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.

A blind-docking server update reorders an HDAC8 ranking of spiro[indoline-tetrazolo[1,5-c]quinazolines]: metal handling, pose accuracy and site selection

Lyudmyla Antypenko

antypenko1@gmail.com

Oles Honchar Dnipro National University

Oleksii Antypenko

Zaporizhzhia State Medical and Pharmaceutical University

Yuliia Kozak

National Academy of Sciences of Ukraine

Volodymyr Salionov

Zaporizhzhia State Medical and Pharmaceutical University

Mykola Maletskyi

Zaporizhzhia State Medical and Pharmaceutical University

Hanna Lysianska

Zaporizhzhia State Medical and Pharmaceutical University

Serhii Kovalenko

Oles Honchar Dnipro National University

Mieko Arisawa

Kyushu University

Article

Keywords: blind docking, scoring-function reproducibility, metal-aware docking, HDAC8, ligand efficiency, spiro[indoline-tetrazolo[1,5-c]quinazoline]

Posted Date: November 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11637897/v1

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

Additional Declarations: No competing interests reported.

A blind-docking server update reorders an HDAC8 ranking of spiro[indoline-tetrazolo[1,5-*c*]quinazolines]: metal handling, pose accuracy and site selection

Lyudmyla Antypenko[1,*], Oleksii Antypenko[2], Yuliia Kozak[3], Volodymyr Salionov[2], Mykola Maletskyi[2], Hanna Lysianska[2], Serhii Kovalenko[1], Mieko Arisawa[4]

[1]Institute of Chemistry and Geology, Oles Honchar Dnipro National University, 72 Nauky Ave., 49010, Dnipro, Ukraine
[2]Department of Pharmaceutical, Organic and Bioorganic Chemistry, Department of Biological Chemistry, Department of Medicines Technology, Zaporizhzhia State Medical and Pharmaceutical University, 26 Maria Prymachenko Blvd., 69035, Zaporizhzhia, Ukraine
[3]Institute of Cell Biology of the National Academy of Sciences of Ukraine, 14/16 Drahomanov Str., 79005, Lviv, Ukraine; Molecular Design Center, Danylo Halytsky Lviv National Medical University, 69 Pekarska Str., 79010, Lviv, Ukraine
[4]Department of Biosciences and Biotechnologies, Graduate School of Bioresources and Bioenvironmental Sciences, Kyushu University, 744 W5-674, Motooka Nishi-ku, Fukuoka, 819-0395, Japan

Correspondence to: antypenkol@gmail.com

Abstract

Blind-docking servers are widely used to rank candidate ligands, and the resulting order is usually reported without the server version. Fifteen ligands — nine spiro[indoline-tetrazolo[1,5-*c*]quinazoline] derivatives, the parent scaffold and five reference compounds — were docked into one HDAC8 structure (PDB 2V5X) under two successive versions of the same server, with ligand input, cavity definitions and search boxes held constant. The rankings do not agree: Spearman $\rho = +0.23$ ($p = 0.42$), root-mean-square difference 2.09 kcal/mol, and shifts of up to thirteen rank positions. A controlled study across eight receptor variants, cut from the same coordinates by deletion of records alone, traces this to retention of heteroatoms in the receptor rather than to any revision of the search or scoring algorithm: deleting the catalytic zinc restores the earlier cavity and makes the score less favourable for eight of ten ligands, by a mean of 1.28 kcal/mol against a run-to-run variation of 0.37, while the same version run on the zinc-free receptor reproduces the earlier one to 0.20 kcal/mol, which bounds the contribution of the engine and template changes. Redocking the co-crystallised ligand shows the better-scoring pose to be the less accurate one under both versions, and the catalytic site proves detectable only through a crystal-packing contact, the deposited entry declaring two monomeric assemblies and published solution measurements of the enzyme being compatible with a monomer. Cytotoxicity and ADMET predictions are new; no HDAC8 inhibition is claimed.

Keywords: blind docking; scoring-function reproducibility; metal-aware docking; HDAC8; ligand efficiency; spiro[indoline-tetrazolo[1,5-*c*]quinazoline]

Introduction

Histone deacetylases remove acetyl groups from lysine residues in histones and non-histone proteins, and aberrant activity contributes to cancer, neurodegenerative disorders and autoimmune disease[1]. Among the eighteen human isoforms, HDAC8 has attracted attention for its characteristic L1 loop, its comparatively open substrate-binding channel[2,3] and its association with Cornelia de Lange syndrome[4]; among its non-histone substrates are SMC3 and estrogen-related receptor alpha[5]. The decisive feature here is catalytic: HDAC8 is a zinc metalloenzyme whose Zn^{2+} ion, coordinated by Asp178, His180 and Asp267, is directly engaged by most established inhibitors and has a catalytic rather than a merely structural role[6]. A docking protocol applied to HDAC8 must therefore make a choice, explicit or otherwise, about how the metal is represented. That choice is the subject of this paper.

HDAC inhibitors are conventionally classified by their zinc-binding group[1,7]. Hydroxamates chelate the metal potently but show limited isoform selectivity[8–10]; benzamides such as entinostat achieve class I selectivity without HDAC8 specificity[11]; cyclic peptides such as romidepsin and apicidin face comparable selectivity challenges[12,13]; PCI-34051 and NCC-149 are selective for HDAC8[14,15]; non-hydroxamate triazoloquinolines show confirmed target engagement[16]; and ligands with no zinc-chelating group have been proposed to occupy a secondary site in an inverted orientation[17]. A reference set drawn from these classes spans chelating and non-chelating ligands across a wide range of size, so a protocol that handles the metal differently should reorder it predictably.

The compounds examined here were not designed de novo against HDAC8. The spiro[indoline-tetrazolo[1,5-*c*]quinazoline] series comes from an in-house library assembled during earlier work on tetrazolo- and triazolo[1,5-*c*]quinazolines[18–23]; four compounds (KB-268, KB-270, KB-271 and KB-292) were synthesised and characterised there[22] and the remainder are analogues examined computationally only. No compound is described as novel. The series shares no zinc-binding group with the reference compounds of Fig. 1 — a liability for intrinsic affinity, but here an advantage, since a non-chelating series docked alongside ligands that depend on metal coordination isolates the effect of the metal-handling protocol: the two groups should move differently if the protocol changes.

Blind-docking web servers are routine instruments in medicinal chemistry, and the rankings they return are frequently reported without recording which version produced them. CB-Dock2[24,25] combines curvature-based cavity detection with AutoDock Vina scoring and a template-based FitDock pipeline. Its successor CB-Dock3[26] retains that cavity-detection algorithm unchanged but adds a refined engine, an expanded template library and a selective heteroatom-preservation mechanism: earlier versions routinely excluded non-amino-acid HETATM records, whereas CB-Dock3 allows functional metal ions to be retained and scores them through metal-aware terms added to Vina 1.2.0 and FitDock, raising the reported CASF-2016 success rate from 56.8% to 67.4% at RMSD ≤ 2.0 Å[26]. For a zinc metalloenzyme this is a change in what the receptor is taken to be.

Two independent findings sharpen the question. Across eleven DUD-E targets docked under AutoDock Vina into site-directed and then blind boxes, ranking degraded only modestly while placement collapsed, the proportion of actives whose top pose lay within 8 Å of the crystallographic site falling from 96% to 48% ($p = 0.002$); the two quantities were unrelated across targets ($r = -0.03$),

and HDAC8 was among those that lost placement severely, retaining a ROC-AUC of 0.66 while placing 43% of actives in site[27].

