Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

Beyond Static Protein–Ligand Fit: A Conceptual Architecture for Incorporating Conformational Ensembles, Binding Kinetics, and Cellular Context into Artificial Intelligence–Assisted Molecular Modeling and Drug Design

Michael Osei[1], Ingrid Larsen[2]*, Rajesh Menon[3], Julia Berg[4]

1. ***Department of Structural Biology and AI-Assisted Drug Design, University of Ghana Business School, Accra, Ghana.***
2. ***Department of Molecular Modeling and Computational Biophysics, BI Norwegian Business School, Oslo, Norway.***
3. ***Department of Pharmacoinformatics and Molecular Dynamics, Indian Institute of Management Ahmedabad, Ahmedabad, India.***
4. ***Department of Drug Design and Molecular Pharmacology, University of Vienna, Vienna, Austria.***

ARTICLE INFO

Received:
01 May 2026
Received in revised form:
30 August 2026
Accepted:
02 September 2026
Available online:
28 October 2026

Keywords: Conformational ensembles, Binding kinetics, Protein–ligand interactions, Molecular dynamics, Artificial intelligence, Target engagement

ABSTRACT

Artificial intelligence has markedly expanded the ability to predict protein structures, protein–ligand complexes, binding poses, and molecular properties, yet much of AI-assisted molecular modeling still treats binding as a predominantly static structural problem. This assumption is limiting because molecular action emerges from distributions of conformational states, transitions among those states, time-dependent association and dissociation, and biological environments that can alter target accessibility, engagement, and selectivity. This conceptual architecture article examines how these information layers should be incorporated without replacing interpretable molecular modeling with untestable complexity. The analysis distinguishes four related but non-equivalent representational problems: geometrical fit, conformational-state occupancy, binding kinetics, and cellular context. It argues that richer modeling is useful only when each added layer changes a testable observation, comparison, or drug-design decision and when validation is matched to the level of inference being made. Conformational ensembles can reveal alternative pockets and ligand-dependent states, but they require experimental or physics-informed constraints. Binding kinetics can distinguish compounds with similar equilibrium affinity, but residence time is not inherently advantageous. Cellular measurements can expose target engagement and polypharmacology that isolated complexes cannot represent, yet cellular engagement does not establish therapeutic efficacy. A dynamic molecular-modeling architecture should therefore preserve these layers as distinct sources of evidence while connecting them selectively through falsifiable cross-level relationships. The resulting approach reframes AI-assisted molecular modeling from prediction of a best-fitting complex toward conditional inference about molecular behavior across structure, time, and biological context.

*This is an **open-access** article distributed under the terms of the Creative Commons Attribution-Non Commercial-Share Alike 4.0 License, which allows others to remix, and build upon the work non commercially.*

To Cite This Article: Osei M, Larsen I, Menon R, Berg J. Beyond Static Protein–Ligand Fit: A Conceptual Architecture for Incorporating Conformational Ensembles, Binding Kinetics, and Cellular Context into Artificial Intelligence–Assisted Molecular Modeling and Drug Design. Pharmacophore. 2026;17(5):17-25. https://doi.org/10.51847/WotZIrvAZL

Introduction

Artificial intelligence has altered the scale and speed at which molecular structures, interactions, and candidate compounds can be generated or prioritized. The central difficulty is no longer simply whether a model can produce a plausible structure or rank compounds against a benchmark. For drug discovery, the harder question is whether the representation being optimized corresponds to the molecular behavior that matters for an experimental or translational decision. Recent assessments of AI-enabled drug discovery therefore distinguish technical model performance from evidence that a model materially improves decision quality across the discovery process [1]. This distinction is especially consequential in molecular modeling because

Corresponding Author: Ingrid Larsen; Department of Molecular Modeling and Computational Biophysics, BI Norwegian Business School, Oslo, Norway. E-mail: ingrid.larsen@bi.no.

apparently precise predictions may still omit biological dimensions that determine whether a ligand binds in the relevant state, for the relevant duration, and in the relevant cellular environment.

The rapid advance of biomolecular structure prediction illustrates both the opportunity and the representational problem. AlphaFold 3 extended learned structural prediction beyond individual proteins toward complexes involving proteins, nucleic acids, small molecules, ions, and modified residues, substantially broadening the range of molecular systems that can be represented within a common predictive framework [2]. RoseTTAFold All-Atom similarly showed that unified all-atom representations can model heterogeneous biomolecular assemblies rather than treating proteins and ligands as fundamentally separate computational objects [3]. These developments reduce important structural uncertainties, particularly when experimentally resolved complexes are unavailable. They do not, however, make a predicted complex equivalent to the complete molecular interaction.

The relevant distinction is between predicting a structurally plausible configuration and representing the physical and biological process through which molecular action occurs. A protein–ligand complex observed or predicted at one coordinate set contains no explicit statement about the probability of alternative receptor states, the transitions connecting those states, the rate at which a ligand associates or dissociates, or whether the same interaction persists in a membrane, organelle, multiprotein environment, or intact cell. Even equilibrium affinity is an aggregate thermodynamic quantity and cannot, by itself, specify how rapidly the bound state is formed or lost. Conversely, a more dynamic representation is not automatically better. An unconstrained ensemble can multiply possibilities without increasing predictive discrimination, and highly detailed simulations can introduce sampling and force-field dependencies that are less transparent than the original static model.

