



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Hidden Water Networks Can Reorder Medicinal Chemistry Priorities Before Affinity Changes Are Visible

Lukas Weber^{1*}, Sophie Janssen², Jonas Richter¹, Elena Fischer³

¹Department of Medicinal Chemistry and Water Networks, Faculty of Pharmacy, University of Stuttgart, Stuttgart, Germany.

²Department of Molecular Recognition and Hydration Effects, Faculty of Pharmaceutical Sciences, Eindhoven University of Technology, Eindhoven, Netherlands.

³Department of Drug Discovery and Solvation, Faculty of Pharmacy, University of Basel, Basel, Switzerland.

*Email: lukas.weber@uni-stuttgart.de

ABSTRACT

Medicinal chemistry commonly ranks closely related compounds by affinity, potency, structural complementarity, or calculated binding free energy. These endpoints can obscure a mechanistically important variable: the organization and response of water at the protein–ligand interface. Interfacial waters differ in persistence, thermodynamic stability, exchange accessibility, network connectivity, and coupling to protein conformation. Ligand modifications may therefore reorganize hydration without immediately producing a large difference in apparent affinity. This Perspective argues that such divergence can become chemically informative before conventional endpoint separation emerges. The central distinction is not between “good” and “bad” waters, but among hydration states that are structurally present, thermodynamically stable, kinetically exchangeable, reorganized by ligand substitution, or coupled to protein motion. Explicit-water simulations, enhanced sampling, room-temperature structural measurements, thermodynamic experiments, and water-aware machine-learning approaches increasingly expose these distinctions, although their outputs are method- and system-dependent. A proposed water-aware prioritization logic is developed in which reproducible hydration-state divergence motivates targeted follow-up rather than automatic compound promotion. The framework remains hypothesis-generating: hydration information should refine medicinal-chemistry interpretation only when its structural, thermodynamic, dynamic, and chemical consequences can be separated from modeling artifacts and other co-varying determinants of binding.

Keywords: Structural waters, Hydration networks, Ligand binding, Binding free energy, Enthalpy–entropy compensation, Medicinal chemistry

Received: 09 March 2025; **Revised:** 30 May 2025; **Accepted:** 04 June 2025

INTRODUCTION

Protein–ligand affinity is a state function, but medicinal-chemistry decisions are often mechanistic. Two compounds with similar measured affinity may reach that endpoint through different combinations of direct contacts, protein reorganization, desolvation, retained interfacial water, and solvent-mediated interactions. Experimental solvent perturbation has demonstrated that water-network organization can alter measured binding thermodynamics even when the protein and ligand remain nominally the same chemical species [1]. This makes hydration more than a background correction to protein–ligand recognition: it is part of the physical state being compared.

A second complication is that the water pattern visible in a deposited structure is not necessarily the hydration state sampled under another experimental condition. Analysis of paired cryogenic and room-temperature structures shows that crystallographic waters differ substantially in their probability of persisting when temperature changes, allowing structural waters to be separated statistically into more persistent and more readily displaced populations [2]. Structural occupancy should therefore not be treated as equivalent to biological permanence, residence time, or favorable contribution to affinity.

Medicinal chemistry has long recognized the potential importance of binding-site water, but simple rules such as “displacing an ordered water is favorable” are thermodynamically incomplete. Water removal can incur or recover enthalpic and entropic terms whose balance depends on the local environment, the substituent introduced, the protein conformation, and the surrounding solvent network [3]. The relevant design problem is consequently not whether a pocket contains water, but whether a ligand perturbation changes a reproducible hydration state in a way that matters to the chemical decision being made.

This Perspective develops that distinction across structural waters, network topology, residence and exchange, ligand-induced reorganization, and enthalpy–entropy compensation. The article does not propose that hydration information should replace affinity. Rather, it argues that affinity-equivalent or weakly differentiated compounds may warrant different experimental or computational follow-up when their hydration states diverge in mechanistically interpretable ways. The resulting prioritization logic is explicitly proposed and remains conditional on validation within the target, chemical series, structural state, and computational regime in which the water signal was obtained.

Structural waters as dynamic participants

A structural water is often described as though crystallographic detection identifies a stable microscopic object with a fixed functional role. That interpretation merges several distinct quantities. Water occupancy describes whether solvent is found at a position. Thermodynamic stability concerns the free-energy cost of moving a water between a confined environment and bulk solvent. Solvent-mediated interaction energy concerns how the surrounding water field changes when ligand or protein chemistry is perturbed. These properties are related but not interchangeable. Computational analyses across idealized cavities, protein hydration shells, and protein–ligand interfaces show that apparently similar water positions can have markedly different free-energy consequences and can mediate both favorable and unfavorable interactions [4–6].