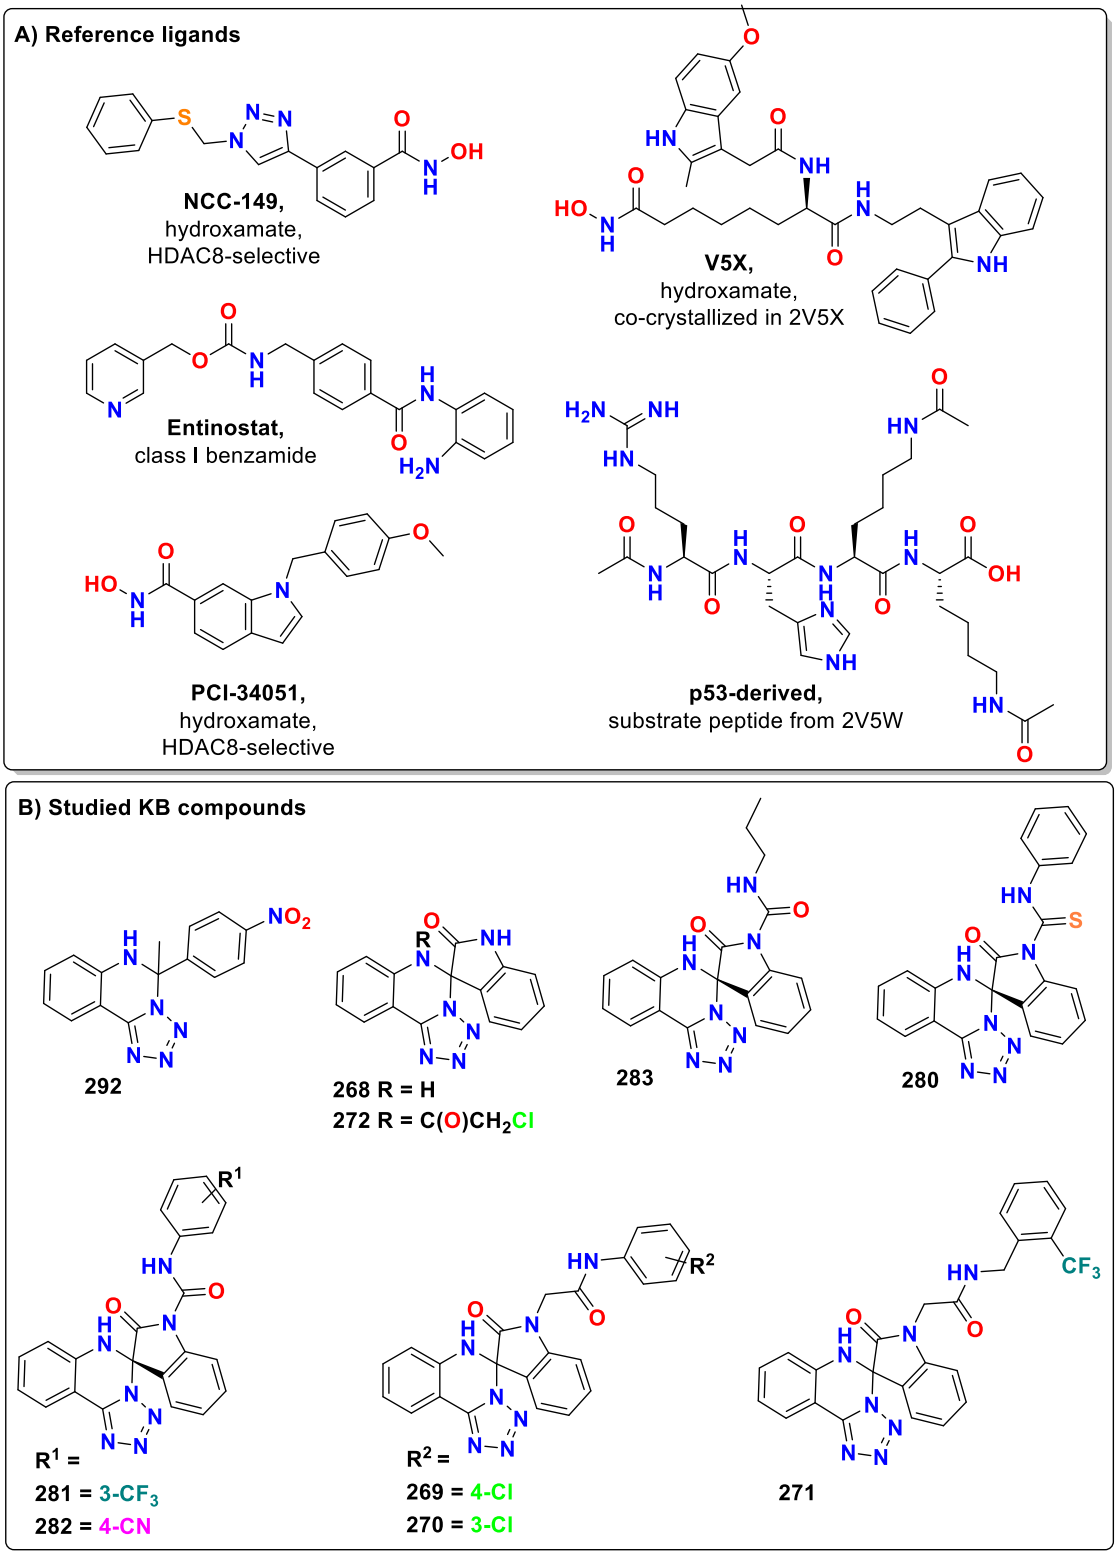

Fig. 1 Molecular structures of the compounds examined. A) Reference HDAC8 ligands: V5X (the hydroxamate co-crystallized in 2V5X), NCC-149, entinostat, PCI-34051, and the p53-derived peptide substrate. B) Studied KB compounds: spiro[indoline-tetrazolo[1,5-*c*]quinazoline] derivatives (KB-268–

KB-272, KB-280–KB-283) and the parent scaffold KB-292. The reference compounds are shown as pharmacological context and as scale anchors for the docking calculations; they share no zinc-binding fragment with the KB series. Of the KB compounds shown, KB-268, KB-270, KB-271 and KB-292 were reported previously[22]; the remainder are designed structures examined computationally only (Methods).

Separately, an audit of 236 open-access docking papers on the SARS-CoV-2 main protease found that only 8.1% reported enough detail to be re-executed[28]; that analysis is a preprint and has not been peer reviewed, but a ranking that cannot be re-executed cannot be asked to survive a server update.

We therefore docked fifteen ligands into one HDAC8 crystal structure (PDB 2V5X[29]) under both versions with ligand input, cavity definitions and search boxes held constant, used the co-crystallised ligand as an internal pose control, and then varied the receptor file alone under a single version to test whether metal handling accounts for the difference. Chemical names are given in Supplementary Table S34.

Results and discussion

Reproducibility of the ranking between server versions

The design of the comparison, including the points at which the two arms differ, is summarised in Supplementary Fig. S1. CurPocket returned five cavities under both versions, closely but not identically defined (Methods), Pocket 1 corresponding to the catalytic site. Of the seventy-five ligand–cavity runs, two returned a positive score on first submission — KB-272 in Pocket 1 (+0.3 kcal/mol, an interpenetration with 281 contacts over twelve residues) and PCI-34051 in Pocket 3 (+1.5) — and both converged on resubmission of identical input, to −9.0 and −7.2 kcal/mol (Supplementary Table S3). That a search should fail grossly and then succeed on identical input is itself a reproducibility observation. One statement of the earlier analysis[30] also requires correction: under CB-Dock2 Pocket 1 returned the most favourable score for every compound, but under CB-Dock3 Pocket 2 outscores it for KB-280 (−10.1 against −9.9 kcal/mol) and KB-292 (−8.2 against −8.0). Pocket 1 scores under both versions are given in Table 1, and the scores across all five cavities in Supplementary Fig. S3.

Table 1 Pocket 1 binding scores under both server versions

Compound	CB-Dock2 (kcal/mol)	CB-Dock3 (kcal/mol)	Δ	Heavy atoms	LE (CB-Dock2)	LE (CB-Dock3)	Zn distance (Å)
p53-derived	−8.4	−13.7	−5.3	49	0.171	0.280	2.00
KB-282	−12.3	−13.1	−0.8	33	0.373	0.397	2.03
KB-270	−10.0	−12.9	−2.9	33	0.303	0.391	3.73
Entinostat	−8.4	−12.7	−4.3	28	0.300	0.454	1.90
KB-269	−10.5	−12.2	−1.7	33	0.318	0.370	3.12
V5X	−10.8	−12.1	−1.3	46	0.235	0.263	2.06
KB-271	−10.4	−10.9	−0.5	37	0.281	0.295	6.25
KB-280	−10.4	−9.9	+0.5	31	0.335	0.319	3.12
NCC-149	−8.9	−9.6	−0.7	23	0.387	0.417	2.36
KB-281	−11.0	−9.2	+1.8	35	0.314	0.263	6.71
KB-272	−8.5	−9.0	−0.5	26	0.327	0.346	4.39

PCI-34051	−8.8	−9.0	−0.2	22	0.400	0.409	2.18
KB-268	−8.8	−8.4	+0.4	22	0.400	0.382	5.41
KB-283	−8.5	−8.4	+0.1	28	0.304	0.300	4.12
KB-292	−8.7	−8.0	+0.7	23	0.378	0.348	2.57

Heavy-atom counts, ligand efficiencies and the minimum ligand-to-zinc distance in the CB-Dock3 pose are given alongside. Compounds are ordered by CB-Dock3 score. LE, ligand efficiency (negative score per heavy atom). Δ, CB-Dock3 score minus CB-Dock2 score; a negative value denotes a more favourable score under CB-Dock3. The CB-Dock3 entries for KB-272 and PCI-34051 are the converged resubmissions described in the text; the superseded values are given in Supplementary Table S3. Zinc distances are the minimum ligand-heavy-atom-to-Zn^{2+} distance across both chains; per-chain values are given in Supplementary Table S1.