The problem addressed here is therefore not whether static molecular modeling should be abandoned. Static structures remain indispensable for recognizing interaction geometries, interpreting chemical complementarity, screening chemical space, and formulating mechanistic hypotheses. The problem is when a static representation is treated as sufficient for an inference that actually depends on state heterogeneity, temporal behavior, or biological context. The article develops this distinction in stages. It first identifies pharmacologically relevant information that is suppressed when protein–ligand binding is represented as a single optimized configuration. It then examines how conformational, kinetic, and cellular information can be integrated while preserving their different evidentiary meanings. The subsequent architecture treats these layers as connected but non-interchangeable and asks what form of validation is required before information can legitimately propagate from one layer to another.

Beyond Static Protein–Ligand Representations

A static protein–ligand representation usually assigns privileged status to one receptor conformation, one ligand pose, or one optimized complex. This simplification can be productive when the underlying system is sufficiently rigid or when the decision being made is narrowly structural. It becomes problematic when ligand binding reorganizes the receptor or when different ligands preferentially stabilize different members of an accessible conformational population. DynamicBind demonstrates the practical significance of this problem by predicting ligand-specific complex structures while allowing substantial protein conformational change rather than assuming that docking must occur into a fixed receptor geometry [4]. The key implication is not that every docking calculation requires unrestricted protein flexibility, but that the receptor structure supplied to a model is itself part of the hypothesis being tested.

State-specific generative modeling makes the distinction more explicit. Deep generative approaches have been used to predict protein structures conditioned on ligand state, thereby acknowledging that the relevant receptor configuration may depend on the interacting molecule [5]. Such methods alter the conceptual unit of prediction. Instead of asking only which ligand fits a protein structure, they permit the question of which protein–ligand configuration is compatible with a particular molecular state. Yet generated structural diversity should not automatically be interpreted as a calibrated conformational ensemble. Analyses of AlphaFold-derived alternative structures indicate that sampled conformations can be strongly shaped by training-set information and representation history, with rare biologically relevant states remaining difficult to recover reliably [6]. Computational multiplicity is therefore not equivalent to physical heterogeneity.

The validation problem becomes more fundamental when ensembles, rather than individual conformations, are the intended representation. Experimental measurements typically constrain projections of an ensemble rather than observing every populated microstate directly. Geometrical analyses of protein conformational ensembles show that multiple structural distributions may be compatible with a given set of experimental observables, making ensemble reconstruction an underdetermined problem unless sufficiently informative constraints are available [7]. Physics-informed approaches can narrow this uncertainty. Energetic frustration analysis combined with structure-prediction methods, for example, has been used to generate alternative protein conformations and motions that are more closely tied to features of an underlying energy landscape [8]. Nevertheless, a predicted pathway or collection of structures should be interpreted as a testable model of conformational accessibility rather than direct evidence of transition probabilities or kinetic rates.

The consequences for ligand discovery are substantial because alternative states can expose interaction opportunities that a dominant or experimentally resolved structure does not show. Cryptic pockets provide a clear example. PocketMiner demonstrated that graph-based learning informed by molecular-dynamics-derived behavior can identify sites that become accessible through protein motion even when they are not apparent in a single static structure [9]. Conversely, the practical success of docking against AlphaFold2-derived structures shows that a single predicted structure can still be useful for

prospective ligand discovery under favorable conditions [10]. These findings are not contradictory. They imply that static structures can be decision-useful without being ontologically complete descriptions of molecular behavior.

Pose evaluation illustrates a second failure of static reduction: structural proximity to a reference geometry is not sufficient if the predicted interaction is physically implausible. PoseBusters showed that AI-based docking methods can produce poses that satisfy conventional pose metrics yet violate basic chemical or physical constraints, particularly when generalizing beyond familiar training distributions [11]. The relevant lesson is that a protein–ligand representation must preserve more than coordinates. It must also preserve the constraints that determine whether those coordinates represent a chemically plausible bound state and, where necessary, whether that state belongs to a realistic accessible ensemble.

These distinctions define four different questions that are often collapsed into a single notion of "binding." The first concerns whether a ligand can form a plausible interaction geometry with a target. The second concerns which target and ligand conformations are populated and accessible. The third concerns how rapidly binding states are entered and exited. The fourth concerns whether the interaction remains observable and selective in its biological environment. A static complex directly addresses only the first of these. It can provide partial information relevant to the others, but inference across them requires additional evidence.

The analytical consequence is summarized in **Table 1**. The table treats the missing information not as a generic list of model limitations, but as distinct molecular phenomena whose omission changes what conclusions can legitimately be drawn from an AI-generated protein–ligand complex.

