



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Molecular Dynamics for Polypharmacology Must Resolve Competing Conformational Basins Across Targets

Zsolt Kovács^{1*}, Katalin Szabo¹, Gabor Nagy², Eszter Toth³

¹Department of Molecular Dynamics and Polypharmacology, Faculty of Pharmacy, Semmelweis University, Budapest, Hungary.

²Department of Conformational Basins and Target Resolution, Faculty of Pharmacy, University of Debrecen, Debrecen, Hungary.

³Department of Computational Polypharmacology, Faculty of Pharmacy, University of Pécs, Pécs, Hungary.

*Email: zsolt.kovacs@semmelweis.hu

ABSTRACT

Polypharmacology is usually represented as a relationship between one chemical structure and several targets, yet this representation suppresses a central physical complication: each target is itself an ensemble of interconverting conformations, and ligand recognition may occur through different subsets of those ensembles. Consequently, apparent multi-target activity can arise from different binding-competent states, different population shifts, and different kinetic pathways rather than from a transferable binding mechanism. This Perspective argues that molecular dynamics for polypharmacology should therefore compare competing conformational basins across targets rather than extend single-structure docking across a larger target panel. Evidence from G-protein-coupled receptors, nuclear receptors, kinases, enzyme isoforms, enhanced-sampling studies, and state-model methodology shows why static structural similarity, affinity, or a single predicted pose cannot by themselves establish cross-target conformational equivalence. We propose a cross-target ensemble logic in which target-local states are resolved first and compared only after their structural, thermodynamic, kinetic, and experimental support has been defined. This framing separates purposeful multi-target pharmacology from unresolved promiscuity while making uncertainty in pose selection, state definition, sampling, and transferability explicit.

Keywords: Polypharmacology, Molecular dynamics, Conformational ensembles, Free-energy landscapes, Binding kinetics, Multi-target drug design

Received: 04 March 2026; **Revised:** 29 May 2026; **Accepted:** 01 June 2026

INTRODUCTION

Polypharmacology has shifted from being regarded mainly as undesirable off-target activity toward an explicit medicinal-chemistry strategy in which a compound is designed to engage more than one biologically relevant target. Rational multitarget ligands, structure–multiple-activity relationships, and systematic analyses of dual-target-directed ligands all support the feasibility of this design space, while also showing that multi-target optimization is intrinsically more constrained than optimizing a single interaction [1–3]. The fundamental object is therefore not simply a molecule with several measured activities. It is a molecule whose activity profile must remain interpretable across distinct binding environments, target functions, and pharmacological objectives. Scale does not remove that constraint. Prospective docking of an approximately 900-million-compound library against three target pairs identified dual binders for selected combinations, but success was target-pair dependent and subsequent potency optimization remained difficult even where structural experiments supported predicted

binding poses [4]. Conversely, machine-learning models can predict broad activity patterns, as demonstrated by experimentally tested kinome-wide polypharmacology profiling [5]. These approaches provide valuable maps of chemical–target relationships, but neither a vector of predicted activities nor a collection of plausible poses necessarily reveals how a ligand is accommodated dynamically by each target.

The missing level is conformational. A protein does not offer one immutable binding site; it occupies populations of states whose accessibility and interconversion can change with ligand chemistry, coupling partners, activation state, solvent conditions, and simulation representation. For polypharmacology, this means that a ligand encounters several landscapes, not merely several structures. A compound may stabilize analogous states at two targets, exploit a pre-existing state at one while inducing or enriching another state at the second, or obtain similar endpoint affinities through kinetically and structurally different routes.

This Perspective develops that problem as a proposed cross-target ensemble framework. Its purpose is not to claim that a universal state mapping already exists. Rather, it argues that target-local conformational basins must be resolved before claims of shared binding logic, selectivity, or purposeful multi-target activity are made. Structural resemblance, thermodynamic favorability, kinetic persistence, and experimental compatibility are treated as related but non-interchangeable evidence. The resulting framework is intended to make explicit where molecular dynamics can strengthen polypharmacology and where simulation uncertainty prevents a medicinal-chemistry ranking from being mechanistically interpretable.

Polypharmacology as a multilandscape problem

Molecular dynamics changes the representational unit of drug recognition from a structure to a trajectory-distributed ensemble. MD can expose fluctuations, transient pockets, coupled motions, and alternative conformations that are absent from a single experimental snapshot, but the usefulness of those states depends on force-field adequacy and sufficient sampling [6]. For a single target, this already complicates interpretation. For several intended targets, the difficulty multiplies because each target may contain its own metastable states, barriers, ligand-entry routes, and binding-competent subsets.

