable contact definitions, collective variables, or clustering procedures, mechanistic uniqueness should not be asserted. Path-classification methods themselves demonstrate that route identity is partly dependent on how trajectory information is organized [8]. A stable residence-time estimate accompanied by unstable route labels may still support a kinetic prediction, but not a unique microscopic narrative. Third, Hamiltonian sensitivity should be examined separately. Changes in force-field parameters, protonation states, partial charges, or local quantum-mechanical treatment can alter predicted kinetics even when trajectory

sampling appears converged. Multiscale residence-time calculations have exposed such sensitivities directly [26]. Reporting only sampling uncertainty can therefore produce false confidence.

Fourth, the reported experimental comparison should identify its level. A method that ranks ligands correctly is performing a different task from one estimating absolute k_{off} . Targeted molecular dynamics, for example, can provide a useful residence-time ranking without converting its driven trajectory clock into physical time [23]. The manuscript therefore proposes that studies label explicitly whether validation concerns rank ordering, absolute kinetics, a calibrated surrogate, or mechanistic consistency.

Finally, discordant cases should be retained rather than removed from the narrative. The HSP90 scaled-dynamics study contained a subset for which residence times were underestimated [10]. Such cases are analytically valuable because they indicate where sampling, representation, or physical-model assumptions may fail. Reporting standards for simulated residence time should therefore include successful cases, failure cases, sensitivity analyses, the definition of the kinetic observable, the initial-state ensemble, route or state convergence evidence, and the limits of chemical- or target-space transfer.

The practical implication is a stricter evidential hierarchy. A visually plausible escape trajectory is pathway evidence. Recurrent trajectories add population evidence. A stable first-passage estimator adds kinetic evidence. Agreement with experimental kinetics adds external validation. Only when these layers are appropriately aligned should a residence-time prediction be used as strong mechanistic evidence or as a quantitative medicinal-chemistry decision variable.

CONCLUSION

Residence time is routinely expressed as a single number, but molecular simulation increasingly shows that the number can emerge from many microscopic histories. Bound-state heterogeneity, pathway multiplicity, transition-state variation, conformational gating, solvent reorganization, rare-event sampling, physical-model specification, and kinetic-estimator assumptions can all affect the mapping from a molecular trajectory to an inferred dissociation timescale.

Treating simulated residence time as an ensemble property preserves the usefulness of the scalar while making its evidential foundation explicit. The proposed interpretation separates path discovery from kinetic weighting, distinguishes sampling uncertainty from Hamiltonian uncertainty, and requires mechanistic claims to remain conditional on state and pathway convergence. It also prevents experimental agreement from being mistaken for proof that a particular microscopic route is uniquely correct.

For molecular simulation to contribute reliably to kinetic drug design, the central question should therefore shift from whether a calculation has produced a residence-time value to whether the ensemble supporting that value has been adequately defined, sampled, tested, and aligned with experiment. That shift provides a more rigorous basis for comparing methods, diagnosing failures, and deciding when simulated kinetics can legitimately inform medicinal chemistry.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. 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
2. Bernetti M, Masetti M, Rocchia W, Cavalli A. Kinetics of drug binding and residence time. *Annu Rev Phys Chem.* 2019;70:143-71. doi:10.1146/annurev-physchem-042018-052340
3. Wang J, Do HN, Koirala K, Miao Y. Predicting biomolecular binding kinetics: A review. *J Chem Theory Comput.* 2023;19(8):2135-48. doi:10.1021/acs.jctc.2c01085
4. Ojha AA, Votapka LW, Amaro RE. Advances and challenges in milestoning simulations for drug-target kinetics. *J Chem Theory Comput.* 2024;20(22):9759-69. doi:10.1021/acs.jctc.4c01108