Table 1. Dynamic information that a single protein–ligand complex does not resolve

Molecular phenomenon	What a static representation can show	Additional representation required	Observable evidence that can constrain the added layer	Common interpretive error if omitted
Ligand-dependent receptor rearrangement	Local geometry of one complex	Alternative receptor states conditioned on ligand binding	Alternative experimental structures, state-sensitive spectroscopy, restrained or physics-informed simulations	Treating one receptor conformation as universally competent for all ligands
Conformational heterogeneity	One selected member of a structural population	Weighted or otherwise constrained ensemble of accessible states	Ensemble-sensitive structural observables and independently validated state distributions	Treating generated structural diversity as a calibrated physical ensemble
Cryptic-pocket accessibility	Pockets visible in the represented structure	State-dependent pocket opening and closure	Dynamic pocket observations, simulation-supported opening events, ligandability experiments	Concluding that an absent pocket is biologically inaccessible
Ligand conformational variability	One bound ligand conformation	Ligand conformer population and receptor-dependent conformational selection	Solution-state conformational measurements, energetics, bound-state comparisons	Assuming the docked ligand pose represents its relevant conformational landscape
Physical plausibility of a predicted pose	Coordinate agreement or geometric complementarity	Chemical, steric, energetic, and interaction-consistency constraints	Orthogonal pose checks, energetic evaluation, experimental structural confirmation	Equating low pose error with a physically valid molecular complex
Transition between molecular states	End-point structures	Pathway or transition representation	Kinetic experiments, enhanced sampling, mechanistic simulations	Inferring transition mechanism from structural endpoints alone

Integrating Conformational, Kinetic and Cellular Information

Moving beyond a static complex requires integration, but integration should not mean collapsing different forms of evidence into one composite score. Conformational states, equilibrium free energies, binding rates, and cellular target engagement answer different questions. Their value lies partly in the fact that they can disagree. A compound may occupy a favorable docked pose while rarely encountering the required receptor state; two compounds may have similar equilibrium affinity but markedly different dissociation rates; a molecule that binds purified protein may fail to engage the same target in cells. A dynamic architecture must therefore preserve distinctions among these layers before attempting to connect them.

The first requirement is richer conformational information on both sides of the interaction. MISATO combines experimentally derived protein–ligand structures with molecular-dynamics trajectories and quantum-mechanical information, providing machine-learning-ready data that retain temporal adaptability beyond a single crystallographic complex [12]. Complementary work on ligand representation shows that explicit profiling of molecular conformational space can improve learned molecular representations across downstream prediction tasks [13]. Together, these approaches imply that "structure" should not be represented solely as one receptor coordinate set interacting with one ligand conformer. Receptor-state diversity and ligand-state diversity are analytically distinct sources of variation, and either can alter the interaction landscape.

Physical information remains necessary even when representation learning is powerful. Interaction-scoring approaches that incorporate physical prior knowledge and data augmentation have shown improved generalization relative to models that rely

more heavily on learned statistical associations [14]. This does not establish that physically informed scores reproduce the entire thermodynamic system. Rather, it shows why an architecture should retain an identifiable physical layer instead of allowing a learned score to become an uninterpretable surrogate for every property of interest.

Equilibrium thermodynamics supplies one such layer. Alchemical free-energy calculations estimate changes in binding free energy under explicitly defined thermodynamic transformations and have become increasingly integrated into medicinal-chemistry workflows [15]. Automated absolute binding free-energy workflows extend this strategy toward scalable prospective use [16]. Enhanced-sampling approaches can further improve the precision of relative binding free-energy estimates when slow degrees of freedom would otherwise remain insufficiently explored [17]. Yet precision in repeated calculations does not eliminate systematic error from force fields, initial states, protonation choices, or inadequate sampling of relevant conformational changes. Thermodynamic predictions therefore require their own validation and should not be used as indirect substitutes for kinetic information.

The biological environment can also change the assumptions under which thermodynamic calculations are interpreted. Membrane proteins illustrate this directly because ligand binding occurs in an anisotropic, compositionally complex environment that differs from soluble-protein conditions. Alchemical affinity calculations applied to membrane targets can be informative, but their validity depends on treatment of membrane composition, receptor state, solvent exposure, and sampling conditions [18]. Cellular context extends the same principle further: an interaction observed for an isolated macromolecule may be altered by localization, competition, post-translational modification, partner proteins, permeability, metabolism, or target abundance.

Binding kinetics introduces a separate temporal dimension. Structural and biophysical studies have shown that ligand binding can proceed through different mixtures of conformational selection and induced fit, and that stabilization of transition states can change which pathway dominates [19]. This means that structurally similar endpoints need not imply identical routes to the bound state. The distinction matters pharmacologically because equilibrium affinity depends on the ratio of association and dissociation behavior, while the individual rate constants determine how rapidly occupancy is established and lost. Reviews of drug–target residence time therefore emphasize that kinetic behavior can contribute information not recoverable from equilibrium affinity alone [20].

The translational interpretation of kinetics must remain conditional. A long residence time may sustain target occupancy under declining free-drug concentrations, but it can also prolong undesirable interactions, complicate reversibility, or become irrelevant when pharmacokinetics, target turnover, compartmentalization, or downstream signaling dominate observed response. The translational value of ligand–receptor kinetics therefore depends on the relationship among molecular binding, exposure, target biology, and the time scale of the pharmacological process being studied [21]. Kinetic information should be treated as an additional coordinate of molecular behavior rather than as a universally superior optimization objective.

