Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Alternate-conformer incompatibility as the dominant bottleneck in no-repair receptor preparation: a frozen compatibility audit of 30 public PDB structures

Andrés Monreal Hernández

andres.monreal@ues.mx

Universidad Estatal de Sonora

Sara Lizbeth Franco Amaya

Universidad de Sonora

Carlos Ivanhoe Martínez Osorio

Universidad de Sonora

Research Article

Keywords: molecular docking, receptor preparation, alternate conformers, AutoDock Vina, Meeko, reproducibility

Posted Date: August 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10707112/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Context:

Reference-pose redocking studies are commonly used to evaluate docking protocols, but the receptor-preparation step that precedes docking is rarely audited on its own terms, and preparation failures are typically discarded silently rather than reported. This work applied a frozen, no-repair receptor-preparation protocol (Meeko 0.7.1) without exception to 30 candidate protein-ligand complexes from the RCSB Protein Data Bank (X-ray, ≤ 2.0 Å resolution, one auditable non-polymer ligand). Of 23 preparation attempts, 7 (30%) produced a receptor accepted without repair, alternate-location selection, or residue deletion; 94% of rejections involved an unresolved alternate side-chain conformer, a pattern that persisted among every candidate with an explicitly recorded sub-1.0 Å resolution. The 7 prepared receptors were carried through identity-mapped, root-mean-square deviation (RMSD)-verified reference-pose docking with AutoDock Vina 1.2.5: five of six RMSD-verified cases recovered the experimental pose within 2.0 Å, one did not (~ 5.6 Å despite successful preparation), and one docking run completed but could not be RMSD-verified due to a pose-count mismatch, reported rather than discarded.

Methods

Candidate registration, structural classification, receptor extraction, and Meeko preparation were scripted with no manual intervention; every command, checksum, and outcome is versioned. Reference-pose RMSD used RDKit's symmetry-aware, substructure-matched heavy-atom alignment, independently cross-checked with a from-scratch NumPy Kabsch implementation. Interface contacts, buried solvent-accessible surface area (Shrake-Rupley), and standard success-threshold classification were computed directly from atomic coordinates. All scripts, manifests, and outputs are publicly archived and independently reproducible.

1 Introduction

Molecular docking is widely used to generate structural hypotheses about how a small molecule may bind a protein target. Reference-pose redocking — attempting to recover an experimentally observed ligand pose from the corresponding receptor coordinates — is a standard way to sanity-check a docking protocol before it is applied to unknown ligands. Public structural archives such as the RCSB Protein Data Bank make the source coordinates for this kind of check explicit and citable [1]. What is less commonly reported is the receptor-preparation step that precedes docking: converting a deposited crystal structure into a chemically well-defined, dockable receptor file is not automatic, and preparation tools routinely reject structures that require a repair decision the analyst has not yet made — for example, choosing among modeled alternate side-chain conformers, or deciding how to treat a non-polymer component that is not a genuine ligand. Recent reviews of open-source and AI-augmented docking workflows continue to list structure-preparation heterogeneity among the practical bottlenecks

limiting reproducibility across current pipelines [2]. These rejections are typically discarded silently by whoever is assembling a benchmark set, so the receptor-preparation failure rate itself is rarely quantified or discussed as a first-class result.

This work is not a docking-accuracy benchmark, a drug-design study, or a source of any biological-activity or binding-affinity claim. It is a frozen, auditable, no-repair receptor-preparation protocol applied without exception to a declared 30-candidate pool drawn from the public RCSB PDB, with every outcome — including every rejection — retained, classified, and reported. Preserving the full outcome, not only the successes, follows established reproducibility practice for computational research [3, 4]. The guiding question is deliberately narrower than “does docking work”: given a public protein-ligand complex, what documentable conditions of provenance, structure, composition, and preparation determine whether it is fit for a reference-pose redocking experiment, and can that determination itself be made auditable and reproducible?

2 Materials and methods

2.1 Candidate registration and inclusion criteria

Candidates were registered from the RCSB Protein Data Bank [1] under four criteria, checked at registration time and never relaxed after the fact: (i) the deposition method is X-ray crystallography; (ii) the reported resolution is ≤ 2.0 Å; (iii) the RCSB entry lists exactly one non-polymer entity, excluding solvent, buffer, crystallization additives, and simple ions — verified per candidate via the RCSB Data API rather than trusted from entry-level summary counts alone, since summary counts alone produced two false positives during discovery; and (iv) the reference ligand pose is directly auditable from the deposited coordinates. The registry was capped at 30 candidates under a pre-declared expansion plan.

2.2 Structural classification

Each registered candidate was automatically classified as clean (single model, single declared polymer chain, single ligand instance, no other non-polymer component) or contextual (any additional structural complexity — multiple chains, multiple ligand instances, extra non-polymer components, glycans, or metals). Clean candidates entered a fixed default preparation pipeline. Contextual candidates were left pending an explicit, case-specific policy and were never processed under the default clean-case rule.

2.3 Frozen no-repair preparation policy

For each processed candidate, the declared receptor chain and reference ligand instance were selected by minimum original-coordinate ligand-to-polymer atom distance. Water and the declared ligand were always removed. The extracted receptor was passed to Meeko 0.7.1 [5] (mk_prepare_receptor, PDB reader; coordinates parsed via Gemmi [6] during upstream classification) with no alternate-location choice, no coordinate repair, no --allow_bad_res, and no residue deletion. A run was scored as prepared only if the tool returned exit code zero and produced a nonempty PDBQT (AutoDock/Vina receptor

coordinate format) file; every other outcome was retained as a failure with its exact error message, normalized into one of three classes.

2.4 The retained-component extension

The default contextual policy strips every non-polymer component that is not the declared reference ligand. Five contextual cases were instead processed under an explicit, chemically justified retained-component policy, added to the extraction script as an optional per-case column rather than a blanket rule change: pilot-008 (3CJO, KSP/Eg5 motor domain with allosteric inhibitor K30) retained its two extra components, ADP and Mg2+, the motor domain's physiological nucleotide cofactor bound at a site distinct from the inhibitor pocket; pilot-006 (3PTB, bovine trypsin with the competitive inhibitor benzamidine, BEN) retained CA, trypsin's well-characterized structural calcium-binding site [7], distinct from the S1 catalytic pocket where benzamidine binds; and expansion-001 (1B9V, influenza neuraminidase with RA2) retained CA, neuraminidase's own structural calcium site, after directly verifying that its other extra component, NAG, carries no covalent bond to the polymer and was therefore stripped under the default rule rather than retained. Two further cases — pilot-002 (1HVR, HIV protease with the cyclic-urea inhibitor XK2) and expansion-019 (9HOO) — had been flagged for an "extra non-polymer component" that, on direct atom-level inspection of the deposited mmCIF (macromolecular Crystallographic Information File), turned out in both cases to be a covalently modified in-chain residue (CSO and CSS respectively) rather than a free ligand or cofactor at all; both were retained as an intrinsic part of the receptor chain under the same no-repair policy that leaves the declared chain otherwise unmodified. One additional contextual case, pilot-003 (1IEP, Abl kinase with imatinib/STI), needed no retention decision: its one extra component, CL, is a genuine crystallization ion and was stripped under the default policy.