5. Lotz SD, Dickson A. Unbiased molecular dynamics of 11 min timescale drug unbinding reveals transition state stabilizing interactions. *J Am Chem Soc.* 2018;140(2):618-28. doi:10.1021/jacs.7b08572
6. Wolf S, Lickert B, Bray S, Stock G. Multisecond ligand dissociation dynamics from atomistic simulations. *Nat Commun.* 2020;11(1):2918. doi:10.1038/s41467-020-16655-1
7. Kokh DB, Wade RC. G protein-coupled receptor-ligand dissociation rates and mechanisms from τ RAMD simulations. *J Chem Theory Comput.* 2021;17(10):6610-23. doi:10.1021/acs.jctc.1c00641
8. Tänzeli V, Jäger M, Wolf S. Learning protein-ligand unbinding pathways via single-parameter community detection. *J Chem Theory Comput.* 2024;20(12):5058-67. doi:10.1021/acs.jctc.4c00250
9. Yu W, Weber DJ, MacKerell AD Jr. Detection of putative ligand dissociation pathways in proteins using site-identification by ligand competitive saturation. *J Chem Inf Model.* 2025;65(6):3022-34. doi:10.1021/acs.jcim.4c01814
10. Bianciotto M, Gkeka P, Kokh DB, Wade RC, Minoux H. Contact map fingerprints of protein-ligand unbinding trajectories reveal mechanisms determining residence times computed from scaled molecular dynamics. *J Chem Theory Comput.* 2021;17(10):6522-35. doi:10.1021/acs.jctc.1c00453
11. Buigues PJ, Gehrke S, Badaoui M, Dudas B, Mandana G, Qi T, et al. Investigating the unbinding of muscarinic antagonists from the muscarinic 3 receptor. *J Chem Theory Comput.* 2023;19(15):5260-72. doi:10.1021/acs.jctc.3c00023
12. Calderón JC, Plut E, Keller M, Cabrele C, Reiser O, Gervasio FL, et al. Extended metadynamics protocol for binding/unbinding free energies of peptide ligands to class A G-protein-coupled receptors. *J Chem Inf Model.* 2024;64(1):205-18. doi:10.1021/acs.jcim.3c01574
13. Coricello A, Chiaravalle AL, Musgaard M, Tehan BG, Elisi GM, Bottegoni G. Adiabatic-bias molecular dynamics simulations reveal the impact of mutations on muscarinic antagonist unbinding kinetics. *J Chem Inf Model.* 2025;65(13):7129-42. doi:10.1021/acs.jcim.5c00601
14. 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
15. Tripathi S, Nair NN. Temperature accelerated sliced sampling to probe ligand dissociation from protein. *J Chem Inf Model.* 2023;63(16):5182-91. doi:10.1021/acs.jcim.3c00376
16. Rubina, Moin ST. Attempting well-tempered funnel metadynamics simulations for the evaluation of the binding kinetics of methionine aminopeptidase-II inhibitors. *J Chem Inf Model.* 2023;63(24):7729-43. doi:10.1021/acs.jcim.3c01130
17. Yu W, Kumar S, Zhao M, Weber DJ, MacKerell AD Jr. High-throughput ligand dissociation kinetics predictions using site identification by ligand competitive saturation. *J Chem Theory Comput.* 2025;21(9):4964-78. doi:10.1021/acs.jctc.5c00265
18. Badaoui M, Buigues PJ, Berta D, Mandana GM, Gu H, Földes T, et al. Combined free-energy calculation and machine learning methods for understanding ligand unbinding kinetics. *J Chem Theory Comput.* 2022;18(4):2543-55. doi:10.1021/acs.jctc.1c00924
19. Lee S, Wang D, Seeliger MA, Tiwary P. Calculating protein-ligand residence times through state predictive information bottleneck based enhanced sampling. *J Chem Theory Comput.* 2024;20(14):6341-49. doi:10.1021/acs.jctc.4c00503
20. Miao Y, Bhattarai A, Wang J. Ligand Gaussian accelerated molecular dynamics (LiGaMD): Characterization of ligand binding thermodynamics and kinetics. *J Chem Theory Comput.* 2020;16(9):5526-47. doi:10.1021/acs.jctc.0c00395
21. Wang J, Miao Y. Ligand Gaussian accelerated molecular dynamics 2 (LiGaMD2): Improved calculations of ligand binding thermodynamics and kinetics with closed protein pocket. *J Chem Theory Comput.* 2023;19(3):733-45. doi:10.1021/acs.jctc.2c01194
22. Wang J, Miao Y. Ligand Gaussian accelerated molecular dynamics 3 (LiGaMD3): Improved calculations of binding thermodynamics and kinetics of both small molecules and flexible peptides. *J Chem Theory Comput.* 2024;20(14):5829-41. doi:10.1021/acs.jctc.4c00502
23. Ziada S, Diharce J, Raimbaud E, Aci-Sèche S, Ducrot P, Bonnet P. Estimation of drug-target residence time by targeted molecular dynamics simulations. *J Chem Inf Model.* 2022;62(22):5536-49. doi:10.1021/acs.jcim.2c00852
24. Blumer O, Reuveni S, Hirshberg B. Short-time infrequent metadynamics for improved kinetics inference. *J Chem Theory Comput.* 2024;20(9):3484-91. doi:10.1021/acs.jctc.4c00170

25. Ojha AA, Votapka LW, Amaro RE. QMrebind: Incorporating quantum mechanical force field reparameterization at the ligand binding site for improved drug-target kinetics through milestoning simulations. *Chem Sci.* 2023;14(45):13159-75. doi:10.1039/D3SC04195F
26. Votapka LW, Ojha AA, Asada N, Amaro RE. Prediction of threonine-tyrosine kinase receptor-ligand unbinding kinetics with multiscale milestoning and metadynamics. *J Phys Chem Lett.* 2024;15(42):10473-8. doi:10.1021/acs.jpclett.4c02332
27. Sinko W, Mertz B, Shimizu T, Takahashi T, Terada Y, Kimura SR. ModBind, a rapid simulation-based predictor of ligand binding and off-rates. *J Chem Inf Model.* 2025;65(1):265-74. doi:10.1021/acs.jcim.4c01805
28. Smith Z, Branduardi D, Lupyan D, D'Arrigo G, Tiwary P, Wang L, et al. Toward automated physics-based absolute drug residence time predictions. *J Chem Inf Model.* 2025;65(24):13360-73. doi:10.1021/acs.jcim.5c01832