The integration problem can consequently be expressed as a sequence of matched questions rather than a hierarchy in which each later layer "proves" the earlier one. Conformational modeling asks which molecular states are accessible and relevant. Thermodynamic analysis asks how favorable defined bound and unbound states are at equilibrium. Kinetic analysis asks how rapidly those states are entered and exited. Cellular testing asks whether the predicted interaction survives the biological constraints surrounding the target. Each layer can revise interpretation of the others, but none can be substituted for another merely because it produces a quantitative output.

Table 2 translates these distinctions into a layer-matched validation logic. The ordering reflects increasing biological context, not increasing certainty: evidence can fail or become ambiguous at any level, and successful validation at one level does not automatically authorize inference at the next.

Table 2. Layer-matched validation requirements for dynamic AI-assisted protein–ligand models

Information layer	Primary modeling question	Appropriate validation target	What successful validation can support	What it still cannot establish
Structural geometry	Is the predicted complex chemically and sterically plausible?	Experimental complex geometry or orthogonal physical plausibility checks	Credible local interaction hypothesis	Population of alternative states, kinetic behavior, cellular engagement
Conformational ensemble	Which receptor and ligand states are accessible and materially populated?	Ensemble-sensitive experimental observables and independently constrained simulations	State-dependent interaction hypotheses and identification of alternative binding opportunities	Unique transition mechanism or binding rates
Equilibrium thermodynamics	How favorable is binding under a defined molecular state and environment?	Prospective affinity measurements under matched conditions	Relative or absolute thermodynamic ranking within the tested domain	Association rate, dissociation rate, residence time, cellular selectivity
Binding kinetics	How rapidly are complexes formed and lost?	Direct or appropriately inferred association and dissociation measurements	Time-dependent occupancy behavior under the tested biochemical conditions	Pharmacological duration in vivo without exposure and system context

Cellular target engagement	Does the compound engage the intended target in the cellular environment?	Direct engagement, binding-site, or competition measurements in relevant cells	Persistence of target interaction under cellular constraints	Therapeutic efficacy, safety, or patient-level benefit
Cellular selectivity and interaction landscape	What intended and unintended molecular interactions occur in cells?	Proteome-scale or targeted selectivity measurements	Context-specific target and off-target landscape	Clinical consequence of each detected interaction
Translational relevance	Does the combination of molecular properties improve a drug-design decision?	Prospective medicinal-chemistry or pharmacological testing linked to a defined decision	Decision-specific usefulness of the integrated model	General superiority of the architecture across targets, modalities, or diseases

The central implication is that dynamic molecular modeling should not be defined by maximal mechanistic detail. It should be defined by retaining the minimum additional information necessary to prevent an important inference error. If ligand ranking changes when receptor state changes, state information is necessary. If compounds with similar affinity behave differently because of dissociation kinetics, temporal information is necessary. If binding measured in purified systems does not persist in cells, cellular evidence is necessary. When none of these distinctions changes the intended decision, a simpler static model may remain adequate. The architecture developed next therefore treats complexity as conditional and falsifiable rather than intrinsically desirable.

Proposed Dynamic Molecular-Modeling Architecture

A dynamic molecular model should preserve distinctions among conformational state, equilibrium thermodynamics, binding kinetics, and cellular realization rather than compressing them into a single prediction. Free-energy and kinetic calculations already illustrate why this separation matters: thermodynamic calculations estimate the relative favorability of defined molecular states, whereas kinetic approaches address the rates and pathways connecting them [22]. These quantities can inform one another, but they are not interchangeable. An architecture that represents them separately can therefore ask whether a ligand is predicted to bind favorably, whether the relevant receptor states are sufficiently accessible, how rapidly the complex forms and dissociates, and whether the interaction remains detectable under cellular constraints.

The first layer of the architecture is a conformational-state field. Instead of assigning one receptor structure the status of ground truth, the model represents a constrained set of receptor and ligand states whose inclusion is justified by experimental structures, ensemble-sensitive measurements, physically informed simulations, or reproducibly generated hypotheses. State identity alone is insufficient: the model must record why each state is included and what evidence could falsify its relevance. This prevents an AI-generated collection of conformers from becoming an unconstrained catalogue in which every downstream prediction can be rationalized after the fact.

A second layer describes the energetic and kinetic relations among states. Evidence from integrated studies of HSP90 inhibitors shows how structural information, molecular dynamics, quantitative structure–kinetics relationships, and experimental kinetic measurements can jointly connect ligand-specific conformational behavior to dissociation kinetics [23]. The importance of such work is methodological rather than target-specific. It demonstrates that a plausible architecture should allow structural states and kinetic observables to constrain one another while preserving the possibility that similar endpoint structures arise from different binding pathways.

For medicinal chemistry, this separation changes the optimization problem. Binding kinetics can be altered through chemical design, but relationships between structural modification and kinetic behavior vary across targets and mechanisms [24]. Affinity and residence time should therefore remain distinct design objectives rather than being merged into a generic binding-quality score. A modification that improves equilibrium affinity may leave dissociation unchanged, whereas another may alter a transition state and substantially modify residence time without producing an equally large change in equilibrium affinity. The architecture must be capable of representing such divergence because it is precisely where a static endpoint description becomes least informative.