The rankings do not agree. Across the fifteen ligands the Pearson correlation between the two score sets is r = +0.37 (p = 0.17), the Spearman rank correlation ρ = +0.23 (p = 0.42) and Kendall's τ = +0.23 (p = 0.25); none approaches significance. The mean signed change is −0.98 kcal/mol (SD 1.91) and the root-mean-square difference 2.09 kcal/mol, of the same order as the entire 3.9 kcal/mol spread of the CB-Dock2 scores. Rank positions are quoted for the full fifteen-ligand set unless stated otherwise, as tied minimum ranks. Displacements reach thirteen places over the set and eight for the largest-moving member of the spiro series: KB-281, ranked second both overall and within the series under CB-Dock2 and nominated there for synthesis, falls to tenth overall and sixth within the series, while KB-270 rises from seventh to third, entinostat from joint last to fourth and the p53-derived peptide from joint last to first. The change is shown ligand by ligand in Fig. 2.

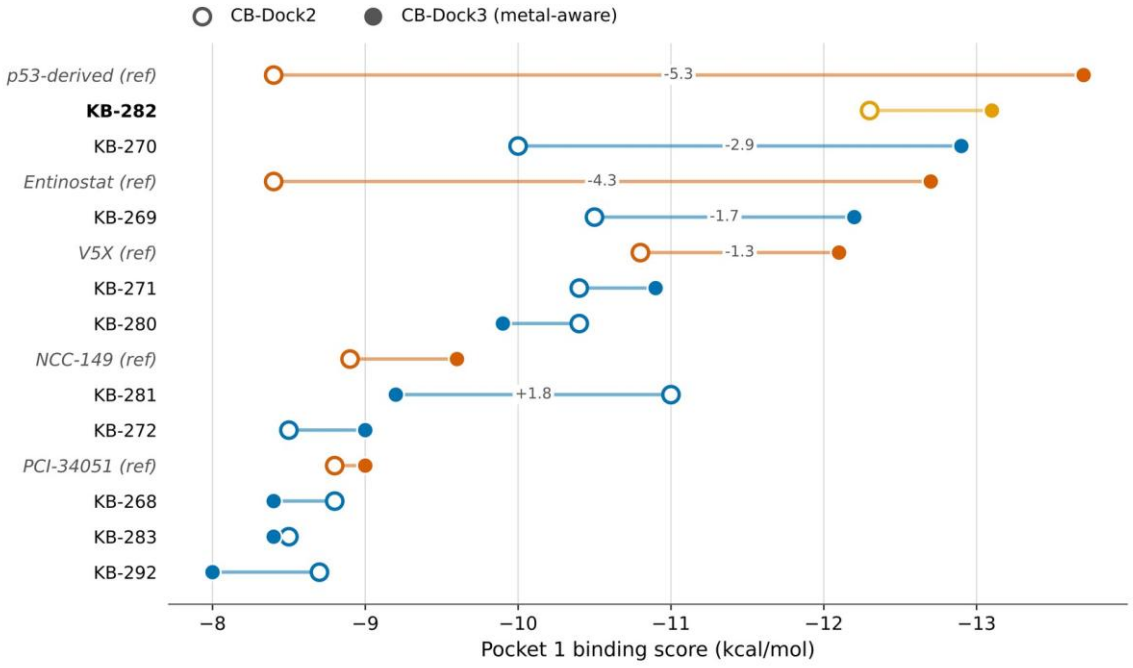

Fig. 2 Pocket 1 binding scores under CB-Dock2 and CB-Dock3. Paired plot for the fifteen ligands, ordered by CB-Dock3 score, with connecting lines indicating the change between versions. Open symbols give the CB-Dock2 score and filled symbols the CB-Dock3 score; reference compounds are distinguished from the KB series by colour and by italic labels, and KB-282 is highlighted. Changes of at least 1.0 kcal/mol are annotated. Data from Table 1.

The practical consequence is direct. The same prioritisation exercise would have nominated KB-282 and KB-281 for synthesis under CB-Dock2 and nominates KB-282 and KB-270 under CB-Dock3; since KB-270 has already been synthesised[22] while KB-281 has not, the two versions imply materially different and differently costly experimental programmes. Nor is the instability confined to the structure-based route: FitDock scores change between versions for every reference ligand for which both are available, and for entinostat the ordering of the two retrieved templates reverses outright (Supplementary Table S18).

Pose accuracy assessed against the internal reference

Because the receptor contains its own ligand in a known pose, the accuracy of each docking mode can be measured rather than inferred. V5X was redocked into its own receptor and the poses compared with the deposited coordinates over 46 heavy atoms (Methods). The two modes give opposite outcomes. Under CB-Dock3 the template-based FitDock pose reproduces the crystallographic conformation at 0.90 Å, within the conventional 2.0 Å criterion, whereas the structure-based Vina pose deviates by 3.55 Å; the alternative atom mapping gives 3.62 Å, so the failure is not an artefact of the symmetry treatment, and 2.98 Å remains after optimal superposition. The FitDock pose nevertheless carries the less favourable score (−8.0 kcal/mol) and the Vina pose the more favourable one (−12.1). The same relation held under CB-Dock2, and more strongly: 0.78 Å at −8.6 kcal/mol against 8.70 Å at −10.8, so the inversion was present in the original calculation, where it went unexamined. The template-based route is stable between versions, and the metal-aware protocol improves structure-based accuracy from 8.70 to 3.55 Å without bringing it within the criterion.

The deposited structure supplies an internal scale: its two independent ligand copies differ by 4.79 Å once the monomers are superposed, so the criterion is applied to a ligand whose own crystallographic copies disagree by more than twice the threshold. The disagreement lies in a distal substituted ring, the nine atoms nearest the metal agreeing to within 0.6 Å; restricted to that reproducible portion the template-based pose deviates by 0.57 Å and the structure-based pose by 1.75 Å (Supplementary Table S4).

In the one case where the correct answer is known, the more favourably scored pose is the less accurate one, under both versions. The comparative ranking above rests on structure-based scores, since template-based scores exist only for ligands with database templates, and that pose is wrong by more than twice the acceptance threshold for the one ligand that can be checked. A correction to the earlier analysis follows[30]: the values reported there as redocking RMSD were taken from the pocket RMSD column of the template panel, which measures query-to-template pocket similarity and serves as a template-selection filter[24]; they carry no information about pose accuracy and are withdrawn.

Isolating the metal term

The version-to-version differences are not randomly distributed across the ligand set; they correspond to the metal-handling protocol, and the correspondence can be shown on the poses rather than inferred from the scores. Under CB-Dock3 the zinc-binding groups of the reference ligands engage the catalytic ion at coordination distance: the hydroxamate oxygen of V5X at 2.06 Å, reproducing the bidentate arrangement observed crystallographically[29]; the 2-aminophenyl nitrogen of entinostat at 1.90 Å[31]; the acetyl-lysine carbonyl nitrogen of the p53-derived peptide at 2.00 Å; and hydroxamate oxygens of

NCC-149 and PCI-34051 at 2.36 and 2.18 Å. None was present in the corresponding CB-Dock2 structure-based poses, where the closest ligand atom lay between 2.60 and 5.35 Å. The template-based arm already placed a hydroxamate oxygen 2.61 Å from the metal under CB-Dock2, so the absence of coordination is a property of the structure-based route on which the ranking rests.

The score changes follow this pattern in direction though not proportionally. Across the fifteen ligands the change correlates with the ligand-to-metal distance of the CB-Dock3 pose ($r = +0.521$, $p = 0.046$; $\rho = +0.606$, $p = 0.017$), a positive coefficient denoting a less favourable score at greater distance; the association rests on two ligands, since excluding the peptide gives $r = +0.481$ ($p = 0.081$) and excluding entinostat as well $r = +0.398$ ($p = 0.178$), so the coefficients describe the trend rather than test it. All five reference species reach the coordination band but only three gain appreciably — the peptide by 5.3 kcal/mol, entinostat by 4.3 and V5X by 1.3 — while NCC-149 and PCI-34051 change by −0.7 and −0.2, and KB-270 and KB-269 gain 2.9 and 1.7 kcal/mol from poses 3.73 and 3.12 Å from the ion. Proximity governs the direction for most of the set without determining the magnitude. The correspondence is plotted in Fig. 3.