This distinction matters because high local density does not necessarily mean that a water is expensive to displace, and low apparent mobility does not uniquely specify its contribution to binding. Confinement can be stabilized by favorable protein–water interactions while simultaneously producing an unfavorable entropy relative to bulk. Conversely, removal of an apparently constrained water can require loss of favorable enthalpic interactions that offset the entropy gained upon release. The meaningful object for medicinal chemistry is therefore not a coordinate but a hydration state defined by the surrounding chemical and conformational environment.

Direct ligand design provides a practical illustration. In the A2A adenosine receptor, a binding-site water was treated as an explicit component of the recognition environment during analogue design. Modifications that altered the ligand’s ability to interact with or replace this water produced corresponding changes in binding behavior, showing that an individual hydration feature can become chemically actionable when the structural hypothesis is tested through an analogue series [7]. Importantly, the lesson is not that displacement itself is desirable. The same site can favor retention, bridging, or substitution depending on the ligand architecture.

A useful interpretation is therefore to regard “structural water” as an ensemble-conditioned label. A detected molecule may belong to a locally persistent network under one temperature, protein conformation, protonation state, or ligand state but exchange rapidly or disappear under another. For drug design, at least three questions should remain separate: is water present, how stable is the associated hydration state, and how strongly is that state coupled to the ligand and protein configuration? The remainder of the article treats these as different evidential dimensions rather than compressing them into a single favorable-versus-unfavorable water classification.

Water-network topology

Binding-site hydration is rarely a collection of independent solvent molecules. Hydrogen-bonded chains, bridging waters, partially connected clusters, and solvent pathways can couple regions of a pocket that appear geometrically separate in a static ligand-contact map. Consequently, a substituent may alter hydration beyond the atoms it directly contacts. The relevant representation is network topology: which sites are occupied, how they connect to

protein and ligand polar groups, whether they communicate with bulk solvent, and how the apo and ligand-bound states differ.

HydraMap v.2 illustrates how such information can be extracted computationally at scale. Statistical potentials derived from structural data can identify hydration sites, estimate desolvation-related properties, and distinguish features such as bridging and potentially displaceable waters across apo and holo environments [8]. Such maps are useful because they convert an otherwise visually intuitive solvent pattern into a reproducible spatial representation. Their output, however, remains a modeled estimate conditioned by structural databases and the underlying statistical assumptions.

Thermodynamic profiling provides an additional layer. WaterKit was developed to characterize hydration sites in terms of thermodynamic properties relevant to ligand-binding interpretation and high-throughput virtual screening [9]. The important distinction is that two sites with comparable coordinates or occupancies need not possess comparable energetic character. Position identifies where water is likely to occur; thermodynamic profiling asks what maintaining or replacing that state may cost.

Recent geometric learning approaches can accelerate this process further. Hydration-site localization and thermodynamic profiling have been learned from explicit-water simulation data and tested against experimental or case-specific observations, reducing the computational cost required to generate water hypotheses during lead discovery [10]. These models are potentially valuable upstream filters, particularly when hundreds or thousands of structures or analogues must be triaged. Their speed should not be confused with increased evidential certainty: predictions inherit the force fields, sampling regimes, training-domain composition, and structural assumptions from which their targets are derived.

Topology is also relevant to selectivity. Grand canonical Monte Carlo analysis across 35 bromodomains showed that water stability varies substantially even among related binding pockets [11]. Such family-level differences can reveal solvent features that sequence comparison or shape complementarity alone may not emphasize. Yet a predicted stability difference is not itself proof that displacement will create selective affinity. It instead defines a testable chemical opportunity: a ligand perturbation can be designed to interact differently with a network whose thermodynamic state is predicted to vary between homologous targets.

The medicinal-chemistry value of topology therefore lies in comparative state definition. A network can be used to ask whether two ligands occupy the same pocket while inducing different hydration connectivity, whether a substituent breaks a bridge to bulk solvent, or whether a homologous target supports an alternative water arrangement. These questions become particularly important when endpoint affinities remain similar because the solvent response may already be diverging before the net free-energy difference becomes large.

Residence, exchange, and displacement

Static hydration topology does not specify how quickly water exchanges. A site can appear occupied in crystallographic density while belonging to an ensemble that changes substantially with temperature or local conformation. In Hsp90 α , ten matched room-temperature and cryogenic structural datasets showed that both water networks and nearby protein conformations repopulate as temperature changes [12]. These observations caution against assigning a single biologically privileged water pattern solely from a low-temperature structure.