2.5 Docking and reference-pose RMSD

Every prepared receptor with a prepared reference ligand was docked with AutoDock Vina 1.2.5 [8, 9] via the Windows Subsystem for Linux (WSL), fixed seed 20260812, one CPU, exhaustiveness 8, num_modes 9, energy_range 3 kcal/mol, with the search box centered on the experimental ligand's coordinate centroid. Recovery of the experimental pose was assessed with an identity-mapped, heavy-atom RMSD: the RCSB Chemical Component Dictionary (CCD) atom order, the RCSB ideal-geometry Structure-Data File (SDF) topology, and the original deposited mmCIF coordinates were bridged and accepted only if CCD-to-SDF element order, deposited heavy-atom completeness, and isomeric heavy-atom SMILES all agreed between the reference and docked molecules, and the docked pose count matched the parsed Vina score count. RMSD itself was computed with RDKit's rdMolAlign.GetBestRMS [10], which resolves the correct atom correspondence between the two independently built molecule representations via symmetry-aware substructure matching before performing a best-fit rigid superposition; this was checked directly during the audit, since a naive same-index RMSD without correct atom correspondence gives values of 5–8 Å even for the best-recovered cases. This concern — that ligand molecular symmetry requires a genuine, verified atom correspondence rather than an assumed one before RMSD is meaningful — is not specific to this audit: it is the same problem addressed independently by DockRMSD

[11], a graph-isomorphism-based atom-mapping tool published for exactly this purpose. This audit did not use DockRMSD directly, but its existence corroborates that the correspondence step, not only the alignment step, is the part of RMSD calculation most prone to silent error. As a fully independent cross-check of the alignment step itself, the same six identity-mapped correspondences were re-aligned with a from-scratch NumPy implementation of the classical Kabsch algorithm [12], with no call into RDKit's own alignment routine; all six values agreed with the RDKit-reported RMSD to machine precision (≤ 0.001 Å). Reported RMSD values are explicitly distinct from Vina's own internal pose-to-best-pose RMSD, which describes pose spread relative to Vina's own top pose and is not a reference-pose recovery metric.

2.6 Structural descriptors of the top-scoring pose

For every completed docking case, receptor-ligand interface contacts (heavy-atom pairs ≤ 4.5 Å) and hydrogen-bond candidates (donor-acceptor N/O pairs ≤ 3.5 Å) were computed directly from the top-scoring pose's coordinates, in a Python reimplementation of the same algorithm and thresholds already validated in the authors' own research-platform tooling (independent implementation, no shared code). The same contact data were rendered as schematic 2D ligand-receptor interaction diagrams following the layout convention established by LigPlot+ [14] and PLIP [15]. Buried interface surface area was computed with a NumPy-vectorized Shrake-Rupley solvent-accessible surface area (SASA) algorithm [13] (Bondi van der Waals radii, 1.4 Å probe, 200 sample points per atom, deterministic Fibonacci-spiral sphere sampling), as ΔSASA = SASA(receptor alone) + SASA(ligand alone) − SASA(complex). RMSD outcomes were additionally classified against the ≤ 1.0/≤2.0/≤3.0 Å thresholds conventionally used in docking-benchmark literature, with ≤ 2.0 Å the threshold most commonly treated as the de facto "correct pose" cutoff in benchmarking campaigns such as CASF [16]. The corresponding top-scoring-pose interaction diagrams are shown in Fig. 7.

2.7 Reproducibility and computational assistance

Every script, configuration, manifest, and generated figure/table referenced in this manuscript is versioned in a public repository, together with the exact commands required to regenerate every reported number from the raw retrieved structures. Independent reproduction was carried out after the frozen results were recorded: all 23 preparation attempts, all 6 verified plus 1 unverified RMSD calculations, a live RCSB re-download and checksum match for 3 representative structures, and one full Vina re-run were repeated from the versioned inputs; every outcome, including the one unverifiable case, reproduced exactly, and the re-run Vina output was byte-identical to the committed result. The protocol workflow and its reported outcome counts are shown in Fig. 6.

A large language model (LLM) was used as a computational aid during this study, under the authors' direction and review. The human authors reviewed the computational evidence, figures, references, and manuscript material and remain responsible for the methods, interpretation, and final submitted text.

3 Results

3.1 Registry outcome

All 30 registered candidates were retrieved and checksummed from RCSB. Structural classification found 12 clean and 18 contextual candidates; no candidate required manual review outside these two classes. The structural inventory is shown in Fig. 1, and the full evidence flow is summarized in Fig. 4.

3.2 Strict preparation census

Of the 30 registered candidates, 23 were carried through strict preparation (all 12 clean candidates and 11 of the 18 contextual candidates, the latter under an explicit declared policy); the remaining 7 contextual candidates have no declared policy yet and were not attempted. Seven of the 23 attempts (30%) produced a receptor accepted by Meeko without repair, alternate-location selection, or residue deletion (Table 1); one was a clean candidate (pilot-001/1STP) and six were contextual candidates processed under the retained-component policy described in Section 2.4. Sixteen attempts failed. Fifteen of the sixteen failures (94%) involved an unresolved alternate side-chain conformer — ten on the pure alternate-location class and five on a combined alternate-location-and-template-matching class — and the remaining failure was a template-matching mismatch unrelated to alternate locations. The complete case-level preparation outcomes are shown in Fig. 2.

A representative rejection (case expansion-010, PDB 6FMC) is quoted directly from the tool's own output: "Residues with alternate location: [25 residues]. Either specify an altloc for each with option wanted_altloc or a general default altloc with option default_altloc." Meeko offers exactly the two flags this protocol declines to use; selecting a conformer on the analyst's behalf is a repair decision, and the frozen policy exists specifically to measure preparation compatibility without making it.