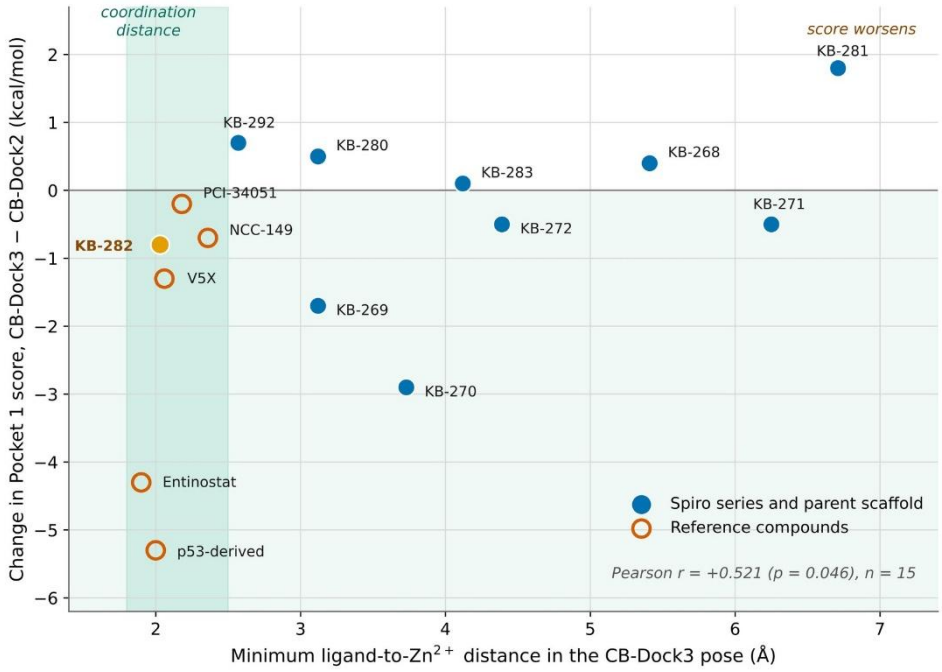

Fig. 3 Minimum ligand-to-zinc distance against the change in binding score between server versions. Each point represents one ligand; the abscissa gives the closest ligand-heavy-atom-to-Zn^{2+} distance in the CB-Dock3 pose and the ordinate the score change Δ from Table 1. Ligands establishing contacts at coordination distance (below 2.5 Å) include the largest score improvements, but the distance does not determine the magnitude of the change. The vertical band marks the 1.8–2.5 Å range conventionally associated with Zn–N and Zn–O coordination, and the shaded lower region the ligands whose score improves under CB-Dock3.

That correspondence is circumstantial, since CB-Dock3 also differs in engine and template library. A separate experiment replaces the inference with a direct manipulation: eight receptor variants were cut from the same deposited coordinates by deletion of records only, so that every variant shares

one coordinate frame, and ten ligands were submitted to each under CB-Dock3 alone with engine, template library and scoring function held fixed (Methods).

The first result concerns cavity detection and identifies the origin of the cavity difference between the two published runs. Deleting the catalytic zinc from the dimeric receptor under CB-Dock3 returns the cavity reported by CB-Dock2: 2020 Å³ centred at (16.465, −0.699, 43.607), reproducing the CB-Dock2 cavity volume exactly and its centre to three decimal places in all three coordinates, against 1887 Å³ at (16.334, −0.699, 43.607) when the zinc is retained. Since CurPocket is unchanged between versions[26], the simplest account is the one the developers describe: CB-Dock2 excluded HETATM records before cavity detection and never saw the zinc, which displaces 133 Å³ of detected volume. The cavity difference is therefore not an algorithmic revision but a consequence of what the receptor was taken to contain.

The second result is the metal effect itself. On the dimeric receptor, deleting the zinc makes the Pocket 1 score less favourable for eight of the ten ligands, by a mean of +1.28 kcal/mol (SD 1.98, range −1.4 to +4.8). The effect appears to decline with distance from the ion: the rank correlation reaches conventional significance ($\rho = -0.685$, $p = 0.029$) while the linear coefficient does not ($r = -0.535$, $p = 0.111$). Grouped, the six ligands within 2.5 Å lose a mean of +2.02 kcal/mol against +0.18 for the four beyond, though with ten points that contrast is not significant ($U = 18.0$, $p = 0.257$). Two ligands run against the trend, KB-281 at 6.71 Å improving by 1.4 kcal/mol and KB-292 at 2.57 Å by 0.6 (Fig. 4C). The p53-derived peptide is the informative case: it carries the largest version-to-version change in the published set and loses +4.3 kcal/mol when the zinc is deleted, so the ligand on which the central comparison most depends is also among the most metal-dependent.

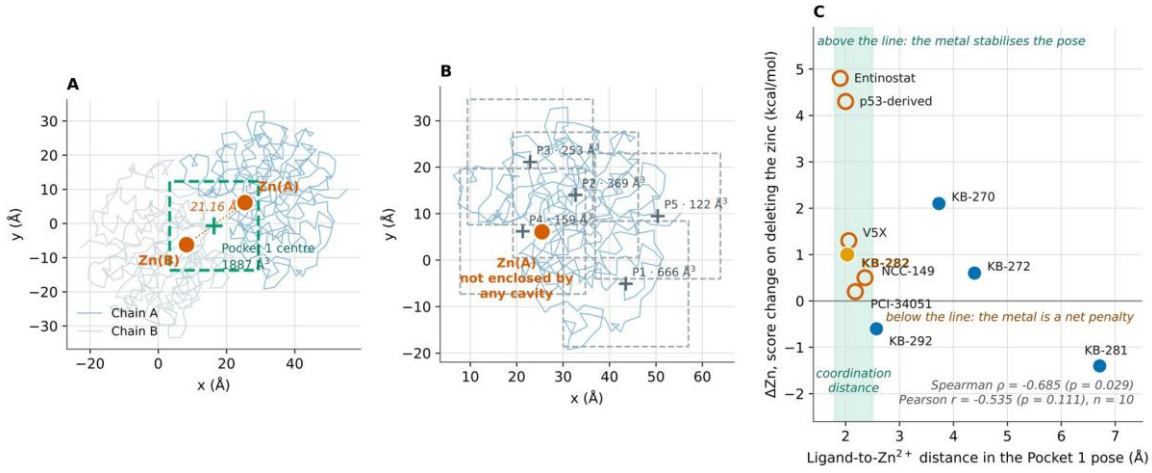

Fig. 4 The catalytic zinc and the crystallographic interface, examined under CB-Dock3 alone. (A) The dimeric receptor viewed down z, with the Pocket 1 search box drawn as a dashed outline. The box encloses both catalytic zinc ions, which lie 21.2 Å apart, so a single cavity definition of 1887 Å^3 covers the catalytic sites of both chains. (B) The chain A monomer, with all five detected cavities and their search boxes, drawn in projection down z. No cavity encloses the catalytic zinc in three dimensions; the Pocket 4 box overlaps the projected position of the ion, from which its centre lies 17.46 Å away, almost all of it along z. The largest cavity, Pocket 1 at 666 Å^3, is centred 21.6 Å from the metal. The same holds for the chain B monomer and for every metal composition tested (Supplementary Table S17). (C) Change in Pocket 1 score on deleting the catalytic zinc, plotted against the minimum ligand-to-zinc

distance of the Pocket 1 pose of the main CB-Dock3 calculation (Table 1); positive values denote a less favourable score once the metal is removed. Reference ligands are drawn as open symbols and KB-282 is highlighted, and the shaded band marks coordination distance. Data from Supplementary Tables S15 and S17.

A control bounds the run-to-run variation the original design could not. The two K^+ ions of the chain A monomer lie more than 8 Å outside the Pocket 1 search box, so arms differing only in their presence pose an effectively identical search problem; pooled over both pairs (n = 20) the mean absolute difference is 0.37 kcal/mol, the largest being 2.5 for entinostat. A reproduction arm repeating the published CB-Dock3 calculation reproduces Pocket 1 to a mean absolute difference of 0.21 kcal/mol. Repeated submission therefore moves the score by a few tenths of a kcal/mol against a root-mean-square version difference of 2.09 and a zinc effect of 1.28, so sampling noise explains neither, though the entinostat outlier shows it is not negligible for every ligand. Reproducibility is worse in the smaller cavities (mean absolute difference 0.41 kcal/mol, maximum 2.7, over all fifty cavity–ligand pairs; Supplementary Tables S15 and S31).

The engine and template changes can be bounded against that noise. Because CB-Dock3 run on the zinc-free dimer returns the CB-Dock2 cavity, the same volume and the same centre to three decimal places, the two arms pose an identical search problem in an identical box and differ only in engine and template library. Over the ten ligands common to both, their Pocket 1 scores agree to a mean absolute difference of 0.20 kcal/mol, with a mean signed difference of 0.00 (SD 0.27) and a maximum of 0.5 (Supplementary Table S36); that is the size of the reproduction arm at 0.21 and smaller than the potassium control at 0.37, so the engine refinement and the expanded template library produce no effect on these scores separable from run-to-run variation. Equivalently, the version difference and the zinc-deletion effect agree ligand by ligand to a mean absolute residual of 0.33 kcal/mol. Two conditions attach. The bound covers the structure-based route only, since the expanded library does change the template-based arm (Supplementary Table S18). And because the developers attribute the rise in CASF-2016 success rate from 56.8% to 67.4% primarily to the refined engine and the updated template library[26], the present bound is specific to scores at one site of one metalloenzyme and is not a statement about pose accuracy across that benchmark.