Residence and exchange also create computational problems. When a cavity is poorly connected to bulk solvent, spontaneous insertion or removal of water may occur on timescales longer than a conventional molecular-dynamics calculation samples reliably. Comparisons of enhanced techniques, accelerated enveloping distribution sampling, and nonequilibrium grand canonical Monte Carlo demonstrate that difficult hydration transitions can require explicit methods for crossing insertion, deletion, or occupancy barriers [13–15]. The choice among these approaches is system-dependent; their common lesson is that absence of a transition during ordinary simulation does not prove that the corresponding hydration state is inaccessible.

Grand canonical methods make this issue explicit by treating water exchange with an external reservoir through an imposed chemical potential. Replica exchange across chemical-potential states can reduce variance in estimates of water occupancy and binding free energy, improving the precision of calculations in which water-state sampling is otherwise slow [16]. Precision, however, remains distinct from accuracy. A precisely sampled state can still be wrong if the force field, protonation state, structural model, or chemical potential is inappropriate.

The problem becomes especially important during alchemical free-energy calculations. When a ligand transformation opens or closes a cavity, a buried water may need to enter or leave before the new ligand state can equilibrate. Free-energy perturbation combined with grand canonical Monte Carlo has shown improved robustness in transformations where conventional sampling fails to reorganize such buried waters adequately [17].

In this setting, water is not simply another interaction term in the final free energy. It is a hidden state variable whose incomplete equilibration can make the apparent ligand free-energy difference depend on simulation history. These observations motivate a separation between displaceability and observed displacement. A water absent from the final ligand structure may have been removed through a thermodynamically favorable process, an unfavorable process compensated elsewhere, or a kinetic pathway inaccessible on the timescale of a simulation. Similarly, a retained water may be strongly stabilized, weakly confined but frequently replenished from bulk, or preserved only because a particular protein conformation remains populated. Medicinal interpretation therefore requires the water state to be placed on both thermodynamic and dynamic coordinates before displacement is assigned a design meaning.

Ligand-induced reorganization

Ligand modification perturbs more than direct protein contacts. Substituent growth can exclude one water while stabilizing another; a new hydrogen-bond donor can recruit a bridge; removal of a polar group can open a solvent pathway; and a seemingly local change can shift protein side chains or pocket conformations that reorganize the surrounding network. The resulting state is therefore a coupled ligand–protein–water ensemble.

Explicit water treatment can improve modeling of such perturbations, but its utility is heterogeneous. In a ten-system evaluation of SILCS-WATER, explicit consideration of ligand-binding-site water networks improved relative-affinity prediction in six systems, with substantial improvement in five, while the remaining systems did not show the same benefit [18]. That heterogeneity is scientifically important. It argues against universal correction factors and instead suggests that sensitivity to hydration modeling can itself identify systems in which water-state assumptions deserve greater scrutiny.

One strategy is to separate water-state sampling from the ligand transformation. A separation-of-states approach for relative binding free energies was developed specifically for cases in which trapped water reorganizes too slowly to equilibrate naturally during an alchemical perturbation [19]. The conceptual advantage is important beyond the method itself: a calculated ligand difference can be decomposed into a chemical transformation and a hydration-state transition rather than forcing both processes to occur simultaneously.

Water-network information can also be encoded directly in predictive representations. A two-state binding-affinity model incorporating apo and holo water networks improved benchmark performance relative to models that lacked equivalent solvent-state information [20]. The result supports the proposition that ligand binding is better represented as a transition between hydrated states than as a ligand inserted into a single dry protein environment. It does not establish that water-aware machine learning will outperform alternative methods prospectively across new chemical space.

Experimental density provides another constraint on reorganization. GOLEM incorporates explicit water molecules during cryo-EM-guided ligand docking, allowing water-mediated configurations to be evaluated alongside ligand placement against experimental maps [21]. In such applications, water modeling is useful because it participates in structural assignment; it should not be interpreted as an independent measurement of binding affinity or mechanism.

Medicinal-chemistry case studies demonstrate why these distinctions matter. In a BRD4 discovery program, conserved-water analysis was incorporated into virtual screening and subsequent optimization, producing the initial inhibitor SQ-1 and supporting development of the more active analogue SQ-17 [22]. The campaign demonstrates that water-aware reasoning can be operationally useful, but it does not isolate water as the sole cause of the activity gain. Ligand chemistry, direct protein contacts, conformational effects, and other physicochemical changes were modified simultaneously.

A complementary pattern appears in optimization of macrocyclic inhibitors of the Keap1–Nrf2 interaction. Structural and thermodynamic analysis showed that major potency improvements were accompanied by changes in binding-site hydration and protein flexibility [23]. Here again, solvent reorganization and conformational adaptation were coupled rather than experimentally separable causes.