Table 1
Strict receptor-preparation outcomes by failure class (23 attempts)

Outcome	Count
Prepared	7
Failed — alternate location requires choice	10
Failed — alternate location + template matching failed	5
Failed — template matching failed	1

3.3 The resolution–alternate-location tension

Ten of the 30 registered candidates carry an explicit resolution figure recorded at registration (all sub-1.0 Å, 0.80–0.94 Å). All ten failed strict preparation, and all ten failed on an alternate-location-related cause. This is read as a corroborating observation over the resolution-tagged subset, not a systematic claim over the full registry, since resolution was not captured in a consistent numeric field for every candidate. Within that caveat, the pattern is mechanistically expected rather than coincidental: the same

crystallographic process that resolves sub-angstrom detail is what makes alternate side-chain conformations visible and modelable in the deposited structure in the first place, so a criterion often used as informal shorthand for “high-quality structure” is, under a no-repair preparation policy, associated with lower rather than higher preparability.

3.4 Completed reference-pose cases

The seven prepared receptors were carried through ligand preparation and reference-pose docking; all seven retained all nine requested Vina poses (Table 2). Six cases produced a verified reference-pose RMSD. Three recovered the experimental pose closely: pilot-001/1STP (biotin, top-score/best 0.689 Å), pilot-007/3D4Q (SM5, 0.769 Å), and expansion-001/1B9V (RA2, top-score 0.930 Å, best 0.591 Å). Two recovered it moderately: pilot-008/3CJO (K30, with the retained ADP-Mg cofactor, top-score 1.504 Å, best 1.154 Å) and pilot-002/1HVR (XK2, with the retained CSO residue, top-score 2.493 Å, best 1.664 Å). One did not recover the pose: pilot-006/3PTB (BEN, with the retained structural Ca2+, top-score 5.594 Å, best 5.578 Å, essentially uniform across all nine poses). Classified against conventional success thresholds, five of the six verified cases (83%) fall within ≤ 2.0 Å; 3PTB fails at every threshold tested. The seventh case, pilot-003/1IEP (imatinib/STI), completed docking and produced Vina scores for all nine requested modes, but its output PDBQT file contained only four pose models — a genuine verification failure confirmed independently by direct parsing and by Meeko’s own PDBQT reader, and reported as unverifiable rather than paired or re-run. The RMSD-versus-top-score relationship for the RMSD-verified cases is shown in Fig. 3.

Table 2
Completed reference-pose docking cases (n = 7)

Case	PDB	Ligand	Top-score affinity (kcal/mol)	Top-score RMSD (Å)	Best RMSD (Å)	RMSD verified?
pilot-001	1STP	BTN	−7.244	0.689	0.689	Yes
pilot-007	3D4Q	SM5	−10.960	0.769	0.769	Yes
expansion-001	1B9V	RA2	−7.501	0.930	0.591	Yes
pilot-008	3CJO	K30	−10.890	1.504	1.154	Yes
pilot-002	1HVR	XK2	−6.033	2.493	1.664	Yes
pilot-006	3PTB	BEN	−6.058	5.594	5.578	Yes
pilot-003	1IEP	STI	−12.680	—	—	No (pose-count mismatch)

Two structural descriptors computed directly from the top-scoring pose’s coordinates (Section 2.6) corroborate two of these outcomes without altering the RMSD-based conclusion. For 3PTB, interface-contact analysis places the docked pose’s nearest receptor residues within trypsin’s genuine S1

specificity pocket, and its buried interface area (355.6 Å²) is unremarkable for a ligand of benzamidine's size — consistent with a pose that occupies the correct pocket in a geometrically imprecise orientation, rather than a pose docked to the wrong site entirely. For 1IEP, whose RMSD could not be verified, the docked pose's closest receptor contact is Thr315, the well-characterized Abl-kinase gatekeeper residue; this is reported strictly as a structural observation about contact geometry, with no biological, pharmacological, or clinical inference drawn from it. The pocket-level receptor renderings for all completed cases are shown in Fig. 5.

3.5 Evidence boundary

The findings above are computational observations for the specified tool versions, protocol, and candidate pool. They estimate no binding affinity, biological activity, or clinical relevance for any target, ligand, or residue discussed, including Thr315 in 1IEP and the S1 pocket in 3PTB.

4 Discussion

Preparation failure was the majority outcome in this audit (16 of 23 attempts, 70%) before any biological question was asked, and that number should not be read as evidence the pipeline is broken. Fifteen of the sixteen failures trace to one specific, mechanistically understood cause — an unresolved alternate side-chain conformer that the frozen no-repair policy deliberately declines to resolve on the analyst's behalf — not to a diffuse mixture of unrelated errors. The resolution–alternate-location tension observed here (Section 3.3) is a useful, actionable, and underreported observation for anyone assembling automated no-repair docking pipelines from "clean-looking" high-resolution public structures: resolution alone is not a reliable proxy for preparability, and in a strict no-repair setting the relationship runs the other way, since the same physical process that yields sub-angstrom resolution is what makes alternate conformations both real and unresolved by default. Benchmarks that report accuracy only on cases that already prepared successfully do not surface this relationship, because the failed cases are never part of what gets reported.

This audit's failure taxonomy can be read productively against the DOCK Blaster feasibility study of Irwin et al. [17], one of the few published accounts that reports receptor-preparation outcomes for an automated pipeline rather than only downstream docking scores. That study likewise found that a substantial fraction of candidate receptors could not be carried through an automated route, for reasons rooted in coordinate ambiguity and non-standard residue records rather than in the docking search itself. The present findings sharpen that observation for one specific, currently maintained tool: under a genuinely no-repair policy, essentially all such incompatibility (94%) traces to one identifiable cause — unresolved alternate side-chain conformers — rather than to a broad mixture of structural defects. This distinction has a practical consequence: a pipeline aiming to raise its no-repair success rate does not need to solve receptor preparation in general, it needs a principled, disclosed policy for alternate-location resolution specifically, which is exactly the decision this audit deliberately declines to make so that the underlying incompatibility rate can be measured rather than absorbed.

The seven completed docking cases illustrate a second point that is easy to elide when only successful, well-recovered cases are shown: successful preparation, successful pose recovery, and successful RMSD verification are three independent outcomes. All seven cases were successfully prepared and successfully docked. Six produced a numerically interpretable RMSD; one did not, for a reason (a PDBQT pose-count mismatch) that is itself a reportable, reproducible technical finding rather than a nuisance to discard. Among the six with a verified RMSD, recovery ranged from 0.591 Å to a complete failure to recover the pose at all (3PTB, ~ 5.6 Å); the case-specific structural descriptors in Section 3.4 add independent, coordinate-derived context to that negative result (correct pocket, imprecise orientation) without changing the RMSD-based conclusion itself. With n = 6–7 completed cases this cannot support a general recovery-rate estimate for any tool, protocol, or target class; it establishes only that, within this audit, preparation success does not guarantee pose recovery, and pose recovery does not guarantee that recovery is even numerically verifiable.