Dependence of the ranking on ligand size

The CB-Dock2 ranking was substantially determined by molecular size, a reported bias of the Vina scoring function traced to the correlation between molecular weight and binding free energy in its training data[32]. Because the compounds span 22 to 49 heavy atoms, size and metal coordination were confounded there. The metal-aware run permits partial separation: across the fourteen non-peptide ligands the dependence weakens from $r = -0.718$ ($p = 0.004$) to $r = -0.602$ ($p = 0.023$), and within the KB series from −0.763 to −0.678. Including the p53-derived peptide reverses the direction outright, from $r = -0.387$ ($p = 0.154$) to $r = -0.693$ ($p = 0.004$), because the peptide is at once the largest ligand and the one whose score improves most; any claim that the protocol reduces size bias is therefore conditional on excluding it (Supplementary Fig. S4). Under ligand efficiency the most efficient ligands are entinostat (0.454), NCC-149 (0.417) and PCI-34051 (0.409), followed by KB-282 (0.397). Normalising by heavy-atom count therefore lifts the two small nanomolar inhibitors above the larger ligands, but it also places entinostat first, and entinostat is the one reference compound without

measurable activity against HDAC8; the metric redistributes the size dependence rather than recovering the potency order (Supplementary Table S19).

Binding poses within the series

Within the series the pose returned under CB-Dock3 for the highest-scoring compound differs materially from that obtained previously, and the difference concerns the metal. KB-282 places the nitrogen of its 4-cyanophenyl nitrile 2.03 Å from the Zn^{2+} of chain B, with the nitrile carbon at 2.33 Å; the C≡N bond length of 1.16 Å in the returned coordinates confirms the assignment. Those three distances place the metal at 89.9° to the C≡N axis measured at the nitrogen, while the nitrile itself is linear at 178.7°, so the metal lies off the axis of an undistorted C≡N group and the contact is side-on rather than the end-on arrangement in which nitriles coordinate zinc. The corresponding CB-Dock2 pose placed the closest ligand atom at 4.78 Å with no metal contact, and made 62 contacts over eight residues against 86 over fifteen. Deleting the zinc from the same receptor under CB-Dock3 reverses the change: the pose moves 6.16 Å over the 33 heavy atoms, the nitrile nitrogen by 12.05 Å, and the closest ligand atom retreats to 4.79 Å from the vacated site, within 0.01 Å of the CB-Dock2 value, while the score changes by 1.0 kcal/mol and the chain A contacts fall from four residues to one (Supplementary Tables S15 and S30). The pose as well as the score tracks the presence of the metal. KB-280 and KB-269 approach to 3.12 Å and KB-270 to 3.73 Å, while KB-268, KB-271, KB-281 and KB-272 remain beyond 5 Å; KB-282 is the only member to reach coordination distance. Residue-level contacts are compared in Fig. 5, which sets the crystallographic pose of V5X (Fig. 5A) against its template-based (Fig. 5B) and structure-based (Fig. 5C) poses and against KB-282 (Fig. 5D).

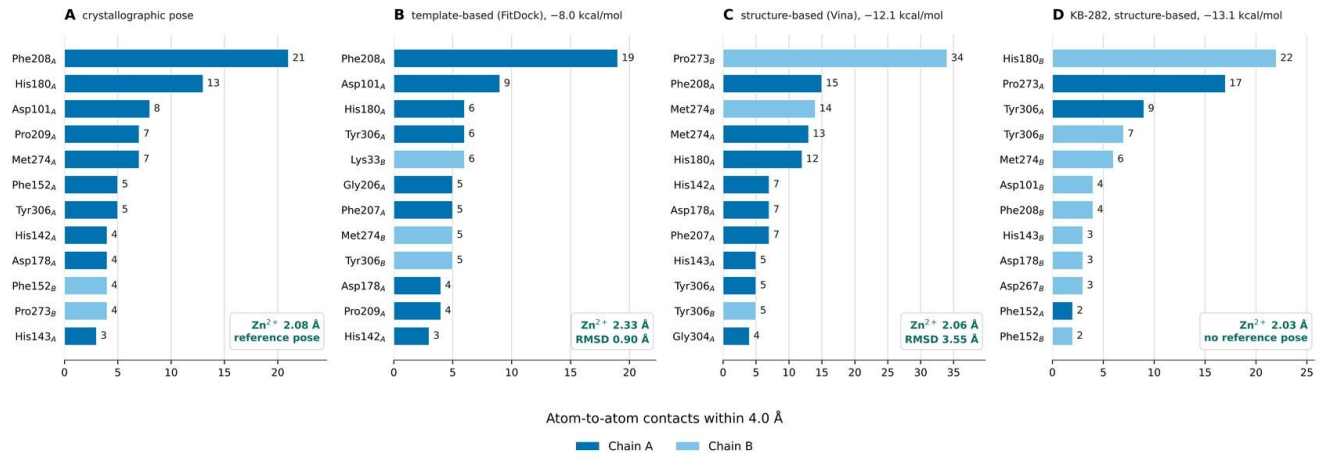

Fig. 5 Residue-level contacts in the HDAC8 catalytic cavity (PDB 2V5X). Bars give the number of ligand-to-protein atom-to-atom contacts within 4.0 Å per residue, coloured by chain (dark, chain A; light, chain B), with the chain also given as a subscript on each residue label; the twelve most contacted residues are shown in each panel. Panels: (A) the crystallographic pose of the co-crystallized 2V5X ligand, taken as the reference and computed directly from the deposited coordinates; (B) the same ligand recovered by template-based (FitDock) docking under CB-Dock3 at −8.0 kcal/mol, 0.90 Å from the crystallographic pose; (C) the same ligand recovered by structure-based (AutoDock Vina) blind docking under CB-Dock3 at −12.1 kcal/mol, the more favourable score of the two; (D) KB-282 under structure-based docking at −13.1 kcal/mol, for which no reference pose exists. The minimum ligand-to-zinc distance is given in each panel; the template-based output carries no metal records, so for panel B

that distance is measured against the deposited zinc position, the two files sharing a coordinate frame. Panels B and C are the same ligand under the two docking modes.

The contacted residues comprise the canonical HDAC8 active-site environment, including all three members of the zinc-coordinating triad, and the predicted profile is dominated by π-interactions with Phe152, Phe208 and Tyr306 together with electrostatic contacts to Asp101[33]. Fifty-seven of the 86 contacts fall in the incomplete chain B and 29 in chain A, so the pose sits at the boundary between the two catalytic sites (Supplementary Tables S1 and S23).

The significance must be stated carefully. Nitriles are recognised zinc-binding groups, but they coordinate zinc end-on, with the metal close to the C≡N axis, and the returned pose does not adopt that geometry. That is what a search carrying no coordination-geometry term would be expected to return, since AutoDock Vina represents the metal only through a generic short-range polar term and contains no coordination-geometry term[34,35], so the score contains no contribution specific to metal binding. The pose places the nitrile within coordination distance of the metal without reproducing a coordination geometry; it does not demonstrate that coordination occurs or contributes to affinity, which would require a method modelling the metal centre explicitly and ultimately an enzymatic assay. It is reported because it identifies a testable hypothesis — whether this nitrile engages the catalytic zinc at all — rather than a geometry to be reproduced.

Structure–score relationships within the series derive from nine compounds and support no conclusion here. Differences among the five substitution classes are not significant under either protocol (Kruskal–Wallis $H = 3.81$, $df = 4$, $p = 0.43$ on CB-Dock3 scores), three classes comprising a single compound. A contrast between compounds carrying an electron-withdrawing group on an N-aryl or N-benzyl substituent and those without is consistent between protocols (Welch $t = -3.39$, $p = 0.016$) but confounded with the presence of the aryl ring itself, three of the four comparison compounds carrying no aryl substituent. The multiple regression reported earlier was unstable under leave-one-out analysis and is not carried forward (Supplementary Tables S12 and S24).

The catalytic site is reached only through a crystal-packing contact

Both published runs used the two-chain file a user obtains by downloading 2V5X, but the deposited A–B pair is crystallographic packing: the entry declares two separate monomeric assemblies, and solution measurements are compatible with a monomer[3]. On either monomer alone, under every metal composition tested and for all ligands, no detected cavity contains the catalytic zinc within its search box. The closest cavity centre lies 15.82 Å from the zinc on the chain A monomer and 16.64 Å on chain B, and the largest cavity of each lies more than 21 Å from the metal, so monomer scores describe a different, non-catalytic site. The cavities of the dimer and of the chain A monomer are drawn in Fig. 4A and Fig. 4B.

The catalytic site is thus detectable by CurPocket only in the two-chain file, where the interface merges the two catalytic clefts into one cavity of roughly 1900 Å³ spanning both. The protocol reaches the biologically interesting site through a contact that the structure's own annotation and the solution data describe as absent. This is not peculiar to CB-Dock, whose developers state that geometric cavity detection is hypothesis-generating and that some surface depressions are packing artefacts[26]; the

present system supplies a concrete instance in which that warning is load-bearing and invisible to a user who inspects only the scores.