The resulting article-level distinction is that ligand-induced hydration divergence can precede decisive affinity divergence. This does not mean that the hydration change should automatically promote one compound. It means that two affinity-similar compounds need not be mechanistically equivalent. If one analogue preserves a stable bridging network while another requires water displacement, protein rearrangement, or a slowly exchanging trapped state, the two compounds may have different vulnerabilities to structural state, force-field assumptions, temperature, target mutation, or later chemical substitution. Hydration divergence therefore becomes a reason to

investigate a compound pair more deeply before assuming that their present affinity ranking contains all of the useful mechanistic information.

The relationships developed above require a visual architecture that separates ligand perturbation from hydration response and from the inference ultimately used for compound decisions.




Figure 1. Ligand perturbations can separate hydration-state trajectories before binding endpoints separate

Enthalpy–entropy compensation

The free energy of binding integrates enthalpic and entropic terms, so equal or similar affinity does not require equal molecular physics. Enthalpy–entropy compensation can arise from real molecular rearrangements, solvent reorganization, conformational effects, or experimental and statistical artifacts [24]. A medicinal-chemistry series can therefore display small differences in ΔG even while ΔH and $-T\Delta S$ change substantially in opposite directions. Endpoint affinity alone cannot determine which microscopic process generated the compensation.

Water is one possible contributor but should not become a default explanation. Controlled aqueous recognition experiments demonstrate that changes in desolvation and retained solvent interactions can redistribute enthalpic and entropic contributions in ways that are not captured by simplistic hydrophobic or water-release narratives [24–26]. These systems provide physical evidence for the principle that association can remain favorable even when solvent interactions are retained and that seemingly modest chemical modifications can alter the balance between desolvation and direct intermolecular interactions. Their quantitative thermodynamics should not be transferred directly to protein–ligand systems, where protein flexibility, heterogeneous surfaces, protonation states, and larger conformational ensembles add additional variables.

For medicinal chemistry, compensation creates a particular interpretive hazard when compounds are ranked only by current affinity. Consider two analogues whose binding free energies are close. One may gain favorable direct interactions while paying an increased desolvation or conformational cost. Another may preserve a water-mediated network and achieve a similar ΔG through a different balance of enthalpy and entropy. If the endpoint is the only recorded variable, these states appear equivalent. Yet later substitution can perturb them differently because the two ligands occupy different positions within the coupled ligand–protein–water energy landscape.

The article therefore proposes that thermodynamic equivalence should not automatically be treated as hydration-state equivalence. A water-aware comparison becomes most informative when three conditions are met. First, a hydration difference is reproducible across an appropriate structural, computational, or experimental method. Second, the difference has a chemically interpretable relationship to ligand substitution rather than appearing only as an unstable model artifact. Third, the proposed consequence can be challenged experimentally or through an

orthogonal calculation. Failure to satisfy these conditions should preserve the compounds as mechanistically unresolved rather than create an artificial ranking.

The distinction is summarized below because the relevant scenarios differ simultaneously in molecular state, thermodynamic interpretation, transferability, expected failure signature, and validation requirement.

Table 1. Thermodynamic equivalence can conceal hydration states with different medicinal-chemistry consequences

Persistent water-mediated ligand state Ligand-driven displacement state Compensated hydration-reorganization state Slow or trapped-water transition Coupled hydration–protein reorganization					
Molecular or chemical state	Persistent water-mediated ligand state	Ligand-driven displacement state	Compensated hydration-reorganization state	Slow or trapped-water transition	Coupled hydration–protein reorganization
Governing thermodynamic condition	Retained solvent interactions contribute materially to the bound-state balance rather than being completely removed on association	Loss of confined water may yield favorable, neutral, or unfavorable net contribution depending on interactions surrendered and gained	Favorable and unfavorable enthalpic/entropic changes offset, producing little change in net binding free energy	Apparent free energy depends on whether hydration states equilibrate during the calculation or experiment being interpreted	Water rearrangement and protein conformational adaptation contribute jointly to the observed thermodynamic state
Consequence for the present compound decision	Do not penalize the compound simply because a crystallographic water remains; determine whether the retained network is chemically reproducible	Do not equate visible water displacement with an affinity advantage	Similar affinity should not end mechanistic comparison when hydration and thermodynamic decompositions diverge	Delay ranking based on small calculated affinity differences until water-state convergence is demonstrated	Avoid attributing an observed gain exclusively to either direct ligand contact or water displacement
Transfer boundary	Structural occupancy alone does not establish favorable thermodynamics or long residence	Host-guest or calculated displacement energy cannot be transferred quantitatively without system-specific validation	Compensation may arise from protein, ligand, solvent, or measurement effects rather than water alone	Simulation convergence does not establish experimental occupancy or residence	Causal separation is difficult when protein and hydration states co-vary
Failure signature	Network disappears with temperature, alternative structure, or adequate sampling	Water leaves structurally but affinity or thermodynamic decomposition does not improve as expected	Large change in component thermodynamics without stable structural or dynamic explanation	Hysteresis, starting-state dependence, or inconsistent free energies between hydration initializations	Affinity trend changes with protein state, structural preparation, or conformational restraint