Thermodynamics and kinetics must also remain distinct. Molecular simulation methods can estimate binding free energies, association and dissociation behavior, and the pathways connecting metastable states, but these quantities answer different questions [7]. Two targets could show similar calculated affinity while differing greatly in residence time or pathway accessibility. Conversely, similar residence behavior need not imply identical binding geometry or equal thermodynamic stabilization. A cross-target analysis that compresses such distinctions into one docking score or one affinity estimate therefore loses precisely the information needed to understand intentional polypharmacology.

Long-timescale MD coupled to Markov modelling has shown that complex association processes can be decomposed into metastable configurations and transition networks at atomic resolution [8]. This provides an important conceptual basis for polypharmacology: target engagement can be viewed as movement through states rather than as arrival at a unique endpoint. Yet the state network is not discovered independently of modelling choices. Markov state models require choices about features, dimensionality, clustering, lag times, and validation, and methodological advances have emphasized objective model selection rather than treating every geometric cluster as a physical state [9].

We therefore propose that a polypharmacological ligand should be represented as interacting with a set of target-local landscapes. Each landscape first requires its own state definitions and uncertainty limits. Cross-target comparison comes afterward. This order matters because the same coordinate value, RMSD threshold, or pocket opening cannot be assumed to possess the same kinetic or pharmacological meaning in two unrelated proteins. The goal is consequently not to force distinct proteins onto a common conformational template, but to determine which state properties are sufficiently comparable to support a medicinal-chemistry inference.

Target-specific conformational ensembles

Experimental receptor studies illustrate why target-local resolution must precede cross-target mapping. ^{19}F NMR analysis of the β_1 -adrenergic receptor showed that ligand efficacy is associated with redistribution among receptor conformations rather than a simple binary switch, while ternary receptor complexes retain region-specific conformational plasticity [10]. An “active” receptor structure is therefore not necessarily a singular dynamical state, and apparent functional similarity can conceal population differences.

The activation pathway itself may contain pharmacologically relevant intermediates. Combined molecular dynamics, Markov modelling, mutagenesis, and functional experiments on the angiotensin II type 1 receptor

identified intermediate conformations, including an intermediate containing a cryptic allosteric site [11]. Such findings do not establish that every target possesses an equivalent cryptic basin. They demonstrate instead that intermediate-state accessibility can be mechanistically consequential and must be resolved at the level of the individual target before a ligand's multi-target behavior is generalized.

Static architectural similarity can be equally deceptive. Cryo-EM structures of β 1-adrenergic receptor–Gs complexes bound to ligands with different efficacies showed broadly similar overall complex architectures while revealing local conformational and stability differences [12]. Thus, global structural resemblance is not sufficient evidence that two ligand conditions occupy equivalent energetic or dynamical states. This point becomes stronger in cross-target comparisons, where sequence, pocket architecture, coupling networks, and activation mechanisms may differ substantially.

Nor should a canonical conformational pathway be assumed simply because targets share a protein family. Work on FFAR1 demonstrated a noncanonical GPCR activation mechanism for particular allosteric agonists, involving conformational changes distinct from conventional descriptions of activation [13]. The significance for polypharmacology is methodological: family membership can guide hypotheses, but it cannot justify transferring a state topology from one target to another without target-specific support.

Ligand chemistry itself can redistribute pre-existing states. A PPAR γ ligand series spanning inverse agonism to agonism showed that subtle structural modifications shifted the receptor ensemble between transcriptionally repressive and active conformations already represented within the native conformational landscape [14]. This result is especially relevant to multi-target medicinal chemistry because it demonstrates how small chemical changes may change not only affinity but the populations of functionally differentiated states. Cross-target optimization must therefore ask which state populations are being changed at each target, not merely whether the ligand binds.

Shared and divergent binding-competent states

The concept of a binding-competent state needs stricter definition in polypharmacology. A conformation predicted to accommodate a ligand is not automatically a thermodynamically populated basin, and a geometrically similar pocket in another target is not automatically functionally equivalent. DynamicBind, for example, was developed to predict ligand-specific protein–ligand structures from unbound protein structures and can model substantial conformational adjustments and cryptic-pocket formation [15]. Such methods expand the structural hypotheses available for flexible targets, but a generated complex remains a structural prediction unless population, kinetics, or experimental evidence establishes more.

Time-resolved structural evidence shows why pathway information can matter. Photoswitch dissociation from the adenosine A2A receptor was accompanied by ligand-chemistry-dependent changes in the binding pocket and extracellular regions [16]. A ligand can therefore interact with a target through a sequence of conformational adjustments that is not recoverable from the initial and final structures alone. In a multi-target setting, two compounds might end at superficially similar complexes while differing in the pathway or barriers by which those complexes are reached.