A consequence follows for every pose reported here. Because one cavity encloses the catalytic sites of both chains, 21.2 Å apart, which of them a pose occupies is a property of the individual run, and it does not reproduce: four of the nine comparable Pocket 1 poses changed chain between two submissions differing only in date and random seed, while their scores moved by 0.3 to 0.6 kcal/mol. For PCI-34051, three submissions of identical input placed the ligand at the chain A zinc, then at the chain B zinc in a near mirror image at comparable score, then at chain A again. The monomer a ligand occupies is therefore not reproducible within one version, one receptor and one cavity, although the score attached to it is stable to a few tenths of a kcal/mol. Across the CB-Dock3 set the poses divide nine to six between the chains, which are not equivalent (Supplementary Tables S1 and S31).

What the reference set does and does not establish

Neither protocol reproduces the reported potency order of the reference compounds. Under CB-Dock3, PCI-34051 (IC_{50} 10 nM, more than 200-fold selective for HDAC8[14]) scores −9.0 kcal/mol and NCC-149 (IC_{50} 70 nM[15]) −9.6, while entinostat, which shows no measurable activity against HDAC8 (IC_{50} > 100 μM[36]), scores −12.7 and the p53-derived substrate peptide −13.7. A ranking that places a substrate and an inactive compound above two nanomolar inhibitors is not a potency ranking. The metal-aware protocol corrects the inversion partially and in the predicted direction, the chelating references rising, but not completely, so the failure is not attributable to metal handling alone. The scores and the published experimental data are set side by side in Table 2.

Table 2 Binding scores under both protocols alongside published experimental data for the reference compounds

Compound	CB-Dock2	CB-Dock3	Reported IC_{50}	Selectivity	Zinc chelation	Zn^{2+} distance in CB-Dock3 pose (Å)
KB-282	−12.3	−13.1	Not determined	Not determined	No formal ZBG	2.03 (nitrile N)
KB-270	−10.0	−12.9	Not determined	Not determined	No formal ZBG	3.73
KB-281	−11.0	−9.2	Not determined	Not determined	No formal ZBG	6.71
PCI-34051	−8.8	−9.0	10 nM	>200-fold for HDAC8	Yes	2.18
NCC-149	−8.9	−9.6	70 nM	Reported HDAC8-selective	Yes	2.36
Entinostat	−8.4	−12.7	>100,000 nM (HDAC8)	Not selective for HDAC8	Yes (benzamide)	1.90
V5X	−10.8	−12.1	Not reported	Not reported	Yes	2.06
p53-derived	−8.4	−13.7	Substrate, not inhibitor	Not applicable	Substrate carbonyl	2.00

Docking scores and experimental IC_{50} values are not on a comparable scale and are tabulated together for reference only. ZBG, zinc-binding group. “Not determined”: no enzymatic inhibition or isoform-selectivity data are available for the KB compounds, and no prediction of selectivity is implied. A fold-selectivity value for NCC-149 traceable to the primary source could not be established and is therefore not quoted. No HDAC8 IC_{50} or selectivity value for the co-crystallized

ligand V5X (PDB chemical component V5X) could be located in the primary literature. Entinostat is selective against, not for, HDAC8[36].

No template was retrieved for the peptide under either version, so its score comes entirely from the structure-based route — the route shown above to misplace the one ligand whose true pose is known — and its ascent from joint last to first should be read with that in mind. No potency claim relative to the established inhibitors can therefore be made for any member of the series, and none is made; selectivity was not assessed, docking having been performed against HDAC8 alone[37]. What the reference set establishes is the calibration of the exercise: a method that scores an inactive benzamide 3.7 kcal/mol better than a 10 nM inhibitor of the target is not a basis on which chemotypes may be compared.

Cytotoxicity and predicted ADMET profile

The cytotoxicity data are new and are the only experimental measurements on any member of this series against mammalian cells. Three compounds were available in isolated form — KB-268, KB-270 and KB-292 — chosen not computationally but because they were already synthesised[22], spanning the unsubstituted spiro parent, one substituted derivative and the scaffold without the spiro centre. KB-270 was the only compound for which IC_{50} values could be estimated, at approximately 35–50 μM against J774.2, Jurkat and MCF-7 and above 50 μM against HCT-116, which is weak cytotoxicity; KB-268 and KB-292 were inactive at every concentration tested, the range extending only to 50 μM. Binding scores and the cytotoxicity data are tabulated together in Table 3.

Table 3 Binding scores and MTT cytotoxicity data for the three compounds evaluated *in vitro*

Compound	CB-Dock2 (kcal/mol)	CB-Dock3 (kcal/mol)	J774.2 (μM)	MCF-7 (μM)	HCT-116 (μM)	Jurkat (μM)	Cytotoxicity at concentrations tested
KB-270	−10.0	−12.9	≈35	≈50	>50	≈43	Weak (IC_{50} ≈ 35–50 μM)
KB-268	−8.8	−8.4	>50	>50	>50	>50	Inactive (IC_{50} > 50 μM)
KB-292	−8.7	−8.0	>50	>50	>50	>50	Inactive (IC_{50} > 50 μM)

Scores are given under both server versions. IC_{50} values are approximate estimates derived from three concentrations in duplicate (Methods).

The pattern is consistent with whatever activity the series possesses residing in the substituent rather than in the heterocyclic core, since the unsubstituted spiro parent and the scaffold lacking the spiro centre were both inactive; with one active compound among three, however, the comparison indicates this rather than establishing it. The pattern across three of four lines is more consistent with a general than a selective mechanism, though four lines cannot establish that. These data cannot test either ranking: the compounds were not chosen computationally, those ranked highest under either protocol were not among them, and general cytotoxicity is not a measure of HDAC8 engagement.

The ADMET predictions are also new; twenty-four endpoints were computed for all fifteen ligands (Supplementary Tables S5–S11) in order to identify liabilities that would make synthesis inadvisable before it is undertaken. Predicted Caco-2 permeability is higher across the series than for V5X, the co-crystallized hydroxamate (0.067), the KB values ranging from 0.239 for KB-271 to 0.930 for KB-292, with 0.619 for KB-282 and 0.350 for KB-270, at roughly two-thirds the molecular weight.

The advantage does not extend to the smaller hydroxamates, since KB-271 (0.239), KB-269 (0.344) and KB-270 (0.350) fall below NCC-149 (0.453); it therefore holds against the two largest reference ligands only and may reflect the size dependence already seen in the docking scores. Against this, the predicted drug-induced liver injury probability for KB-282 (0.737) and KB-270 (0.734) exceeds that of the hydroxamate (0.572), and predicted reproductive toxicity is markedly higher for KB-282 (0.962) than for any reference compound, the largest adverse difference in the panel; KB-270 is more moderate, so the liability is not uniform. All values are model outputs for compounds whose target activity has not been demonstrated; the hepatic and reproductive signals are the liabilities to characterise before the series is advanced, and the differences between candidate leads are smaller than the uncertainty of the models (Supplementary Fig. S2; Supplementary Tables S26–S29).

Limitations

Three qualifications constrain the methodological comparison. Each ligand was submitted once to each version in the original design, so the differences cannot be decomposed into version and run-to-run components from that design alone; Vina is a stochastic search, and the two runs returning a positive score on first submission show the variation to be occasionally gross. The receptor-variant study supplies the missing bound, and against 0.37 kcal/mol the root-mean-square version difference of 2.09 is large but not overwhelmingly so; the case for a systematic version effect rests additionally on the changes tracking the metal protocol, on their correspondence with qualitative pose changes, and on the metal term reproducing them when varied directly.

Second, the two versions differ in more than metal handling, but the contribution of the engine refinement and the expanded template library can be bounded, because CB-Dock3 run on the zinc-free dimer returns the CB-Dock2 cavity and therefore poses an identical search problem: over the ten ligands common to both arms the Pocket 1 scores agree to a mean absolute difference of 0.20 kcal/mol, within run-to-run variation. That bound applies to the structure-based scores on which the ranking rests and not to the template-based arm, which the expanded library does change. An earlier form of this analysis stated that isolating the metal term was impossible through the interface, and a later one that the remainder could not be bounded; both statements were wrong.

Third, the receptor file is imperfect. Chain B lacks residues 87 to 93, the chains are therefore not equivalent, and the largest cavity spans both catalytic sites; part of the score change may accompany relocation between the two sites rather than closer engagement of one, and several of the highest-scoring poses, including that of KB-282, occupy the incomplete chain. The two zinc positions coincide to 0.11 Å, but the cavities around them do not, since one chain is incomplete and it is the cavity that is scored; scores obtained at the chain B site are therefore not interchangeable with those obtained at chain A, and the comparison is not one within a single site; repeating the calculation on one complete chain would remove both ambiguities and also remove the catalytic site from the detected cavity set. Concerning the series the constraints are more severe: no HDAC8 assay was performed and cytotoxicity cannot substitute for one; six of the ten KB structures, including the highest-ranked, have not been synthesised; the receptor was rigid, so the mobile L1 loop is unrepresented[38]; and the nitrile–zinc contact requires confirmation by a method that models coordination.