Experimental counterpart	Room-temperature structural comparison, solvent-sensitive thermodynamics, or designed analogue perturbation	Analogue series with matched structural and thermodynamic measurements	Calorimetry combined with structural or simulation evidence	Alternative initialization, enhanced water sampling, or experimentally constrained structure	Structural ensemble comparison plus thermodynamic measurements
Prospective test	Modify the ligand interaction while preserving geometry and test whether the network and binding response change together	Compare substitutions designed to displace versus retain the same hydration region	Test an analogue predicted to perturb the implicated network selectively	Repeat the ligand comparison with a method that explicitly samples insertion/removal or separates hydration states	Design a perturbation expected to preserve direct contacts while altering hydration or conformational coupling
Hidden-water article-level test	Proposed: prioritize mechanistic follow-up only if network persistence and chemical sensitivity agree	Proposed: treat displacement as informative only when its energetic consequence can be separated from co-varying ligand contacts	Proposed: affinity-equivalent compounds enter follow-up only when hydration divergence is independently supported	Proposed: unresolved hydration-state dependence is a computational uncertainty state, not a compound advantage	Proposed: elevate the mechanism only when competing conformational explanations have been challenged

The table intentionally contains no numerical score. Its function is to separate situations in which similar affinity is accompanied by different hydration physics from situations in which apparent water differences are insufficiently validated to alter medicinal-chemistry interpretation.

Why similar affinity can conceal different water physics

A small difference in apparent affinity can be chemically reassuring while remaining mechanistically ambiguous. This is particularly important for congeneric series in which a substituent perturbs a buried or partially enclosed hydration site. If water rearrangement is slow relative to the computational protocol, two ligand states can appear artificially similar, artificially different, or dependent on the hydration state from which the calculation was initiated. Nonequilibrium relative binding free-energy calculations developed for systems containing trapped waters illustrate the magnitude of this problem. Across eight test systems, explicitly managing water rearrangement yielded relative binding free energies within 1.1 kcal/mol of experiment, with reported statistical errors below 0.4 kcal/mol [27]. These results should not be generalized into a universal accuracy guarantee; they demonstrate instead that a hidden hydration coordinate can become a material source of error when ligand transformations require water displacement that ordinary sampling does not accomplish reliably.

The deeper medicinal-chemistry implication is that affinity equivalence and state equivalence are different propositions. Two analogues can occupy overlapping crystallographic poses and produce similar ΔG while differing in whether binding retains a water bridge, empties a confined site, changes solvent connectivity to bulk, or couples hydration to a protein rearrangement. Enthalpy–entropy compensation can further compress these differences at the level of net free energy. Thus, a narrow affinity interval does not justify concluding that two analogues are mechanistically interchangeable.

This distinction is especially consequential early in optimization. When affinity differences are already large, water analysis may help explain an established SAR trend. When affinities remain close, hydration-state divergence can instead identify where the next discriminating experiment should be placed. A compound whose predicted advantage exists only after a trapped water is artificially removed should not automatically outrank an analogue whose water state equilibrates robustly. Conversely, a reproducible network difference that survives alternative starting structures, enhanced sampling, and experimental structural constraints may deserve greater mechanistic attention even if the current potency values remain nearly indistinguishable.

The proposed concept of water-physics divergence before affinity divergence therefore refers to an investigational state, not a validated biomarker. It contains no universal score, cutoff, or claim that the compound with the more favorable calculated water displacement will later become the superior lead. Its purpose is narrower: when the observable endpoint compresses several microscopic contributions into a similar ΔG , solvent organization can provide an additional axis along which apparently equivalent compounds are challenged.

Water-aware compound prioritization

Large-scale structural analysis now makes such challenges feasible earlier than would have been practical when hydration interpretation depended mainly on manual inspection. Analysis of more than 9,000 protein–ligand complexes showed that water distributions and ligand-associated displacement patterns can be predicted systematically, although positional recovery remains imperfect: approximately 38% of predicted waters were within 1.0 Å of experimental positions and about 83% within 1.5 Å [28]. These values demonstrate useful structural recoverability rather than affinity-prediction accuracy. A correctly positioned water can still have uncertain residence, thermodynamic stability, or relevance to a particular ligand transformation.

For prioritization, water information should therefore be treated as a sequence of increasingly demanding questions rather than as one score. The first question is whether an analogue changes a hydration feature reproducibly. The second is whether that feature remains stable to reasonable variation in structural preparation, temperature, protonation state, force field, or sampling method. The third is whether the change is chemically interpretable—for example, whether a substituent removes a water bridge, opens a solvent pathway, or creates a new coupled hydration–protein state. Only after these questions are addressed does it become useful to ask whether the hydration difference explains or predicts a medicinal-chemistry consequence.