Association provides an additional multiscale challenge. Protein–ligand encounter begins at spatial and temporal scales different from those governing short-range rearrangement within the binding site. Multiscale calculations combining Brownian dynamics with molecular dynamics have been developed to estimate association rates by separating diffusional encounter from atomistically resolved stages near the target [25]. This implies that a single trajectory representation is not necessarily the correct unit for the entire binding process. The architecture should instead permit different computational descriptions at different scales, with explicit interfaces where information passes from one level to another.

The third layer is therefore not another molecular coordinate system but a cross-level coupling structure. Molecular-state predictions influence which kinetic transitions are considered plausible; kinetic behavior can identify states whose functional importance would be underestimated from occupancy alone; and cellular measurements can invalidate an apparently convincing molecular account when accessibility, competition, localization, or off-target binding changes the realized interaction. These dependencies should be reciprocal where the evidence is reciprocal. Cellular evidence can motivate revision of molecular hypotheses just as molecular hypotheses can determine which cellular measurements are most informative.

This structure also places a limit on complexity. A conformational, kinetic, or cellular layer should be added only when it can change a testable inference. If a ligand series is robustly ranked across plausible receptor states, additional ensemble complexity may not alter the medicinal-chemistry decision. If compounds with similar equilibrium affinity show materially different dissociation behavior, kinetic information becomes decision-relevant. If purified-protein predictions fail to explain cellular engagement, the contextual layer becomes indispensable. Dynamic modeling is therefore an adaptive representation problem rather than a demand to maximize simulation detail.

These dependencies are represented in **Figure 1** as separable information fields connected only where a testable relationship exists. The figure deliberately avoids treating richer dynamics as an automatically superior endpoint; instead, it marks where evidence can support, qualify, or fail to support inference across molecular scales.

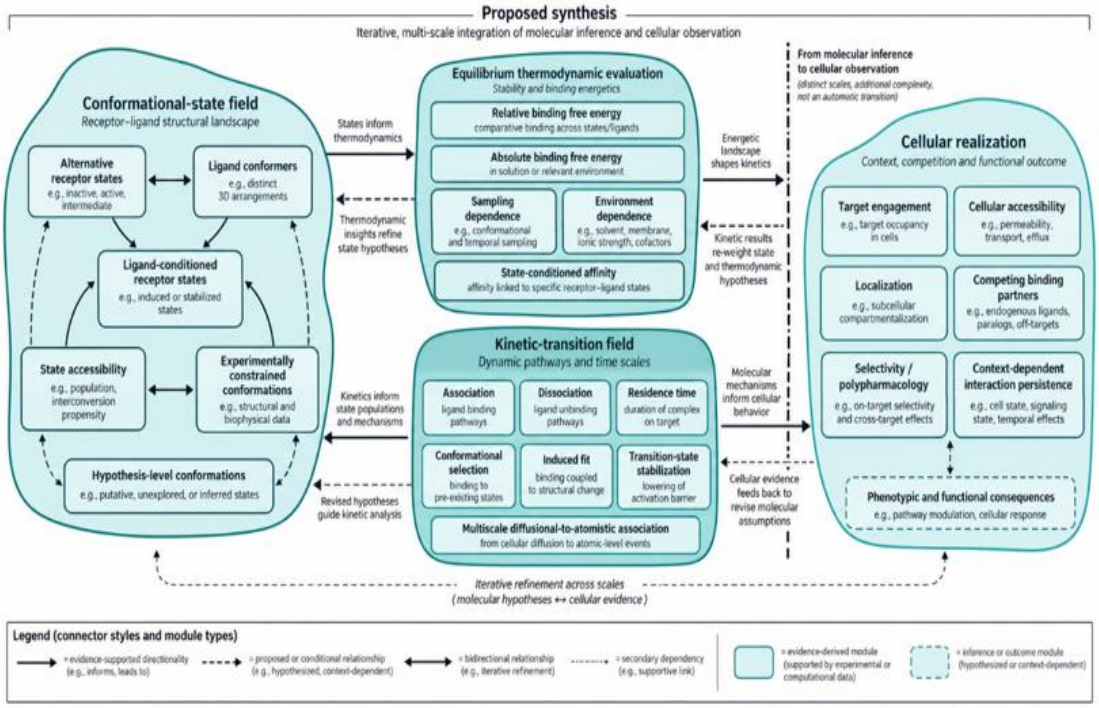


Figure 1. Cross-level dependencies governing dynamic protein–ligand action.

Validation and Translational Implications

The value of a dynamic architecture depends on whether its additional layers can be tested independently. Target-engagement assays illustrate why validation should be matched to the type of inference being made. Contemporary assay strategies span biochemical, biophysical, cellular, and proximity-based measurements, each interrogating a different aspect of target interaction [26]. A structural prediction may therefore be validated as a binding geometry without establishing residence time; a kinetic measurement may establish dissociation behavior without proving intracellular engagement; and cellular engagement may be demonstrated without establishing therapeutic effect.

Cellular validation is especially important because the environment surrounding a target can introduce interaction possibilities absent from purified molecular systems. Chemical-proteomic approaches have expanded the ability to map small-molecule binding sites directly in cells, exposing interactions under native or near-native biological conditions [27]. Such measurements provide a qualitatively different form of evidence from docking or free-energy calculation. They test whether a proposed molecular interaction survives cellular accessibility, competition, protein state, and local environment. Failure at this stage should not be interpreted merely as experimental noise; it may indicate that the molecular representation omitted a biologically decisive variable.