Four further constraints bear on generality. The receptor-variant study was performed on one crystal structure of one target, so what generalises is the dependence of the result on the composition of the receptor file rather than the magnitude of 1.28 kcal/mol. The four compounds carrying a stereogenic spiro centre were drawn and submitted in the (S) configuration alone, so the scores reported for them refer to one enantiomer and no statement about stereoselectivity is possible. The ADMET endpoints come from a single prediction suite with no external validation for this chemotype, and several predicted plasma protein binding values exceed 100%, which is model extrapolation rather than a physical quantity, so the panel serves to flag liabilities and not to rank candidates. Finally, repeated submission of identical input changes which monomer a pose occupies while leaving the score stable to a few tenths of a kcal/mol, so every pose-derived statement here, the KB-282 nitrile contact included, is conditional on a single run.

Implications for practice

Five recommendations follow, addressed to users of blind-docking servers generally. The server version and date of calculation should be reported as routinely as any other software version, since a CB-Dock2 and a CB-Dock3 ranking on identical input are different results, and such reporting is at present uncommon[28]. Where the target is a metalloenzyme, the treatment of the metal should be stated explicitly, since this single choice moves scores by up to 5.3 kcal/mol here. Where the receptor carries a co-crystallised ligand, that ligand should be redocked and the deviation reported: it costs one calculation and is the only internal check on pose accuracy available without external data, and the pocket RMSD of the template panel does not serve the purpose. Where the search is blind, the distance from the returned pose to the intended site should be reported alongside the score, since ranking and placement carry no information about each other[27]. Finally, where the file contains more than one chain, the biological assembly should be checked before submission.

Conclusions

Fifteen ligands docked into the same HDAC8 crystal structure under two successive versions of the same blind-docking server, with ligand input, cavity definitions and search boxes held constant, give rankings that do not agree (Spearman $\rho = +0.23$, $p = 0.42$; root-mean-square difference 2.09 kcal/mol), individual compounds moving by up to thirteen rank positions. The disagreement is systematic: under the metal-aware protocol the zinc-binding groups of the reference inhibitors engage the catalytic ion at coordination distance, contacts absent from the CB-Dock2 structure-based poses, and the ligands establishing them are those whose scores improve most. Against the co-crystallised ligand as an internal reference, the template-based pose is recovered within 0.90 Å while the structure-based pose carrying the more favourable score falls outside the 2.0 Å criterion, so score and pose accuracy are inversely related in the one case where the answer is known.

The attribution to metal handling can be replaced by a measurement. Holding version, engine and template library fixed and varying only the receptor file, deletion of the catalytic zinc restores the CB-Dock2 cavity to three decimal places in all three coordinates and makes the score less favourable for eight of ten ligands by a mean of 1.28 kcal/mol, against a run-to-run bound of 0.37. Run on that zinc-free receptor, CB-Dock3 reproduces the CB-Dock2 scores to a mean absolute difference of 0.20 kcal/mol, so the engine refinement and the expanded template library contribute nothing to these scores

that is separable from run-to-run variation. The cavity difference between the two published runs is therefore not an algorithmic revision but a consequence of which atoms the receptor was taken to contain. The same experiment shows the catalytic site to be detectable only in the two-chain file, although the entry declares the biological unit a monomer: the site the calculation finds is reached through a crystal-packing contact, and which of the two sites a pose occupies is not reproducible between submissions of identical input.

Neither protocol reproduces the reported potency order of the reference compounds, and no claim of HDAC8 inhibition, potency or isoform selectivity is made for the spiro series. Two results stand independently of the docking: KB-270 was weakly cytotoxic across three of four cell lines ($IC_{50} \approx 35$–50 μM) while the unsubstituted parent and the non-spiro scaffold were inactive, placing activity in the substituent rather than the core; and the predicted ADMET profile shows a consistent absorption advantage over the hydroxamate chemotype set against elevated hepatic and reproductive-toxicity probabilities. Under the metal-aware protocol KB-282 places its 4-cyanophenyl nitrile 2.03 Å from the catalytic zinc, a testable structural hypothesis rather than a demonstrated interaction; synthesis followed by enzymatic assay, with counter-screening against a further class I isoform, would settle it.

A blind-docking score ranking reported without the server version, without an explicit statement of metal handling, without an internal pose control and without the distance from the returned pose to the intended site is not a reproducible result, and should not be treated as one.

Methods

Receptor preparation and its limitations

The crystal structure of human HDAC8 in complex with a hydroxamate inhibitor (PDB 2V5X, resolution 2.25 Å) served as the receptor for all docking calculations[29]. The entry carries a single engineered substitution, Ser39Asp in both chains, recorded in the _struct_ref_seq_dif category of the deposited mmCIF file; it lies outside the catalytic pocket, and the sequence is otherwise that of the wild-type enzyme. The provenance is stated explicitly because several secondary summaries of this entry instead report the Tyr306Phe substitution of the companion substrate complex 2V5W; Tyr306 is present in 2V5X and is contacted by most of the poses reported in the Results. The deposited coordinates contain two protein chains, but the biological unit is a monomer: the entry declares two separate monomeric assemblies, one per chain, under both author and software assignment, and size-exclusion chromatography together with light scattering showed HDAC8 to be monomeric in solution as both apoprotein and inhibitor complex[3]. The A–B pair in the asymmetric unit is therefore crystallographic packing rather than a physiological dimer, a point of consequence for the analysis of site selection reported above. Each chain carries one catalytic Zn^{2+} ion, coordinated by Asp178, His180 and Asp267 and by the two hydroxamate oxygens of the co-crystallized ligand at 1.95–2.25 Å, together with two K^+ ions. The catalytic Zn^{2+} was retained in the receptor file for both docking runs. Crystallographic water molecules were removed and the co-crystallized inhibitor was deleted prior to docking; hydrogen addition and side-chain treatment were performed by the server-side preparation routine of the respective CB-Dock version.

Two properties of this receptor file constrain the interpretation of the results reported above, and both are set out here rather than in the limitations alone because they affect the primary data.

First, the deposited chain B is incomplete. Residues 87 to 93 are absent, so the chain runs from Gly86 directly to Ile94; chain A is complete over the modelled range 14–376. The CB-Dock3 submission interface returns an explicit chain-break warning for this file. The missing segment contains Asp87, Asp88 and Asp89, which are consequently present in chain A and absent from chain B. The two chains are therefore not equivalent receptors, and any residue-contact statistic pooled across them mixes a complete binding-site environment with an incomplete one. Contact counts are reported pooled, as returned by the server, but a per-chain breakdown for every ligand is provided in Supplementary Table S1 so that the composition of each pose can be inspected directly.

Second, because the arrangement is a crystallographic dimer, the largest detected cavity spans the interface region and contains the catalytic site of both chains. Which of the two sites a given pose occupies is therefore a property of the individual result rather than of the cavity definition, and is tabulated per ligand in Supplementary Table S1.

Ligand preparation

Molecular structures were drawn in ChemDraw Professional 15.0 (PerkinElmer Informatics, Waltham, MA, USA) and exported in MOL format. The MOL files were uploaded to the docking server together with the receptor file, and all subsequent three-dimensional structure generation, protonation, charge assignment and geometry optimization were performed by the internal preparation routines of the CB-Dock server. The identical procedure was used for both the CB-Dock2 and the CB-Dock3 calculations, so that ligand input geometry is held constant between the two runs and does not contribute to the differences reported above. No local conformer generation, energy minimization or charge assignment was performed by the authors. The spiro center of KB-280, KB-281, KB-282 and KB-283 is stereogenic; the (*S*)-configured enantiomer was drawn and submitted in each case, and the scores reported for these compounds accordingly refer to that enantiomer alone.

Fifteen ligands were prepared in this way: nine spiro[indoline-3,5′-tetrazolo[1,5-*c*]quinazoline] derivatives (KB-268–KB-272, KB-280–KB-283); KB-292, the parent tetrazolo[1,5-*c*]quinazoline scaffold lacking the spiro-indoline moiety, included as a structural control; and five reference compounds, namely the HDAC8-selective inhibitors PCI-34051 and NCC-149, the class I inhibitor entinostat, the hydroxamate V5X co-crystallized with 2V5X, and the p53-derived substrate peptide taken from the companion complex 2V5W.

Docking calculations under two server versions

Blind docking was performed twice on the same receptor and the same fifteen ligand files, first with CB-Dock2[24,25] and subsequently with CB-Dock3[26].

Both versions identify binding cavities with CurPocket, a protein-surface-curvature algorithm, dock into the five highest-ranked cavities, merge pockets sharing more than 50% of binding residues, and score structure-based poses with AutoDock Vina[34]. Both additionally retrieve homologous templates from a structure database by FP2 fingerprint similarity (threshold 0.4) and, where templates are found, perform template-based docking with FitDock[39] after superposition of the query protein onto the template.