The framework proposed here consequently uses water information to redirect attention, not to declare a winner. A robust hydration divergence can justify additional analogue design, room-temperature structural characterization, calorimetry, water-sensitive simulation, or free-energy calculations that explicitly permit solvent exchange. A weak or method-dependent signal should instead remain an uncertainty. This approach deliberately preserves discordance. If one explicit-water method predicts a decisive displacement while another converged method retains the water, the scientifically appropriate state is unresolved rather than an averaged or selectively preferred answer.

Chemical perturbation remains central. Computational hydration analysis becomes most informative when paired with analogues that test the hypothesized mechanism. If a proposed affinity advantage depends on replacing a particular water, a matched substitution that preserves geometry while altering the capacity for displacement provides a more informative test than simply extending the same scoring workflow. Likewise, if a predicted network controls selectivity between homologous targets, chemically addressing that network should generate a target-dependent response before strong mechanistic language is justified.

In this sense, water-aware prioritization is not an alternative to conventional medicinal chemistry. It is a way of deciding which uncertainty to attack next. Affinity, selectivity, permeability, metabolic stability, synthetic tractability, safety, and exposure remain independent decision dimensions. Hydration information becomes valuable when it distinguishes compounds that conventional binding endpoints have not yet separated and when the distinction can be translated into a falsifiable chemical or experimental test.

Experimental and simulation validation

Validation must match the claim being made. A predicted water position requires a different test from a predicted residence time, displacement free energy, ligand pose, or thermodynamic consequence. No single technique observes all of these quantities, and agreement among methods is meaningful only when the methods interrogate comparable states.

Room-temperature structural measurements provide one important counterweight to exclusive reliance on cryogenic coordinates. Serial crystallography applied in a drug-discovery setting to soluble epoxide hydrolase revealed temperature-dependent ligand-binding behavior and improved definition of flexible regions relative to previously available cryogenic structures [29]. Such observations do not establish that every room-temperature crystal structure is closer to the complete biological ensemble. They establish that temperature can materially alter the structural state from which medicinal-chemistry hypotheses are generated.

Dynamic simulation supplies a complementary form of evidence. In Photosystem II, room-temperature XFEL crystallography and crystalline molecular-dynamics simulations produced compatible water-density patterns, while the trajectories additionally exposed coordinated water exchange that could not be recovered from static

density alone [30]. The biological system is not a drug-binding pocket, so its water kinetics cannot be transferred quantitatively to medicinal chemistry. Its methodological lesson is nevertheless relevant: agreement in average water location does not imply that a static experiment determines exchange dynamics, and a dynamic simulation should not be considered validated merely because it reproduces one time-averaged density feature.

For the present perspective, validation should therefore be modular. Structural persistence can be challenged by matched temperature states or independent structural measurements. Hydration topology can be tested by comparing predictions across experimentally constrained protein conformations. Exchange limitations can be probed through enhanced sampling or alternative initial hydration states. Thermodynamic interpretations can be challenged through calorimetry or solvent-sensitive experiments where appropriate. Pose hypotheses can be tested against crystallographic or cryo-EM density. Relative binding free-energy predictions that depend on buried-water transitions should be repeated with protocols that explicitly permit insertion, deletion, or separation of hydration states.

The proposed decision architecture does not require every compound to undergo every test. That would turn mechanistic resolution into an impractical screening burden. Instead, the intensity of validation should increase with the importance of the water-dependent inference. A minor hydration difference used only to generate an analogue hypothesis may justify computational cross-checking. A water mechanism used to explain an unexpected selectivity inversion requires stronger structural and thermodynamic support. A hydration feature proposed as a prospective basis for advancing one lead over another demands still stronger evidence, ideally including a designed chemical perturbation capable of falsifying the proposed mechanism.

Failure is informative in this framework. If a water-network difference disappears at room temperature, collapses under adequate exchange sampling, depends strongly on one force field, or cannot survive a matched analogue test, the appropriate conclusion is not that water is unimportant. It is that the specific hydration hypothesis was insufficiently robust to carry the medicinal-chemistry decision. Conversely, convergence across structural, thermodynamic, dynamic, and chemical evidence can justify progressively stronger mechanistic interpretation while still remaining confined to the studied target and chemical series.

CONCLUSION

Binding affinity compresses a multidimensional molecular process into a single thermodynamic endpoint. Hidden hydration networks are one reason that compounds with similar affinity need not occupy equivalent mechanistic states. Water molecules can differ in persistence, connectivity, exchange accessibility, displacement cost, and coupling to protein conformation, and ligand modification can reorganize these properties before a large difference in net binding free energy becomes visible.