Recent guidance on docking reproducibility converges on the same point this audit was designed to demonstrate empirically rather than assert. Martis and Téletchéa's ten-tip guide to meaningful, reproducible docking calculations [18] identifies undocumented receptor-preparation choices as one of the most common sources of irreproducible or misleading docking results, alongside score misinterpretation and inadequate pose validation — concerns echoed in the workflow review cited in the Introduction [2], which lists structure-preparation heterogeneity among the practical bottlenecks limiting reproducibility across current open-source and AI-augmented docking pipelines. The present audit operationalizes that guidance directly for one specific tool and policy, by reporting the preparation step's own outcome distribution — successes and rejections alike — rather than only the docking results that happen to survive it.

A related, independently parameterized audit of 102 DUD-E targets by the same authors is currently under review elsewhere. That study and this one share two PDB identifiers as reference-pose cases (3D4Q and 3CJO) but are otherwise distinct: different registries (a frozen 30-candidate RCSB clean-discovery ranking here, versus 102 DUD-E target/PDB pairs there), different central findings (a compatibility census with case-specific retained-component policy here, versus a three-condition source-and-reader comparison there), and different, independently frozen Vina configurations (exhaustiveness 8, seed 20260812, here; exhaustiveness 16, seed 20260809, there). The two studies report different RMSD values for 3D4Q and 3CJO as a direct consequence of independent runs under independently frozen protocols, and neither value should be read as a replication of, or a discrepancy with, the other.

This audit's candidate pool is small and skewed toward the sub-1-Å resolution stratum of the discovery ranking it was drawn from, because that stratum dominates the ranking; findings describe compatibility with one specific tool version (Meeko 0.7.1) under one specific no-repair policy, not a statement about Meeko, Vina, or receptor-preparation tools in general. Seven of the eighteen contextual candidates remain unprocessed, with no declared retention policy yet; this is reported as future work, not as a gap in the results presented here. Reproducibility is treated as a first-class contribution rather than an

afterthought: every number in this manuscript, including the one negative docking result and the one unverifiable case, can be regenerated from the versioned public repository.

5 Conclusions

This work reports a frozen, no-repair receptor-preparation protocol applied without exception to a declared 30-candidate registry from the RCSB Protein Data Bank, together with every outcome — successes and rejections alike. Thirty percent of 23 preparation attempts succeeded, and 94% of the rejections trace to one specific, mechanistically understood cause: an unresolved alternate side-chain conformer that the protocol deliberately declines to resolve on the analyst's behalf. Within the subset of candidates with an explicit recorded resolution, all ten sub-1.0 Å candidates failed for that same reason, suggesting resolution is not a reliable preparability proxy under a no-repair policy. Seven prepared receptors were carried through reference-pose docking; the results show preparation success, pose-recovery success, and RMSD-verifiability to be three separable outcomes rather than one. Every reported number is reproducible from a versioned, publicly available repository.

Declarations

Competing interests

The authors have no relevant financial or non-financial interests to disclose.

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Author Contribution

A.M.H. and S.L.F.A. contributed to conceptualization. A.M.H. and C.I.M.O. contributed to methodology. A.M.H. performed data curation, formal analysis, investigation, software development, and visualization. All authors contributed to writing, review, and editing, and approved the final manuscript.

Data Availability

All code, pinned workflow protocols, manifests, non-coordinate result tables, computed metrics, and software-version records are publicly available in the Zenodo archive at https://doi.org/10.5281/zenodo.21928743 and the GitHub repository at https://github.com/sircalch/docking-reference-audit. To respect the RCSB PDB coordinate-data policy,

raw coordinate files and intermediate docking logs are not redistributed; the manuscript and archive provide PDB identifiers, source URLs, commands, configurations, filenames, and checksums needed for independent retrieval and reproduction.

References

1. Burley SK, Bhikadiya C, Bi C, Bittrich S, Chao H, Chen L, Craig PA, Crichlow GV, Dalenberg K, Duarte JM, Dutta S, Fayazi M, Feng Z, Flatt JW, Ganesan S, Ghosh S, Goodsell DS, Green RK, Guranovic V, Henry J, Hudson BP, Khokhriakov I, Lawson CL, Liang Y, Lowe R, Peisach E, Persikova I, Piehl DW, Rose Y, Sali A, Segura J, Sekharan M, Shao C, Vallat B, Voigt M, Webb B, Westbrook JD, Whetstone S, Young JY, Zalevsky A, Zardecki C (2023) RCSB Protein Data Bank (RCSB.org): delivery of experimentally-determined PDB structures alongside one million computed structure models of proteins from artificial intelligence/machine learning. Nucleic Acids Res 51:D488–D508. https://doi.org/10.1093/nar/gkac1077
2. Azam F, Almahmoud SA (2026) Open-Source Molecular Docking and AI-Augmented Structure-Based Drug Design: Current Workflows, Challenges, and Opportunities. Int J Mol Sci 27:3302. https://doi.org/10.3390/ijms27073302
3. Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, Blomberg N, Boiten J-W, da Silva Santos LB, Bourne PE, Bouwman J, Brookes AJ, Clark T, Crosas M, Dillo I, Dumon O, Edmunds S, Evelo CT, Finkers R, Gonzalez-Beltran A, Gray AJG, Groth P, Goble C, Grethe JS, Heringa J, ’t Hoen PAC, Hooft R, Kuhn T, Kok R, Kok J, Lusher SJ, Martone ME, Mons A, Packer AL, Persson B, Rocca-Serra P, Roos M, van Schaik R, Sansone S-A, Schultes E, Sengstag T, Slater T, Strawn G, Swertz MA, Thompson M, van der Lei J, van Mulligen E, Velterop J, Waagmeester A, Wittenburg P, Wolstencroft K, Zhao J, Mons B (2016) The FAIR Guiding Principles for scientific data management and stewardship. Sci Data 3:160018. https://doi.org/10.1038/sdata.2016.18
4. Sandve GK, Nekrutenko A, Taylor J, Hovig E (2013) Ten Simple Rules for Reproducible Computational Research. PLoS Comput Biol 9:e1003285. https://doi.org/10.1371/journal.pcbi.1003285
5. Forli Lab (2025) Meeko 0.7.1. Python Package Index. https://pypi.org/project/meeko/0.7.1/
6. Wojdyr M (2022) GEMMI: A library for structural biology. J Open Source Softw 7:4200. https://doi.org/10.21105/joss.04200
7. Bode W, Schwager P (1975) The refined crystal structure of bovine beta-trypsin at 1.8 Å resolution. II. Crystallographic refinement, calcium binding site, benzamidine binding site and active site at pH 7.0. J Mol Biol 98:693–717. https://doi.org/10.1016/S0022-2836(75)80005-2
8. Trott O, Olson AJ (2010) AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. J Comput Chem 31:455–461. https://doi.org/10.1002/jcc.21334
9. Eberhardt J, Santos-Martins D, Tillack AF, Forli S (2021) AutoDock Vina 1.2.0: New Docking Methods, Expanded Force Field, and Python Bindings. J Chem Inf Model 61:3891–3898.