Chemical similarity is likewise insufficient evidence for preserved selectivity. In a series of nearly 50 benzenesulfonamide derivatives evaluated across the 12 catalytically active human carbonic anhydrase isozymes, small chemical changes could produce very large shifts in affinity and selectivity [17]. The important point is not that every such selectivity shift is necessarily caused by a conformational-basin change. Rather, scaffold continuity cannot justify assuming cross-target equivalence, and unexpected activity differences should trigger examination of alternative binding states, local interactions, solvation, protonation, and other competing explanations.

Conformational-state preference can nevertheless be directly relevant to selectivity. Quantitative profiling of Aurora kinase inhibitors showed that inhibitors differ in their preference for activation states and linked those preferences to selectivity behavior [18]. This provides an empirical example in which ligand ranking depends on the state of the target being considered. It also exposes the central cross-target problem: a ligand may encounter a state that is readily available at one target but rare, differently stabilized, or kinetically inaccessible at another. For this Perspective, a shared binding-competent state is therefore proposed to mean a pair or set of target-local states that are sufficiently comparable for a specified inference—not states that merely superimpose geometrically. The required comparison may involve pocket interaction geometry, ligand orientation, state population, transition accessibility, thermodynamic stabilization, kinetic persistence, or functional coupling.

Which dimensions are necessary depends on the medicinal-chemistry question. A state can be shared for one analytical purpose and divergent for another.

A compound-ranking decision consequently requires explicit treatment of competing explanations. Similar activity across two targets may arise through analogous binding states, distinct states with convergent affinities, different kinetic routes, or assay-level effects that do not establish a common molecular mechanism. **Table 1** separates these cases so that transfer across targets is not assumed merely from an attractive structural correspondence.

Table 1. Cross-target decision cases for distinguishing shared binding competence from superficial structural agreement

Decision context	Evidence and limits	Permitted action
Compound-ranking case / Governing state condition	Mechanistic interpretation and competing explanation	Uncertainty and validation discriminator / Decision consequence and transfer boundary / Article-level test
Compound-ranking case: Similar ligand pose and similar local pocket arrangement Governing state condition: Each target supports a ligand-compatible local state	May indicate analogous recognition, but similar coordinates can arise without similar populations or transition kinetics	Uncertainty and validation discriminator: Test state occupancy, alternative poses, local dynamics, and orthogonal structural evidence Decision consequence and transfer boundary: Mechanistic transfer remains provisional until target-local dynamics are supported Article-level test: Does the apparent correspondence persist when each target is treated as an ensemble?
Compound-ranking case: Similar measured activity with divergent target conformations Governing state condition: Different target-local states can support comparable endpoint activity	Convergent pharmacology may arise through structurally different recognition mechanisms	Uncertainty and validation discriminator: Compare ligand interactions, state populations, pathway accessibility, and assay context Decision consequence and transfer boundary: Preserve multi-target activity as an observation without assigning a common molecular mechanism Article-level test: Can equal activity be explained without forcing state equivalence?
Compound-ranking case: Shared scaffold but sharply different target preference Governing state condition: One chemical modification alters compatibility with target-local states	Selectivity may reflect state preference, local interaction changes, or other physicochemical effects	Uncertainty and validation discriminator: Use analogue series, structural data, dynamics, and alternative-state testing Decision consequence and transfer boundary: Do not extrapolate selectivity from scaffold identity Article-level test: Which target-local change best discriminates the altered profile?
Compound-ranking case: State preference observed in one target but unresolved in another Governing state condition: Evidence is asymmetric across the target pair	The characterized target supplies a mechanistic hypothesis, not a transferable state assignment	Uncertainty and validation discriminator: Require independent sampling and validation for the unresolved target Decision consequence and transfer boundary: Withhold a shared-state claim Article-level test: Has the second target been resolved independently rather than fitted to the first?
Compound-ranking case: Multiple plausible binding states within either target Governing state condition: Compound ranking depends on unresolved basin identity or weighting	An apparently favorable state may be one member of a competing ensemble	Uncertainty and validation discriminator: Test alternative modes, kinetics, state stability, and sensitivity to representation Decision consequence and transfer boundary: Delay mechanistic ranking when state uncertainty could reverse the comparison Article-level test: Is the proposed between-target difference larger than the unresolved within-target ambiguity?

The same logic requires a visual separation between structural support, ensemble support, and cross-target inference. **Figure 1** therefore places the shared-state decision boundary between target-local landscapes and the eventual claim of transferable binding competence.




Figure 1. Structural resemblance is insufficient to establish a shared binding-competent basin across targets

Ligand-induced population shifts

The importance of conformational ensembles does not imply that every ligand-induced change should be interpreted from simulation alone. One route to reducing ambiguity is to constrain conformational hypotheses with orthogonal measurements. Hydrogen-exchange-informed molecular dynamics improved ligand-binding predictions in a benchmark involving PAK1 and MAP4K1 ligands, particularly where protein conformational adjustment challenged conventional flexible docking [19]. The broader lesson is that a simulated ligand-induced state becomes more credible when an independent observable can discriminate it from competing structural explanations.