The practical consequence is not that medicinal chemistry should rank compounds by a universal water score. Such a score is not supported by the evidence. Water-aware interpretation is most useful when it identifies a reproducible difference that conventional affinity measurements have not resolved and converts that difference into a falsifiable experiment or chemically precise next step.

The proposed framework therefore treats hydration divergence as an evidence-dependent reason for differential follow-up. Structural occupancy must remain distinct from residence and thermodynamic stability; calculated displacement must remain distinct from experimentally established benefit; and improved performance in a water-aware computational method must remain distinct from prospective superiority in chemical optimization. Discordant results are uncertainty states rather than failures to be hidden.

When these boundaries are preserved, hidden water networks provide more than retrospective explanation. They can expose mechanistic distinctions between apparently equivalent compounds early enough to influence which analogues are synthesized, which simulations require stronger sampling, which structural states deserve experimental interrogation, and which affinity similarities should not yet be interpreted as chemical equivalence. In that restricted but consequential sense, water physics can reorder medicinal-chemistry priorities before affinity changes become obvious.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Szalai TV, Bajusz D, Börzsei R, Zsidó BZ, Ilaš J, Ferenczy GG, et al. Effect of water networks on ligand binding: Computational predictions vs experiments. *J Chem Inf Model*. 2024;64(23):8980-98. doi:10.1021/acs.jcim.4c01291
2. Seffernick JT, Fischer M. An experimental proxy of water displaceability for ligand discovery. *Nat Methods*. 2025;22(7):1476-85. doi:10.1038/s41592-025-02724-0
3. Spyraakis F, Ahmed MH, Bayden AS, Cozzini P, Mozzarelli A, Kellogg GE. The roles of water in the protein matrix: A largely untapped resource for drug discovery. *J Med Chem*. 2017;60(16):6781-827. doi:10.1021/acs.jmedchem.7b00057
4. Setiadi J, Biedermann F, Nau WM, Gilson MK. Thermodynamics of water displacement from binding sites and its contributions to supramolecular and biomolecular affinity. *Angew Chem Int Ed Engl*. 2025;64(35):e202505713. doi:10.1002/anie.202505713
5. Shi J, Cho JH, Hwang W. Heterogeneous and allosteric role of surface hydration for protein-ligand binding. *J Chem Theory Comput*. 2023;19(6):1875-87. doi:10.1021/acs.jctc.2c00776
6. Cui D, Zhang BW, Matubayasi N, Levy RM. The role of interfacial water in protein-ligand binding: Insights from the indirect solvent mediated potential of mean force. *J Chem Theory Comput*. 2018;14(2):512-26. doi:10.1021/acs.jctc.7b01076
7. Matricon P, Suresh RR, Gao ZG, Panel N, Jacobson KA, Carlsson J. Ligand design by targeting a binding site water. *Chem Sci*. 2021;12(3):960-8. doi:10.1039/D0SC04938G
8. Li Y, Zhang Z, Wang R. HydraMap v.2: Prediction of hydration sites and desolvation energy with refined statistical potentials. *J Chem Inf Model*. 2023;63(15):4749-61. doi:10.1021/acs.jcim.3c00408
9. Eberhardt J, Forli S. WaterKit: Thermodynamic profiling of protein hydration sites. *J Chem Theory Comput*. 2023;19(9):2535-56. doi:10.1021/acs.jctc.2c01087
10. Hinz FB, Masters MR, Nguyen JT, Mahmoud AH, Lill MA. Accelerated hydration site localization and thermodynamic profiling. *J Chem Inf Model*. 2025;65(6):2794-805. doi:10.1021/acs.jcim.4c02349
11. Aldeghi M, Ross GA, Bodkin MJ, Essex JW, Knapp S, Biggin PC. Large-scale analysis of water stability in bromodomain binding pockets with grand canonical Monte Carlo. *Commun Chem*. 2018;1:19. doi:10.1038/s42004-018-0019-x
12. Stachowski TR, Vanarotti M, Seetharaman J, Lopez K, Fischer M. Water networks repopulate protein-ligand interfaces with temperature. *Angew Chem Int Ed Engl*. 2022;61(31):e202112919. doi:10.1002/anie.202112919
13. Ge Y, Wych DC, Samways ML, Wall ME, Essex JW, Mobley DL. Enhancing sampling of water rehydration on ligand binding: A comparison of techniques. *J Chem Theory Comput*. 2022;18(3):1359-81. doi:10.1021/acs.jctc.1c00590
14. Gracia Carmona O, Gillhofer M, Tomasiak L, de Ruiter A, Oostenbrink C. Accelerated enveloping distribution sampling to probe the presence of water molecules. *J Chem Theory Comput*. 2023;19(11):3379-90. doi:10.1021/acs.jctc.3c00109
15. Deng J, Cui Q. Efficient sampling of cavity hydration in proteins with nonequilibrium grand canonical Monte Carlo and polarizable force fields. *J Chem Theory Comput*. 2024;20(5):1897-911. doi:10.1021/acs.jctc.4c00013
16. Ross GA, Bruce Macdonald HE, Cave-Ayland C, Cabedo Martinez AI, Essex JW. Replica-exchange and standard state binding free energies with grand canonical Monte Carlo. *J Chem Theory Comput*. 2017;13(12):6373-81. doi:10.1021/acs.jctc.7b00738
17. Ross GA, Russell E, Deng Y, Lu C, Harder ED, Abel R, et al. Enhancing water sampling in free energy calculations with grand canonical Monte Carlo. *J Chem Theory Comput*. 2020;16(10):6061-76. doi:10.1021/acs.jctc.0c00660
18. Kumar A, Goel H, Yu W, MacKerell AD Jr. Modeling ligand binding site water networks with site identification by ligand competitive saturation: Impact on ligand binding orientations and relative binding affinities. *J Chem Theory Comput*. 2024;20(24):11032-48. doi:10.1021/acs.jctc.4c01165