https://doi.org/10.1021/acs.jcim.1c00203
10. Landrum G (2024) RDKit: Open-source cheminformatics. https://www.rdkit.org. https://doi.org/10.5281/zenodo.591637
11. Bell EW, Zhang Y (2019) DockRMSD: an open-source tool for atom mapping and RMSD calculation of symmetric molecules through graph isomorphism. J Cheminform 11:40. https://doi.org/10.1186/s13321-019-0362-7
12. Kabsch W (1976) A solution for the best rotation to relate two sets of vectors. Acta Crystallogr A 32:922–923. https://doi.org/10.1107/S0567739476001873
13. Shrake A, Rupley JA (1973) Environment and exposure to solvent of protein atoms. Lysozyme and insulin. J Mol Biol 79:351–371. https://doi.org/10.1016/0022-2836(73)90011-9
14. Laskowski RA, Swindells MB (2011) LigPlot+: multiple ligand-protein interaction diagrams for drug discovery. J Chem Inf Model 51:2778–2786. https://doi.org/10.1021/ci200227u
15. Salentin S, Schreiber S, Haupt VJ, Adasme MF, Schroeder M (2015) PLIP: fully automated protein-ligand interaction profiler. Nucleic Acids Res 43:W443–W447. https://doi.org/10.1093/nar/gkv315
16. Su M, Yang Q, Du Y, Feng G, Liu Z, Li Y, Wang R (2019) Comparative Assessment of Scoring Functions: The CASF-2016 Update. J Chem Inf Model 59:895–913. https://doi.org/10.1021/acs.jcim.8b00545
17. Irwin JJ, Shoichet BK, Mysinger MM, Huang N, Colizzi F, Wassam P, Cao Y (2009) Automated Docking Screens: A Feasibility Study. J Med Chem 52:5712–5720. https://doi.org/10.1021/jm9006966
18. Martis EAF, Téletchéa S (2025) Ten quick tips to perform meaningful and reproducible molecular docking calculations. PLoS Comput Biol 21:e1013030. https://doi.org/10.1371/journal.pcbi.1013030

Figures

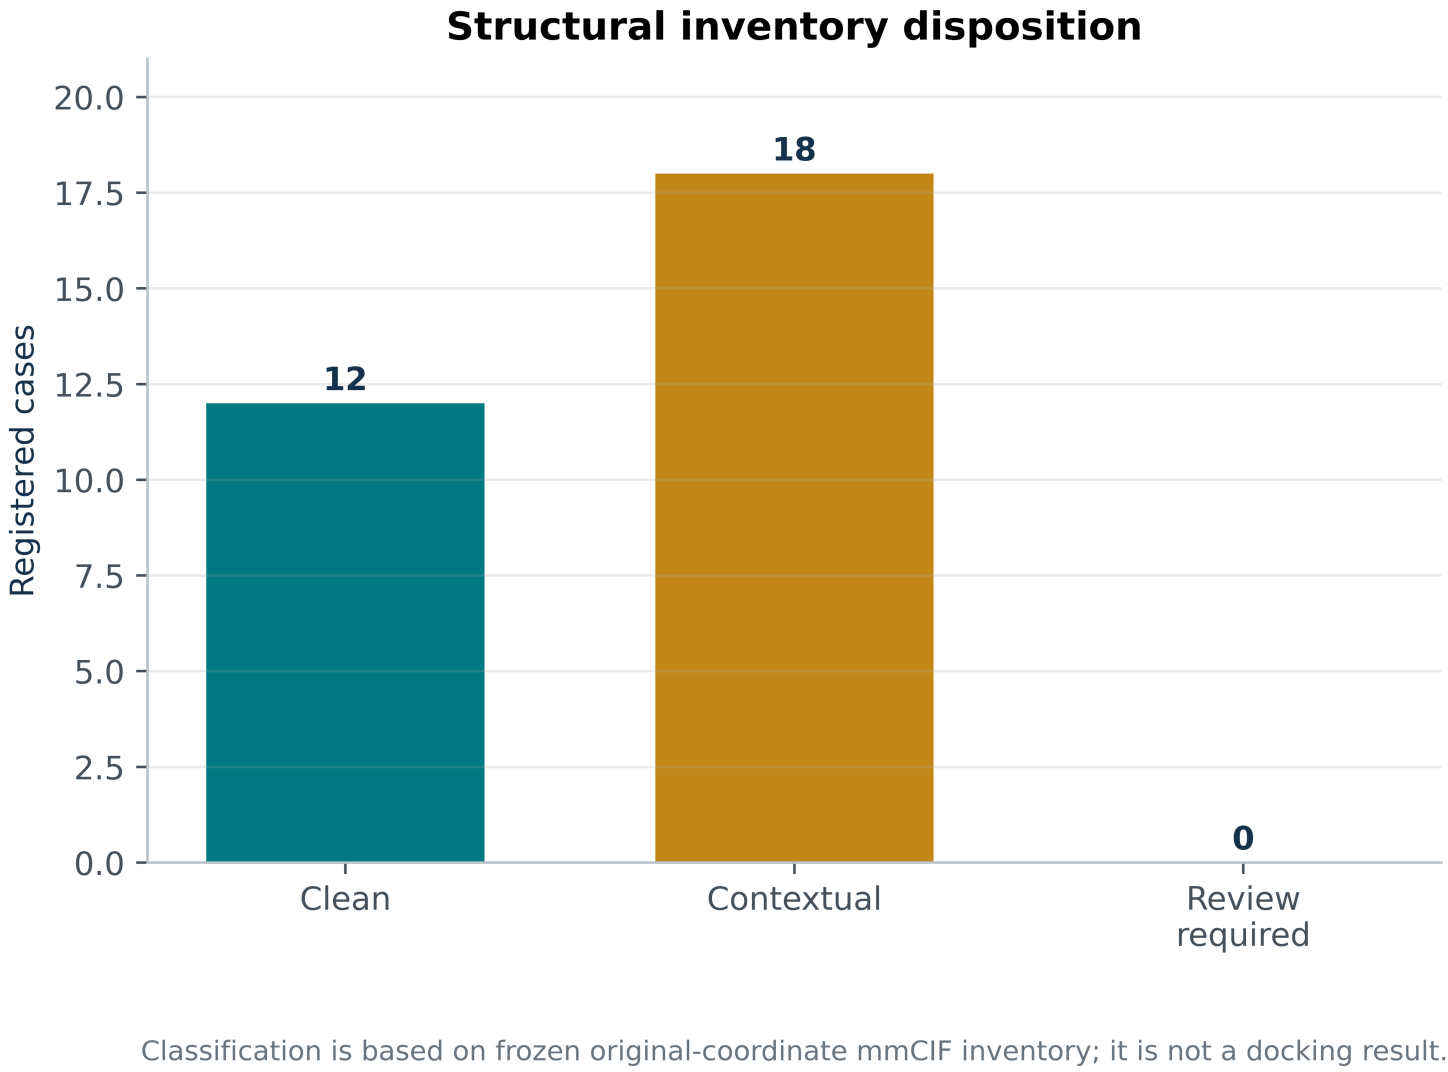


Figure 1

Structural inventory of the 30 registered candidates by classification stratum (clean vs. contextual)