Generative models add another capability but also another boundary. Generative deep learning has been used to broaden sampling of the highly dynamic Aβ42 conformational ensemble and to produce experimentally testable conformations beyond the original training snapshots [20]. Such exploration is potentially useful for targets whose accessible state space is difficult to cover with direct trajectories. However, generation and thermodynamic weighting are different tasks. The frequency with which a generative model produces a conformation should not automatically be interpreted as its equilibrium population, nor should a generated state be assumed to occupy the same energetic position in different targets.

For polypharmacology, ligand-induced population shifts should therefore be treated as target-local hypotheses that can become cross-target evidence only after their observability and weighting are established. The relevant question is not simply whether the ligand “moves” two proteins. It is whether the ligand measurably redistributes functionally relevant states in each target, whether the affected states are comparable for the intended inference, and whether alternative causes of the apparent shift have been excluded adequately.

Comparing free-energy basins across targets

A cross-target comparison requires more than identifying visually plausible states; the relative stability and accessibility of those states must also be considered. Enhanced-sampling methods such as LiGaMD3 can generate repeated ligand binding and unbinding events and support estimation of thermodynamic and kinetic quantities in validated systems [21]. Such approaches demonstrate that state populations and transitions can become computationally accessible, but they do not remove the need to establish convergence separately for each target. State definition itself can dominate the comparison. In the T4 lysozyme L99A model system, kinetics-based analysis resolved closed, intermediate, and open pocket states and showed why state assignments derived simply from structural similarity can fail to represent the underlying dynamics [22]. This has a direct conceptual consequence for multi-target analysis: basin boundaries cannot be exported from one protein to another merely

because both contain an apparently corresponding pocket motion. Each target first requires a state definition linked to its own dynamics.

Collective variables introduce a second uncertainty. Enhanced-sampling work has shown that free-energy reconstruction can remain feasible with imperfect collective variables under appropriately designed sampling schemes, while simultaneously emphasizing that the choice of coordinate influences convergence and exploration [23]. A seemingly stable cross-target free-energy difference should therefore be challenged by changing or perturbing the coordinates that define the relevant transition.

Reaction-coordinate fidelity matters for mechanistic interpretation as well as efficiency. Enhanced sampling based on reaction coordinates inferred from energy relaxation has demonstrated large acceleration of conformational transitions in selected systems while preserving natural transition pathways more faithfully than poorly chosen empirical coordinates [24]. The transferable lesson is not the reported acceleration factor; it is that a free-energy landscape can appear computationally converged while still being distorted by a coordinate that does not represent the actual transition mechanism.

State models add another layer. Coarse molecular states may violate the Markov assumption when intrastate and interstate timescales are insufficiently separated, producing biased kinetic estimates if transitions are counted naively [25]. Weighted-ensemble trajectories can likewise introduce trajectory-merging bias into long-lag Markov models unless the effect is corrected explicitly [26]. Thus, neither enhanced sampling nor state modelling should be treated as an uncertainty-free route to a cross-target ranking.

We propose that free-energy basins should be compared across targets only after a target-local adequacy step. At minimum, the analyst should know what defines the basin, whether the relevant transition is sampled, whether alternative state definitions change the result, whether the chosen coordinates preserve meaningful pathways, and whether the kinetic model is compatible with the generated trajectories. Only then does a difference between targets become interpretable as a potential molecular-design signal rather than a difference between modelling failures.

Selectivity versus intentional multi-target activity

Selectivity and polypharmacology are not opposites. A deliberately multi-target compound is selective in a different sense: it should preferentially engage a chosen pattern of targets while avoiding pharmacologically undesirable interactions. The challenge is that a cross-target simulation can manufacture apparent precision if each target is represented by one preferred pose or one insufficiently tested state. Recent benchmarking of docking protocols against alternative binding modes reinforces this concern, showing that recovery of viable alternative poses depends on the docking protocol and available structural information [27]. A top-ranked pose is therefore a starting hypothesis, not a unique definition of the ligand's binding basin.

This distinction changes how simulation evidence should be interpreted. If a compound appears favorable at two intended targets, the result may represent genuine designed multi-target recognition. It may instead reflect two different binding mechanisms that happen to yield similar endpoint affinities, or it may arise because one target's alternative poses or conformations were sampled more completely than the other's. Conversely, a ligand that looks selective against one static structure may lose that selectivity when a previously inaccessible binding-competent state becomes available in the nominal off-target.