19. Wagle S, Merz PT, Ge Y, Bayly CI, Mobley DL. Leveraging a separation of states method for relative binding free energy calculations in systems with trapped waters. *J Chem Theory Comput.* 2024;20(24):11013-31. doi:10.1021/acs.jctc.4c01145
20. Qu X, Dong L, Luo D, Si Y, Wang B. Water network-augmented two-state model for protein-ligand binding affinity prediction. *J Chem Inf Model.* 2024;64(7):2263-74. doi:10.1021/acs.jcim.3c00567
21. Zhao Z, Tajkhorshid E. GOLEM: Automated and robust cryo-EM-guided ligand docking with explicit water molecules. *J Chem Inf Model.* 2024;64(14):5680-90. doi:10.1021/acs.jcim.4c00917
22. Zhong H, Wang X, Chen S, Wang Z, Wang H, Xu L, et al. Discovery of novel inhibitors of BRD4 for treating prostate cancer: A comprehensive case study for considering water networks in virtual screening and drug design. *J Med Chem.* 2024;67(1):138-51. doi:10.1021/acs.jmedchem.3c00996
23. Begnini F, Geschwindner S, Johansson P, Wissler L, Lewis RJ, Danelius E, et al. Importance of binding site hydration and flexibility revealed when optimizing a macrocyclic inhibitor of the Keap1-Nrf2 protein-protein interaction. *J Med Chem.* 2022;65(4):3473-517. doi:10.1021/acs.jmedchem.1c01975
24. Fox JM, Zhao M, Fink MJ, Kang K, Whitesides GM. The molecular origin of enthalpy/entropy compensation in biomolecular recognition. *Annu Rev Biophys.* 2018;47:223-50. doi:10.1146/annurev-biophys-070816-033743
25. Tobajas-Curiel G, Sun Q, Sanders JKM, Ballester P, Hunter CA. Substituent effects on aromatic interactions in water. *Chem Sci.* 2023;14(23):6226-36. doi:10.1039/D3SC01027A
26. Tobajas-Curiel G, Sun Q, Sanders JKM, Ballester P, Hunter CA. Aromatic interactions with heterocycles in water. *Chem Sci.* 2023;14(40):11131-40. doi:10.1039/D3SC03824F
27. Wagle S, Bayly CI, Mobley DL. Advancing binding affinity calculations: A non-equilibrium simulations approach for calculation of relative binding free energies in systems with trapped waters. *J Chem Theory Comput.* 2025;21(15):7593-604. doi:10.1021/acs.jctc.5c00758
28. Morningstar-Kywi N, Wang K, Asbell TR, Wang Z, Giles JB, Lai J, et al. Prediction of water distributions and displacement at protein-ligand interfaces. *J Chem Inf Model.* 2022;62(6):1489-97. doi:10.1021/acs.jcim.1c01266
29. Dunge A, Phan C, Uwanguo O, Bjelcic M, Gunnarsson J, Wehlander G, et al. Exploring serial crystallography for drug discovery. *IUCrJ.* 2024;11(Pt 5):831-42. doi:10.1107/S2052252524006134
30. Doyle MD, Bhowmick A, Wych DC, Lassalle L, Simon PS, Holton J, et al. Water networks in Photosystem II using crystalline molecular dynamics simulations and room-temperature XFEL serial crystallography. *J Am Chem Soc.* 2023;145(27):14621-35. doi:10.1021/jacs.3c01412