Selectivity creates a related cross-level test. Large-scale chemical-proteomic characterization of kinase inhibitors has revealed broad target landscapes that cannot be inferred reliably from one modeled target complex alone [28]. This matters because apparent success at the intended target can coexist with substantial off-target engagement. Dynamic molecular modeling should therefore treat cellular selectivity as an external constraint on interpretation rather than a property that can simply be extrapolated from a predicted pose or affinity score.

Some targets make the limits of static validation especially visible. Intrinsically disordered proteins and highly heterogeneous binding systems cannot be represented naturally by choosing a single stable receptor conformation. Ensemble docking has been used to address such systems by retaining multiple target states during ligand evaluation [29]. The broader implication is that validation depth should reflect the physical behavior of the target. A comparatively rigid binding site may be adequately represented by a small state set, whereas a disordered or strongly allosteric system may require substantially richer

conformational evidence. The architecture should adapt to this difference rather than impose the same dynamic complexity on every target.

Translational value also depends on how molecular evidence enters medicinal-chemistry decisions. Contemporary optimization rarely maximizes a single property; it balances potency, selectivity, physicochemical properties, exposure, safety, synthetic feasibility, and project-specific biological requirements [30]. Dynamic molecular predictions therefore acquire decision value only when they alter one of these choices in a defensible manner. A calculated kinetic advantage that does not survive cellular testing, a conformational hypothesis that does not change compound ranking, or an elaborate simulation that cannot discriminate among competing design options may be scientifically interesting without being operationally useful.

The same caution applies to AI-enabled molecular generation. Generative models can enlarge accessible chemical space and propose compounds conditioned on structural or molecular objectives, but generated structures still require chemical, physical, biological, and experimental filtering [31]. Dynamic information can enrich these objectives—for example by penalizing candidates whose predicted binding depends on an implausible state or by distinguishing compounds with different kinetic hypotheses—but it cannot remove the need for prospective testing.

Earlier deep-learning docking systems already demonstrated that learned scoring can improve pose evaluation within defined structural settings [32]. Large-scale deep docking similarly showed how machine learning can accelerate screening of enormous chemical libraries by reducing the number of compounds requiring expensive downstream evaluation [33]. More broadly, machine learning has been applied across target identification, molecular design, property prediction, and development activities [34]. These advances remain important components of the architecture developed here. The proposed shift is therefore not away from high-throughput static prediction but toward preventing efficiency gains at one representational level from being mistaken for validation at another.

A translational claim becomes increasingly demanding as it crosses representational boundaries. Predicting a physically credible pose requires less evidence than claiming a populated binding state; claiming a populated state requires different evidence from assigning a residence time; and cellular engagement requires evidence that neither structural nor kinetic modeling supplies automatically. Patient-level or clinical utility lies further downstream still and cannot be inferred solely because a dynamic molecular model has passed computational, biophysical, and cellular tests.

Figure 2 therefore represents translation as a field of evidential permissions and failures rather than as a linear validation pipeline. A prediction may gain support, lose support, or require a qualitatively different test when it moves from molecular representation toward a drug-design decision.

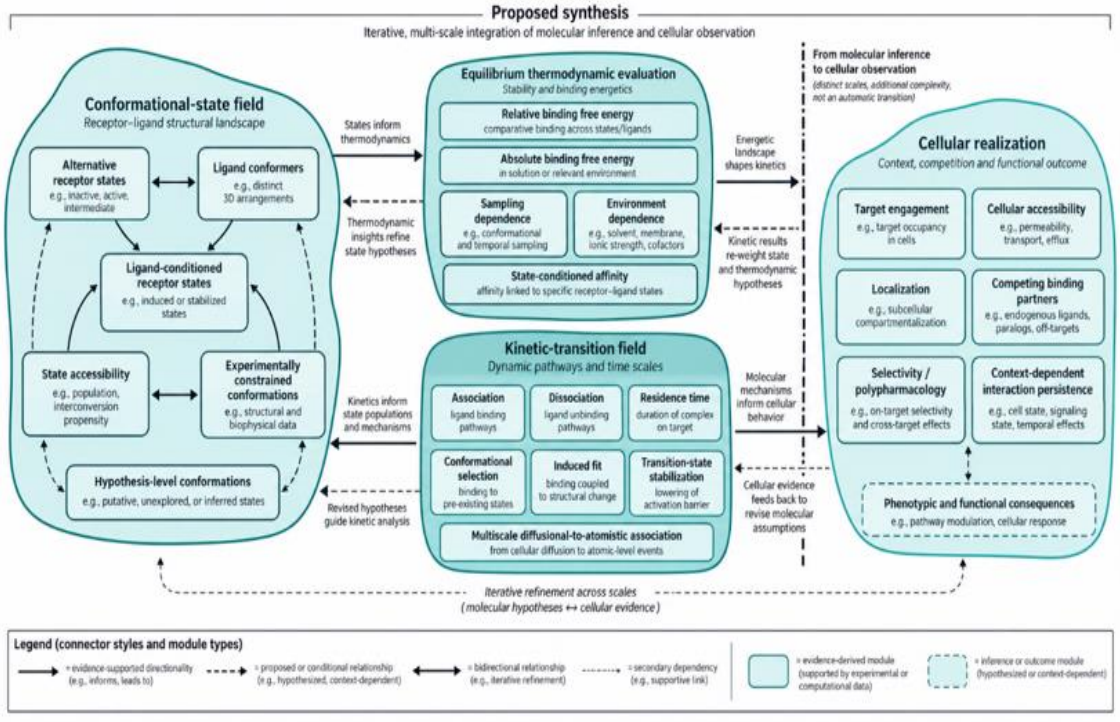


Figure 2. Evidence needed before dynamic molecular predictions can guide drug-design decisions.