The proposed criterion for intentional multi-target activity is therefore not "binding to more than one target." It is a sufficiently supported pattern of engagement in which the target set is pharmacologically motivated and the molecular evidence remains coherent after target-local conformational uncertainty has been examined. This criterion deliberately leaves room for mechanistic divergence. A desirable ligand need not bind all intended targets through the same state or pathway. What matters is that the evidence distinguishes deliberate activity from unresolved promiscuity and that mechanistic similarity is claimed only where it has actually been demonstrated.

This interpretation also prevents affinity from functioning as an unqualified decision rule. The carbonic-anhydrase and Aurora-kinase examples show that relatively small changes in ligand chemistry or target conformational state can alter selectivity patterns substantially. The large-library polypharmacology study likewise shows that finding credible dual-target poses and initial activities does not guarantee straightforward optimization. Together, these cases motivate a decision architecture in which structural, ensemble, and pharmacological evidence can agree or disagree without being collapsed into a single score.

Table 2 therefore treats concordance and discordance as informative outcomes. A compound may be retained as a promising intentional multitarget ligand even when its two target mechanisms differ, provided those mechanisms

are independently supported and the pharmacological target combination remains justified. Conversely, visually elegant pose similarity should not upgrade an uncertain interaction into purposeful polypharmacology.

Table 2. Evidence-reconciliation grid for separating intentional multi-target engagement from unresolved cross-target promiscuity

Evidence configuration	Interpretive alternatives	Adjudication requirement
Cross-target evidence pattern / Molecular or chemical state / Evidence producing the pattern	Inference not justified	Inference justified / Discriminating prospective test and transfer boundary
Cross-target evidence pattern: Activity and target-local state evidence are concordant at both intended targets Molecular or chemical state: Independently supported binding-competent states Evidence producing the pattern: Compatible assays, target-local structural evidence, and ensemble support	The two targets necessarily use the same mechanism	Inference justified: The compound has evidence consistent with intentional engagement of the specified target pair Discriminating prospective test and transfer boundary: Test whether cross-target state properties remain comparable under orthogonal structural or kinetic measurements
Cross-target evidence pattern: Similar activity but clearly divergent binding states Molecular or chemical state: Distinct target-local basins Evidence producing the pattern: Comparable endpoint pharmacology with different structural or conformational support	A shared binding mechanism or transferable state topology	Inference justified: Multi-target activity can remain intentional despite mechanistic divergence Discriminating prospective test and transfer boundary: Validate each target independently and preserve divergence in medicinal-chemistry interpretation
Cross-target evidence pattern: Similar predicted poses but asymmetric ensemble support Molecular or chemical state: One target resolved, the other dominated by modelling uncertainty Evidence producing the pattern: Docking agreement without equivalent dynamical or experimental evidence	Shared binding competence or equivalent affinity mechanism	Inference justified: Structural correspondence is a hypothesis Discriminating prospective test and transfer boundary: Improve sampling or obtain orthogonal evidence for the unresolved target before cross-target transfer
Cross-target evidence pattern: Apparent selectivity depends on one preferred pose Molecular or chemical state: Alternative modes remain plausible in at least one target Evidence producing the pattern: Ranking changes with pose or state initialization	A robust selectivity mechanism	Inference justified: Selectivity is conditional on unresolved structural assumptions Discriminating prospective test and transfer boundary: Propagate alternative binding modes into dynamics or free-energy testing
Cross-target evidence pattern: Multi-target activity extends beyond the pharmacologically intended set Molecular or chemical state: Several targets are compatible with the ligand Evidence producing the pattern: Broad activity profile without adequate state-level discrimination	That the profile represents purposeful or clinically advantageous multi-target design	Inference justified: The compound is polypharmacological in a descriptive sense Discriminating prospective test and transfer boundary: Resolve intended versus unintended target engagement within the relevant exposure and biological context
Cross-target evidence pattern: Cross-target ranking changes with state definition or sampling protocol Molecular or chemical state: Competing basins remain incompletely resolved	A medicinal-chemistry preference based on the current simulation alone	Inference justified: The ranking is method-sensitive Discriminating prospective test and transfer boundary: Return analysis to pose selection, state definition, sampling, and validation before prioritization

Evidence producing the pattern: Discordant MD, docking, or free-energy outcomes

Figure 2 makes the same distinction dynamically. Instead of treating multi-target activity as a single branch from ligand to several proteins, it represents each target as an independent state-transition route and places intended polypharmacology downstream of an evidence-reconciliation boundary.




Figure 2. Distinguishing designed multi-target engagement from state-dependent promiscuity

A cross-target ensemble framework

The practical problem is not how to place several independently simulated targets into one larger workflow, but how to decide when their state descriptions can legitimately be compared. Methods for accelerating protein–ligand binding and unbinding increasingly improve accessible timescales and throughput, yet their performance remains dependent on the system, collective variables, state representation, and validation strategy [28]. Cross-target analysis therefore requires a common comparison logic, not necessarily a common simulation protocol.