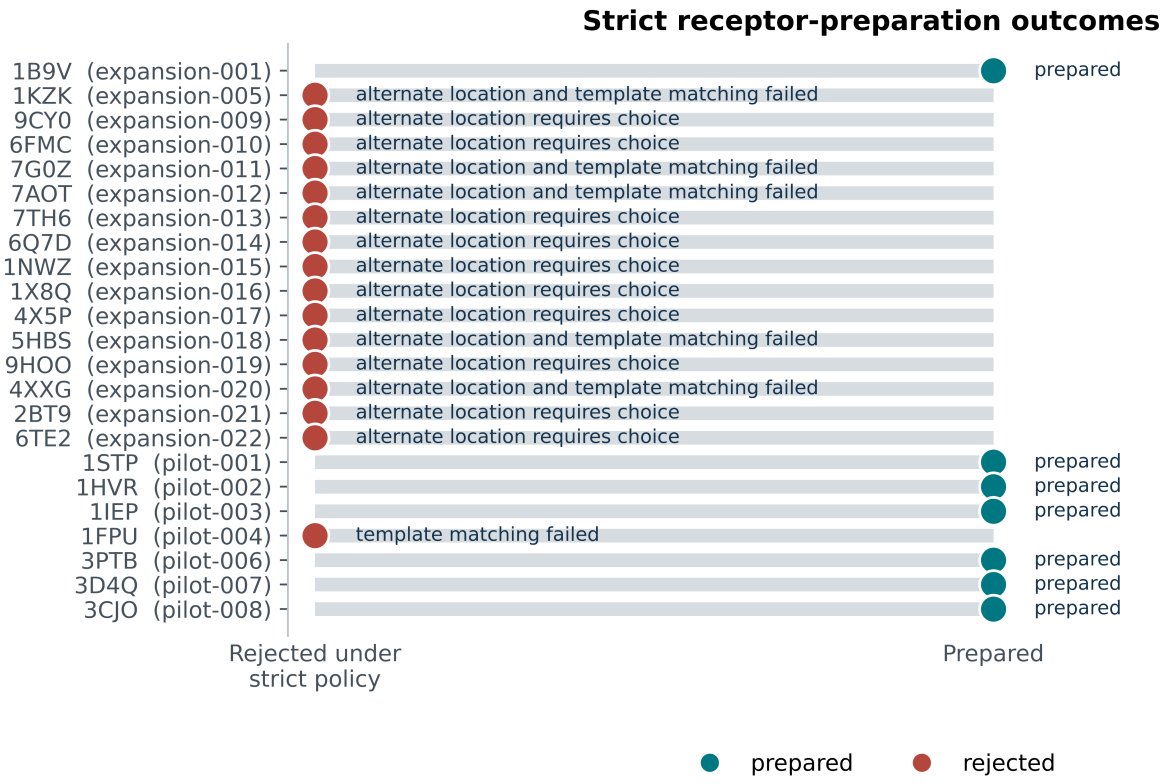
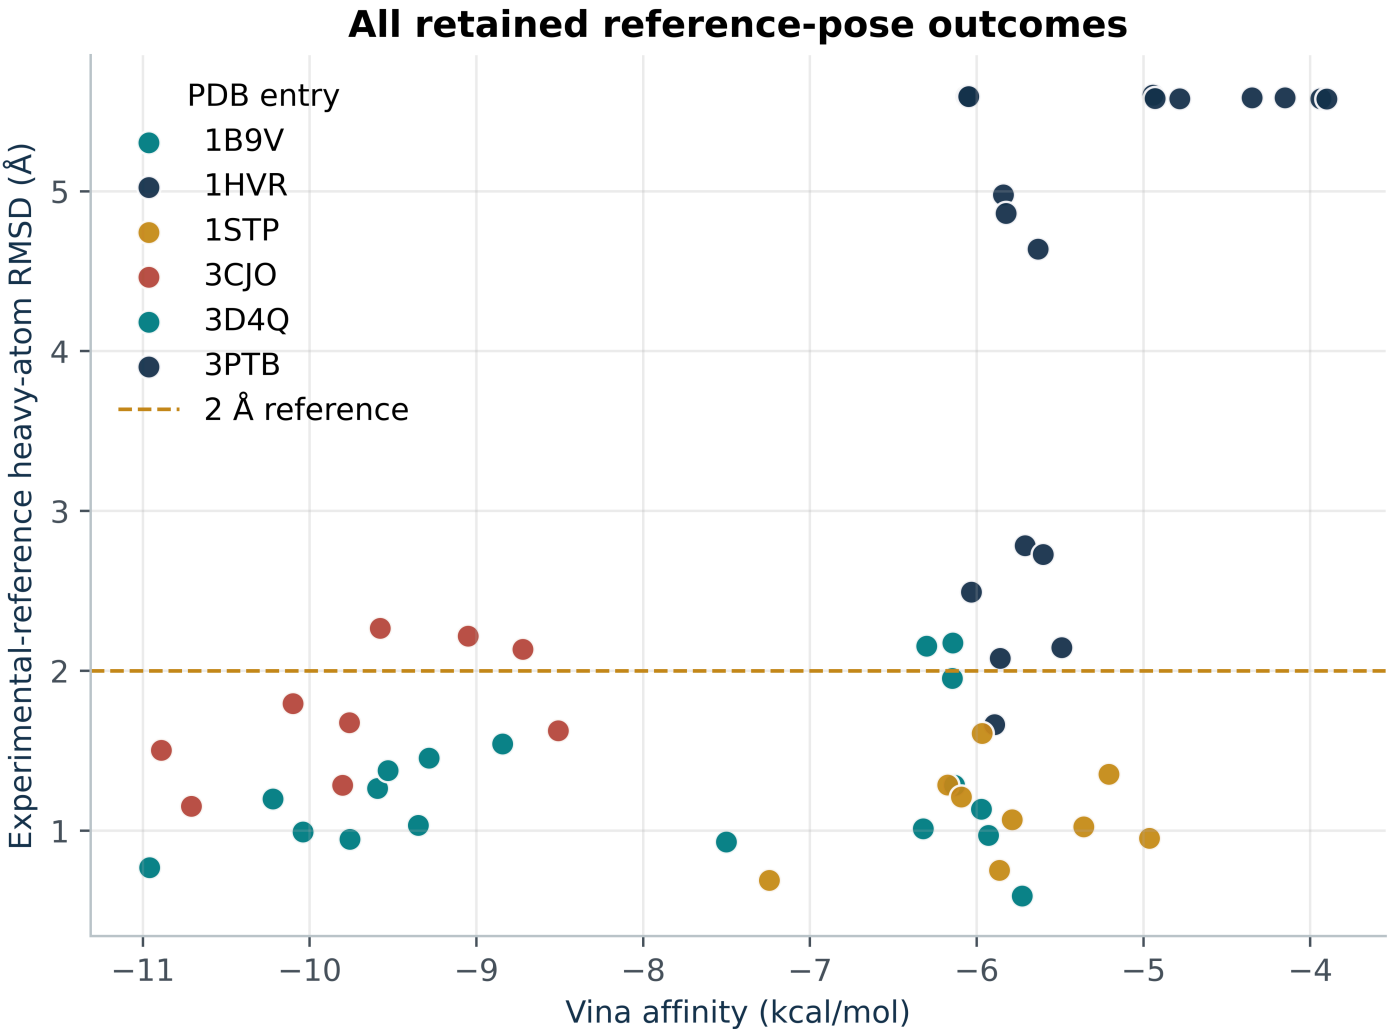