Conclusion

Static protein–ligand models remain useful because molecular geometry matters, and efficient structure-based prediction will continue to be central to drug design. Their limitation emerges when structural fit is asked to stand in for phenomena that occur across different states, time scales, or biological environments. A single complex cannot determine the distribution of

accessible conformations, the rates connecting bound and unbound states, or whether the predicted interaction persists amid cellular competition, localization, and off-target binding.

A more defensible approach is therefore not to replace static modeling with maximal molecular complexity, but to add dynamic information selectively. Conformational ensembles are warranted when state dependence changes the interaction being modeled. Kinetic information is warranted when association or dissociation changes the pharmacological interpretation of compounds that appear similar at equilibrium. Cellular evidence becomes necessary when molecular predictions are used to support claims about target engagement or selectivity in biological systems. Each added layer must have an identifiable evidentiary function and a corresponding route to falsification.

The central implication is that dynamic molecular modeling should be organized around conditional inference across levels. Structural, thermodynamic, kinetic, and cellular evidence can constrain one another, but no layer automatically inherits the validity of another. The strongest AI-assisted molecular model is therefore not the one containing the greatest number of predicted variables. It is the one that preserves the distinctions required by the drug-design question, exposes where inference crosses an evidentiary boundary, and adds complexity only when doing so changes a testable scientific or medicinal-chemistry decision.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Bender A, Thomas MC, Scannell JW, Shaywitz DA, Ghiandoni GM, Greener JG, et al. Artificial intelligence in drug discovery—what it is, where we stand and the path forward. Nat Rev Drug Discov. 2026. doi:10.1038/s41573-026-01496-2
2. Abramson J, Adler J, Dunger J, Evans R, Green T, Pritzel A, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024;630(8016):493-500. doi:10.1038/s41586-024-07487-w
3. Krishna R, Wang J, Ahern W, Sturmfels P, Venkatesh P, Kalvet I, et al. Generalized biomolecular modeling and design with RoseTTAFold All-Atom. Science. 2024;384(6693). doi:10.1126/science.adl2528
4. Lu W, Zhang J, Huang W, Zhang Z, Jia X, Wang Z, et al. DynamicBind: Predicting ligand-specific protein-ligand complex structure with a deep equivariant generative model. Nat Commun. 2024;15(1):1071. doi:10.1038/s41467-024-45461-2
5. Qiao Z, Nie W, Vahdat A, Miller TF III, Anandkumar A. State-specific protein-ligand complex structure prediction with a multiscale deep generative model. Nat Mach Intell. 2024;6(2):195-208. doi:10.1038/s42256-024-00792-z
6. Lazou M, Khan O, Nguyen T, Padhorny D, Kozakov D, Joseph-McCarthy D, et al. Predicting multiple conformations of ligand binding sites in proteins suggests that AlphaFold2 may remember too much. Proc Natl Acad Sci U S A. 2024;121(48). doi:10.1073/pnas.2412719121
7. Fiorucci L, Schiavina M, Felli IC, Pierattelli R, Ravera E. Are protein conformational ensembles in agreement with experimental data? A geometrical interpretation of the problem. J Chem Inf Model. 2024;64(14):5392-401. doi:10.1021/acs.jcim.4c00582
8. Guan X, Tang QY, Ren W, Chen M, Wang W, Wolynes PG, et al. Predicting protein conformational motions using energetic frustration analysis and AlphaFold2. Proc Natl Acad Sci U S A. 2024;121(35). doi:10.1073/pnas.2410662121
9. Meller A, Ward M, Borowsky J, Kshirsagar M, Lotthammer JM, Oviedo F, et al. Predicting locations of cryptic pockets from single protein structures using the PocketMiner graph neural network. Nat Commun. 2023;14(1):1177. doi:10.1038/s41467-023-36699-3
10. Lyu J, Kapolka N, Gumpper R, Alon A, Wang L, Jain MK, et al. AlphaFold2 structures guide prospective ligand discovery. Science. 2024;384(6702). doi:10.1126/science.adn6354
11. Buttenschoen M, Morris GM, Deane CM. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. Chem Sci. 2024;15(9):3130-9. doi:10.1039/D3SC04185A
12. Siebenmorgen T, Menezes F, Benassou S, Merdivan E, Didi K, Mourão ASD, et al. MISATO: Machine learning dataset of protein-ligand complexes for structure-based drug discovery. Nat Comput Sci. 2024;4(5):367-78. doi:10.1038/s43588-024-00627-2
13. Wang L, Wang S, Yang H, Li S, Wang X, Zhou Y, et al. Conformational space profiling enhances generic molecular representation for AI-powered ligand-based drug discovery. Adv Sci (Weinh). 2024;11(40). doi:10.1002/advs.202403998
14. Cao D, Chen G, Jiang J, Yu J, Zhang R, Chen M, et al. Generic protein-ligand interaction scoring by integrating physical prior knowledge and data augmentation modelling. Nat Mach Intell. 2024;6(6):688-700. doi:10.1038/s42256-024-00849-z