We propose that each target first pass through a target-local stage comprising pose uncertainty, conformational sampling, basin definition, and validation. Only then should states enter a cross-target comparison. A shared comparison anchor might be ligand-contact geometry, pocket hydration, state population, free-energy stabilization, transition accessibility, residence behavior, or an experimentally observable conformational response. The anchor must be chosen according to the medicinal-chemistry question rather than assumed from structural similarity. Experimental receptor systems already show that similar global structures can coexist with ligand-dependent local dynamics and population shifts [10, 12, 14].

The framework also treats disagreement as diagnostically useful. If activity, docking, state populations, free energies, and experimental observations disagree, the discordance should return the analysis to the level capable of producing it—pose selection, force-field representation, sampling, state definition, or assay interpretation. Averaging incompatible evidence into a single cross-target score would obscure rather than solve the mechanistic problem.

Table 3 expresses this sequence explicitly. It is a **proposed analytical ledger**, not a validated scoring system, and deliberately contains no universal numerical cutoff.

Table 3. Sequence-aware ledger for deciding when target-local ensembles become comparable across a polypharmacological target set

State and entry conditions	Evidence through time	Transition and reassessment
----------------------------	-----------------------	-----------------------------

Analytical stage / Molecular or chemical state	Cross-target comparison anchor / Experimental counterpart	Observable readout	Prospective test / Condition that interrupts progression / Manuscript-specific consequence
Analytical stage: Structural hypothesis Molecular or chemical state: One or more plausible ligand poses per target	Cross-target comparison anchor: Interaction geometry that is relevant to the intended comparison Experimental counterpart: Structural or binding-site evidence where available	Pose stability and competing orientations	Prospective test: Challenge the preferred pose against credible alternatives Condition that interrupts progression: Ranking depends strongly on initialization or omitted modes Manuscript-specific consequence: Return to pose generation before assigning state correspondence
Analytical stage: Target-local ensemble resolution Molecular or chemical state: Interconverting binding-compatible and alternative conformations	Cross-target comparison anchor: State properties defined independently within each target Experimental counterpart: Conformational probes or structural ensembles where feasible	Populations, transitions, local dynamics	Prospective test: Reproduce target-local state behavior with independent trajectories or methods Condition that interrupts progression: Apparent state exists only under one representation or sampling setup Manuscript-specific consequence: Do not export the state definition to another target
Analytical stage: Basin adequacy Molecular or chemical state: Metastable states with tested boundaries	Cross-target comparison anchor: Thermodynamic or kinetic property appropriate to the design question Experimental counterpart: Affinity, kinetic, or conformational measurements where applicable	Relative stability, transition accessibility, or persistence	Prospective test: Perturb state definitions and sampling coordinates Condition that interrupts progression: Between-target result changes materially with the basin definition Manuscript-specific consequence: Treat cross-target ordering as method-sensitive
Analytical stage: Cross-target reconciliation Molecular or chemical state: Independently supported target-local states	Cross-target comparison anchor: Comparable function of the state rather than coordinate identity Experimental counterpart: Orthogonal measurements across the target set	Concordant or deliberately divergent state behavior	Prospective test: Test whether the proposed correspondence survives independent validation Condition that interrupts progression: Similarity is purely geometric or evidence is asymmetric Manuscript-specific consequence: Preserve targets as mechanistically distinct
Analytical stage: Medicinal-chemistry interpretation Molecular or chemical state: Chemical series evaluated across the supported ensembles	Cross-target comparison anchor: Pattern of intended target engagement Experimental counterpart: Analogue-series pharmacology	State-sensitive activity or free-energy trend	Prospective test: Prospectively test compounds selected because of the ensemble hypothesis Condition that interrupts progression: Ranking is smaller than unresolved pose, parameter, or sampling uncertainty Manuscript-specific consequence: Withhold mechanistic prioritization until uncertainty is reduced

Computational requirements and validation

A cross-target ensemble analysis is only as transferable as its molecular representation. Force-field parameter coverage is therefore not a technical afterthought. Expansion and validation of CGenFF against a substantially enlarged organic-molecule training set improved several tested molecular-property benchmarks, illustrating that chemical-space coverage influences the robustness of atomistic modelling [29]. A ligand whose parameters are weakly supported should not be compared mechanistically with a well-parameterized analogue as though the uncertainty were equivalent.

Sampling of the ligand itself is equally important. Terminal-Flip Monte Carlo accelerated difficult torsional transitions and improved relative binding free-energy convergence in several benchmark systems [30]. The specific performance gains are method- and system-dependent, but the general implication is broader: an apparent protein-state difference may actually originate from insufficient ligand conformational equilibration.