Figure 2

Strict receptor-preparation outcomes for all 23 attempted candidates, labeled by failure class where applicable



Each point is one retained Vina pose. RMSD is identity-mapped and aligned to the experimental ligand pose.

Figure 3

Reference-pose RMSD versus Vina top-score affinity for all six RMSD-verified completed cases

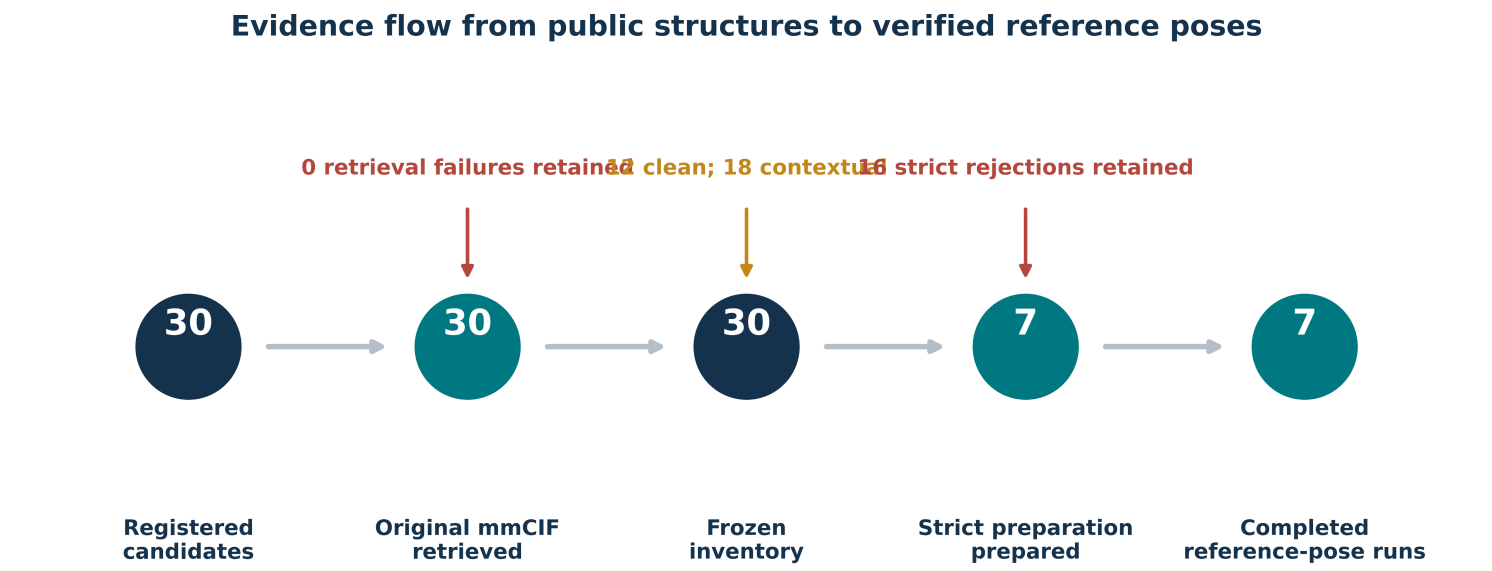


Figure 4

Evidence flow from candidate registration through classification, preparation, and completed docking

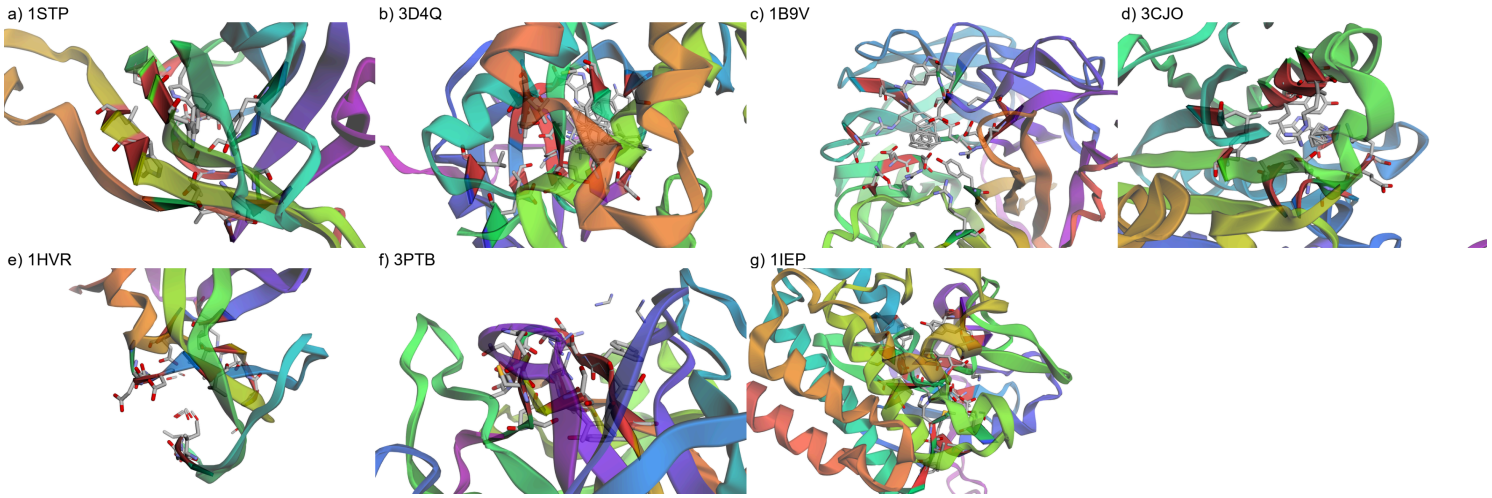


Figure 5

Receptor cartoon renders (3Dmol.js), zoomed to the binding pocket, with interface-contact residues highlighted (green: hydrogen-bond candidate; red: other contact ≤4.5 Å), for each of the seven completed docking cases: (a) 1STP, (b) 3D4Q, (c) 1B9V, (d) 3CJO, (e) 1HVR, (f) 3PTB, (g) 1IEP