15. Qian R, Xue J, Xu Y, Huang J. Alchemical transformations and beyond: Recent advances and real-world applications of free energy calculations in drug discovery. J Chem Inf Model. 2024;64(19):7214-37. doi:10.1021/acs.jcim.4c01024
16. Ries B, Alibay I, Anand NM, Biggin PC, Magarkar A. Automated absolute binding free energy calculation workflow for drug discovery. J Chem Inf Model. 2024;64(14):5357-64. doi:10.1021/acs.jcim.4c00343
17. Tsai HC, Xu J, Guo Z, Yi Y, Tian C, Que X, et al. Improvements in precision of relative binding free energy calculations afforded by the alchemical enhanced sampling (ACES) approach. J Chem Inf Model. 2024;64(18):7046-55. doi:10.1021/acs.jcim.4c00464
18. Zhang H, Im W. Ligand binding affinity prediction for membrane proteins with alchemical free energy calculation methods. J Chem Inf Model. 2024;64(14):5671-9. doi:10.1021/acs.jcim.4c00764
19. Stenström O, Diehl C, Modig K, Akke M. Ligand-induced protein transition state stabilization switches the binding pathway from conformational selection to induced fit. Proc Natl Acad Sci U S A. 2024;121(14). doi:10.1073/pnas.2317747121
20. Knockenhauer KE, Copeland RA. The importance of binding kinetics and drug-target residence time in pharmacology. Br J Pharmacol. 2024;181(21):4103-16. doi:10.1111/bph.16104
21. Liu H, Zhang H, IJzerman AP, Guo D. The translational value of ligand-receptor binding kinetics in drug discovery. Br J Pharmacol. 2024;181(21):4117-29. doi:10.1111/bph.16241
22. Adediwura VA, Koirala K, Do HN, Wang J, Miao Y. Understanding the impact of binding free energy and kinetics calculations in modern drug discovery. Expert Opin Drug Discov. 2024;19(6):671-82. doi:10.1080/17460441.2024.2349149
23. Xu C, Zhang X, Zhao L, Verkhivker GM, Bai F. Accurate characterization of binding kinetics and allosteric mechanisms for the HSP90 chaperone inhibitors using AI-augmented integrative biophysical studies. JACS Au. 2024;4(4):1632-45. doi:10.1021/jacsau.4c00123
24. Qiao H, Yuan Z, Xing L, He BX, Ma LY, He ZX. Innovative medicinal chemistry strategies for improving target binding kinetics in drug discovery. J Med Chem. 2025;68(21):22116-44. doi:10.1021/acs.jmedchem.5c02299
25. Muñiz-Chicharro A, Ganotra GK, Wade RC. A multiscale simulation approach to compute protein-ligand association rate constants by combining Brownian dynamics and molecular dynamics. J Chem Inf Model. 2025;65(20):11215-31. doi:10.1021/acs.jcim.5c01488
26. St John-Campbell S, Bhalay G. Target engagement assays in early drug discovery. J Med Chem. 2025;68(12):12331-68. doi:10.1021/acs.jmedchem.4c03115
27. Wozniak JM, Li W, Governa P, Chen LY, Jadhav A, Dongre A, et al. Enhanced mapping of small-molecule binding sites in cells. Nat Chem Biol. 2024;20(7):823-34. doi:10.1038/s41589-023-01514-z
28. Reinecke M, Brear P, Vornholz L, Berger BT, Seefried F, Wilhelm S, et al. Chemical proteomics reveals the target landscape of 1,000 kinase inhibitors. Nat Chem Biol. 2024;20(5):577-85. doi:10.1038/s41589-023-01459-3
29. Dhar A, Sisk TR, Robustelli P. Ensemble docking for intrinsically disordered proteins. J Chem Inf Model. 2025;65(13):6847-60. doi:10.1021/acs.jcim.5c00370
30. Rácz A, Mihalovits LM, Beckers M, Fechner N, Stiefl N, Sirockin F, et al. The changing landscape of medicinal chemistry optimization. Nat Rev Drug Discov. 2025;24(11):870-87. doi:10.1038/s41573-025-01225-1
31. Du Y, Jamasb AR, Guo J, Fu T, Harris C, Wang Y, et al. Machine learning-aided generative molecular design. Nat Mach Intell. 2024;6(6):589-604. doi:10.1038/s42256-024-00843-5
32. McNutt AT, Francoeur P, Aggarwal R, Masuda T, Meli R, Ragoza M, et al. GNINA 1.0: Molecular docking with deep learning. J Cheminform. 2021;13(1):43. doi:10.1186/s13321-021-00522-2
33. Gentile F, Agrawal V, Hsing M, Ton AT, Ban F, Norinder U, et al. Deep docking: A deep learning platform for augmentation of structure based drug discovery. ACS Cent Sci. 2020;6(6):939-49. doi:10.1021/acscentsci.0c00229
34. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5