Validation should therefore interrogate the entire inferential chain. Collective-variable selection, reaction-coordinate fidelity, and state-model assumptions can each alter the apparent landscape [23–25]. Where feasible, structural measurements, hydrogen exchange, spectroscopy, kinetic experiments, mutational tests, or other orthogonal observables should discriminate competing simulated explanations. Agreement with one endpoint quantity is valuable but should not automatically validate every intermediate state or transition pathway.

The proposed framework consequently treats validation as layered rather than binary. A state can be structurally plausible yet thermodynamically uncertain; thermodynamically stable yet kinetically inaccessible on the relevant timescale; or computationally reproducible yet unsupported experimentally. Cross-target inference becomes strongest when the evidence required for the intended claim is explicit and comparable across targets.

Medicinal chemistry implications

For medicinal chemistry, the central change is from optimizing interactions with isolated structures to optimizing a pattern of state engagement. This does not require every intended target to share an identical binding mechanism. A useful multi-target ligand may stabilize different target-local conformations if those states independently support the desired pharmacology.

The converse is equally important. Minor chemical modifications can sharply change cross-target selectivity, as demonstrated across carbonic-anhydrase isoforms [17], while conformational-state preference can contribute directly to kinase selectivity [18]. Such cases argue against assuming that a scaffold carries a fixed target profile independent of state accessibility.

A practical ranking rule follows from this perspective: when unresolved pose, parameterization, state-definition, or sampling uncertainty is large enough to plausibly reverse a proposed cross-target difference, the mechanistic ranking should be withheld. The appropriate response is not to manufacture a confidence score but to identify the uncertainty source and design the next discriminating computation or experiment.

The value of molecular dynamics in polypharmacology therefore lies less in producing more poses than in exposing when target engagement depends on competing states, pathways, and timescales. Used this way, simulation becomes a means of deciding which medicinal-chemistry conclusions are supportable and which remain conditional.

CONCLUSION

Polypharmacology is a multilandscape problem because every intended target contributes its own conformational states, transition barriers, and validation uncertainties. Similar affinity, activity, or static binding geometry cannot establish that two targets are engaged through equivalent molecular mechanisms. The proposed cross-target ensemble framework therefore resolves target-local states before comparing them and treats structural, thermodynamic, kinetic, and experimental evidence as complementary rather than interchangeable. Its central discipline is also its principal limitation: cross-target interpretation should stop when unresolved modelling uncertainty can still explain the apparent difference. This perspective does not provide a universally validated ranking algorithm. It provides a testable architecture for determining when molecular dynamics has earned a mechanistic polypharmacology claim.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Proschak E, Stark H, Merk D. Polypharmacology by design: A medicinal chemist's perspective on multitargeting compounds. *J Med Chem.* 2019;62(2):420-44. doi:10.1021/acs.jmedchem.8b00760
2. López-López E, Medina-Franco JL. Toward structure-multiple activity relationships (SMARTs) using computational approaches: A polypharmacological perspective. *Drug Discov Today.* 2024;29(7):104046. doi:10.1016/j.drudis.2024.104046