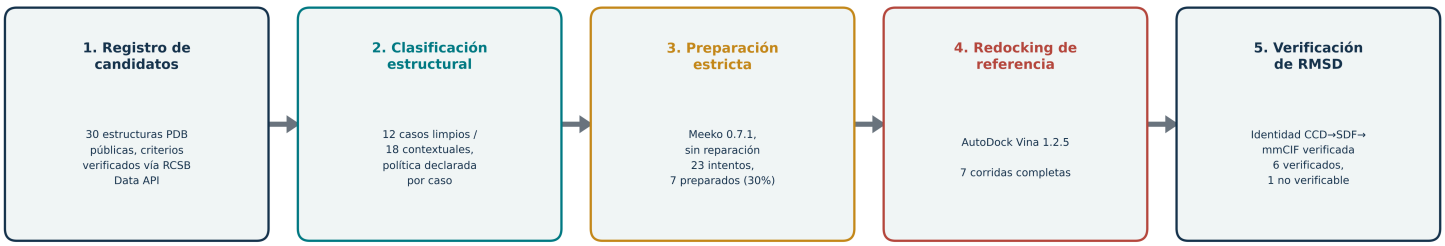


Figure 6

Protocol workflow diagram: candidate registration, structural classification, strict preparation, reference-pose docking, and RMSD verification, each stage annotated with its reported outcome count

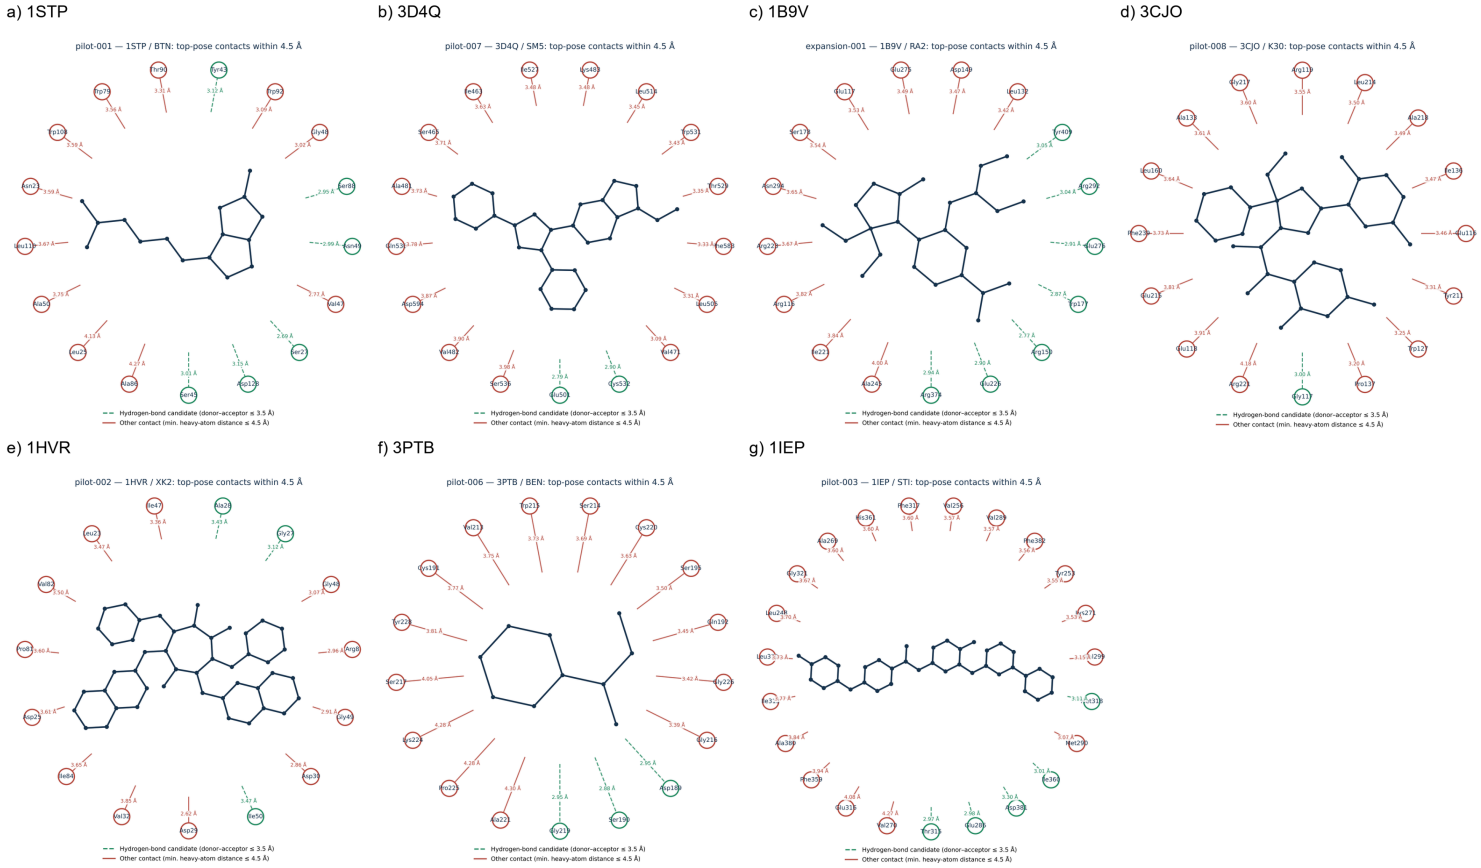


Figure 7

Schematic 2D ligand-receptor interaction diagrams (LigPlot+/PLIP-style layout) for the top-scoring pose of each of the seven completed docking cases: (a) 1STP, (b) 3D4Q, (c) 1B9V, (d) 3CJO, (e) 1HVR, (f) 3PTB, (g) 1IEP