3. Lembo V, Bottegoni G. Systematic investigation of dual-target-directed ligands. *J Med Chem.* 2024;67(12):10374-85. doi:10.1021/acs.jmedchem.4c00838
4. Wu Y, Vigneron S, Braz J, Srinivasan K, Fink EA, Huang XP, et al. Large library docking for polypharmacology. *J Med Chem.* 2026;69(5):6210-29. doi:10.1021/acs.jmedchem.5c03810
5. Li X, Li Z, Wu X, Xiong Z, Yang T, Fu Z, et al. Deep learning enhancing kinome-wide polypharmacology profiling: Model construction and experiment validation. *J Med Chem.* 2020;63(16):8723-37. doi:10.1021/acs.jmedchem.9b00855
6. Hollingsworth SA, Dror RO. Molecular dynamics simulation for all. *Neuron.* 2018;99(6):1129-43. doi:10.1016/j.neuron.2018.08.011
7. 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
8. 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
9. 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
10. Frei JN, Broadhurst RW, Bostock MJ, Solt A, Jones AJY, Gabriel F, et al. Conformational plasticity of ligand-bound and ternary GPCR complexes studied by 19F NMR of the β 1-adrenergic receptor. *Nat Commun.* 2020;11(1):669. doi:10.1038/s41467-020-14526-3
11. Lu S, He X, Yang Z, Chai Z, Zhou S, Wang J, et al. Activation pathway of a G protein-coupled receptor uncovers conformational intermediates as targets for allosteric drug design. *Nat Commun.* 2021;12(1):4721. doi:10.1038/s41467-021-25020-9
12. Su M, Paknejad N, Zhu L, Wang J, Do HN, Miao Y, et al. Structures of β 1-adrenergic receptor in complex with Gs and ligands of different efficacies. *Nat Commun.* 2022;13(1):4095. doi:10.1038/s41467-022-31823-1
13. Powers AS, Khan A, Paggi JM, Latorraca NR, Souza S, Di Salvo J, et al. A non-canonical mechanism of GPCR activation. *Nat Commun.* 2024;15(1):9938. doi:10.1038/s41467-024-54103-6
14. MacTavish BS, Zhu D, Shang J, Shao Q, He Y, Yang ZJ, et al. Ligand efficacy shifts a nuclear receptor conformational ensemble between transcriptionally active and repressive states. *Nat Commun.* 2025;16(1):2065. doi:10.1038/s41467-025-57325-4
15. 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
16. Glover H, Saßmannshausen T, Bertrand Q, Trabuco M, Slavov C, Bacchin A, et al. Photoswitch dissociation from a G protein-coupled receptor resolved by time-resolved serial crystallography. *Nat Commun.* 2024;15(1):10837. doi:10.1038/s41467-024-55109-w
17. Zakšauskas A, Paketurytė-Latvė V, Jankūnaitė A, Čapkauskaitė E, Becart Y, Smirnov A, et al. Affinity and selectivity of protein-ligand recognition: A minor chemical modification changes carbonic anhydrase binding profile. *J Med Chem.* 2025;68(16):17752-73. doi:10.1021/acs.jmedchem.5c01421
18. Lake EW, Muretta JM, Thompson AR, Rasmussen DM, Majumdar A, Faber EB, et al. Quantitative conformational profiling of kinase inhibitors reveals origins of selectivity for Aurora kinase activation states. *Proc Natl Acad Sci U S A.* 2018;115(51). doi:10.1073/pnas.1811158115
19. Walters BT, Patapoff AW, Kiefer JR, Wu P, Wang W. Integrating hydrogen exchange with molecular dynamics for improved ligand binding predictions. *J Chem Inf Model.* 2025;65(12):6144-54. doi:10.1021/acs.jcim.5c00397
20. Ruzmetov T, Hung TI, Jonnalagedda SP, Chen SH, Fasihianifard P, Guo Z, et al. Sampling conformational ensembles of highly dynamic proteins via generative deep learning. *J Chem Inf Model.* 2025;65(5):2487-502. doi:10.1021/acs.jcim.4c01838
21. 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
22. Zhang C, Osato M, Mobley DL. Kinetics-based state definitions for discrete binding conformations of T4 L99A in MD via Markov state modeling. *J Chem Inf Model.* 2024;64(23):8870-79. doi:10.1021/acs.jcim.4c01364

23. Ray D, Rizzi V. Enhanced sampling with suboptimal collective variables: Reconciling accuracy and convergence speed. *J Chem Theory Comput.* 2025;21(1):58-69. doi:10.1021/acs.jctc.4c01231
24. Li H, Ma A. Enhanced sampling of protein conformational changes via true reaction coordinates from energy relaxation. *Nat Commun.* 2025;16(1):786. doi:10.1038/s41467-025-55983-y
25. Sartore S, Teichmann F, Stock G. Markov-type state models to describe non-Markovian dynamics. *J Chem Theory Comput.* 2025;21(5):2757-65. doi:10.1021/acs.jctc.4c01630
26. Bose S, Kilinc C, Dickson A. Markov state models with weighted ensemble simulation: How to eliminate the trajectory merging bias. *J Chem Theory Comput.* 2025;21(4):1805-16. doi:10.1021/acs.jctc.4c01141
27. Kubincová A, Çınaroğlu SS, Ongsioco J, Hahn DF, Gapsys V, Tresadern G, et al. Benchmarking docking protocols on predicting alternative binding modes. *J Chem Inf Model.* 2026;66(14):8269-83. doi:10.1021/acs.jcim.6c00356
28. Zia SR, Coricello A, Bottegoni G. Increased throughput in methods for simulating protein ligand binding and unbinding. *Curr Opin Struct Biol.* 2024;87:102871. doi:10.1016/j.sbi.2024.102871
29. Croitoru A, Kumar A, Lambry JC, Lee J, Sharif S, Yu W, et al. Increasing the accuracy and robustness of the CHARMM General Force Field with an expanded training set. *J Chem Theory Comput.* 2025;21(6):3044-65. doi:10.1021/acs.jctc.5c00046
30. Wang X, Ding Y, Zou R, Qu T, Zheng H. Terminal-Flip Monte Carlo: Accelerating the convergence of molecular dynamics and alchemical free energy calculations. *J Chem Theory Comput.* 2026;22(11):5849-57. doi:10.1021/acs.jctc.6c00177