The versions differ in four respects relevant here. CB-Dock3 employs a refined docking engine and an expanded template library, drawn from BioLiP2 in its 2025.04.23 release[26]. It adds metal-aware scoring, through the native metal parameterisation of AutoDock Vina 1.2.0 and through metal–atom

interaction terms added to FitDock and fitted against metal-containing complexes of PDBbind 2020[26]. Most consequentially for a metalloenzyme, it changes what reaches the calculation at all: earlier versions of the pipeline routinely excluded non-amino-acid HETATM records, which the developers identify as a source of error for metalloenzymes, whereas CB-Dock3 flags such records in the preparation interface and allows them to be retained[26]. The reported CASF-2016 success rate rises accordingly, from 56.8% under CB-Dock2 to 67.4% at RMSD ≤ 2.0 Å[26]. The protocol change is visible in the returned files themselves: the CB-Dock3 structure-based complexes carry the two zinc ions as explicit coordinate records, whereas the CB-Dock2 complexes and the template-based outputs of both versions carry none. The cavity-detection algorithm itself, CurPocket, is not reported as having changed between the two versions[26], which matters for the interpretation of the cavity difference described next.

Cavity definitions were closely comparable between the two runs but not identical. The catalytic cavity, designated Pocket 1 throughout, was returned with a volume of 2020 Å^3 and centre (16.465, −0.699, 43.607) by CB-Dock2 and 1887 Å^3 and centre (16.334, −0.699, 43.607) by CB-Dock3; the remaining four cavities were returned with identical volumes and centres by both versions. Since CurPocket is unchanged between the versions, the difference must originate in the receptor presented to it rather than in the detection algorithm, and the receptor-variant experiment shows directly that it does. Search-box dimensions are assigned by the server as a function of ligand size and were identical between runs for every ligand except V5X, for which CB-Dock2 used 27 × 27 × 27 Å and CB-Dock3 26 × 26 × 26 Å. Box dimensions for all ligands under both versions are listed in Supplementary Tables S2 and S3. Sampling geometry can therefore be excluded as a material source of the differences reported in the Results, with the single noted exception.

Each ligand was submitted once to each server version. Replicate submissions were not performed as part of the original design. Consequently the version-to-version differences reported above cannot be decomposed into a version component and a run-to-run stochastic component from that design alone, a limitation set out above; the receptor-variant study was undertaken in part to supply the missing bound, since it contains pairs of runs that differ only outside the scored region and therefore behave as replicates.

Docking validation against the internal reference pose

Protocol validation was performed against the co-crystallized ligand of the receptor structure itself, which provides an internal reference pose of known geometry. The deposited coordinates of the 2V5X ligand (PDB chemical component V5X) were extracted from the crystal structure, and the root-mean-square deviation between those coordinates and each docked pose of V5X was computed over the 46 heavy atoms.

Four pose sets were assessed: the structure-based (AutoDock Vina) Pocket 1 pose and the template-based (FitDock) pose returned by each of the two server versions for the same ligand input. Both pose files and the deposited structure share a common coordinate frame, so RMSD was computed in place, over the 46 heavy atoms and without superposition. The atom correspondence was taken as the mapping of the ligand onto itself that preserves the molecular graph and minimizes the deviation; the molecule admits two such mappings, which give 0.90 and 1.12 Å for the template-based pose and 3.55 and 3.62 Å for the structure-based pose, so neither choice alters the comparison against the 2.0 Å threshold. Deviations after optimal superposition are reported alongside the in-place values in

Supplementary Table S4, since the difference between the two separates conformational error from rigid-body displacement. Results are reported above and in Supplementary Table S4. The deposited structure contains two independent copies of the ligand, one in each monomer. The copy adjacent to the docked poses (author chain A, label_asym_id F) was used as the reference throughout. Superposing the monomers on 356 Cα atoms (0.48 Å) places the two zinc ions 0.11 Å apart but leaves the two ligand copies 4.79 Å apart, a difference confined to one distal substituted ring; the nine heavy atoms within 8.5 Å of the metal agree between copies to between 0.14 and 0.57 Å. Deviations are therefore reported both for the whole molecule and, separately, for that zinc-binding subset, which the crystal structure defines reproducibly.

It should be stated explicitly that the RMSD values reported in this section are ligand-pose deviations. The "pocket RMSD" quantity displayed in the template information panel of the CB-Dock interface is a different measure: it expresses the structural similarity between the binding pocket of the query protein and that of the retrieved template protein, and serves as a template-selection filter, templates being retained when sequence identity exceeds 40% and pocket RMSD is at most 4 Å[24]. It carries no information about the accuracy of the docked ligand pose and is not used as validation here.

Pose analysis

Docked complexes were analyzed directly from the returned coordinate files. Residues were counted as contacting when any heavy atom lay within 4.0 Å of a ligand heavy atom, and contacts were resolved by chain. Distances from each ligand to both catalytic Zn^{2+} ions were computed as the minimum ligand-heavy-atom-to-metal distance. Protein–ligand interaction types were assigned with Discovery Studio Visualizer 2017 R2 (Dassault Systèmes BIOVIA) using CASP contact criteria (sum of van der Waals radii plus a tolerance of 0.5 Å); recorded interaction types comprised conventional and carbon–hydrogen bonds, parallel and T-shaped π–π stacking, and π–σ, π–alkyl, π–anion, halogen and π–sulfur contacts.

A caution applies to the interpretation of short ligand–metal distances. The AutoDock Vina scoring function comprises a van der Waals-like potential, a non-directional hydrogen-bond term, a hydrophobic term and a conformational entropy penalty; it contains no explicit electrostatic term and no dedicated metal-coordination term, and metal ions are represented through the generic short-range polar term[34]. Force fields incorporating explicit zinc-coordination geometry have been developed for this reason[35]. A ligand atom placed at coordination distance from the zinc in a Vina-scored pose therefore indicates that the pose is geometrically compatible with coordination; it does not demonstrate coordination, and no energetic contribution specific to metal binding is represented in the reported score.

Provenance of the compounds and characterization data

The ten KB compounds fall into two groups. Four were synthesized and characterized previously and reported as b1, b2, b3 and c8[22]: KB-268 (b1), KB-270 (b2), KB-271 (b3) and KB-292 (c8). The remaining six — KB-269, KB-272, KB-280, KB-281, KB-282 and KB-283 — are designed structures examined by computational methods only, for which no experimental data are reported.

Synthetic procedures, yields, melting points, elemental analyses, IR, ^{1}H NMR and LC-MS data for the four previously reported compounds appear in full in the original publication[22] and are not reproduced here. The analytical conditions under which that characterization was performed are

summarized for completeness: melting points in open capillary tubes on a Mettler Toledo MP 50 apparatus, uncorrected; elemental analysis (C, H, N) on a vario EL Cube analyzer, within ±0.3% of theoretical values; ^{1}H NMR at 400 MHz on a Varian-Mercury 400 spectrometer in DMSO-d_6 with TMS as internal standard; LC-MS on an Agilent 1100 Series system with diode-array and mass-selective detection under atmospheric pressure chemical ionization; electron impact mass spectra on a Varian 1200 L instrument at 70 eV. Purity exceeding 95% was confirmed by ^{1}H NMR and LC-MS. No new synthetic or analytical work was undertaken for the present study.

Cytotoxicity assessment

Cytotoxicity data for KB-268, KB-270 and KB-292 were obtained in a parallel investigation of the anticancer potential of tetrazolo[1,5-*c*]quinazoline derivatives and have not been published elsewhere. They are presented as independent experimental context and not as validation of any docking prediction: the three compounds were not selected computationally, and general cytotoxicity in tumour cell lines is not a measure of HDAC8 engagement.

Testing was performed at the Molecular Design Center, Danylo Halytsky Lviv National Medical University, using the MTT assay (Sigma-Aldrich). Cell lines were MCF-7 (RRID:CVCL_0031), HCT-116 (RRID:CVCL_0291), Jurkat (RRID:CVCL_0065) and J774.2 (RRID:CVCL_0357), all verified mycoplasma-free. Cells (6,000 per well) were seeded in 96-well plates in 100 µl DMEM or RPMI-1640 per ATCC recommendations; adherent lines were allowed to attach overnight and the suspension line Jurkat was treated directly after seeding. Compounds were added at 1, 10 and 50 µM in duplicate and incubated for 72 h at 37 °C in 5% CO_2. Medium was removed, MTT was added from a 5 mg/ml stock, and cells were incubated a further 4 h. Formazan crystals were dissolved in DMSO and absorbance read at 570 nm (TECAN Infinite M Plex). IC_{50} values were estimated in GraphPad Prism 8 by interpolation between the three concentrations; a four-parameter logistic model is not identifiable from three concentration levels and was not fitted. Because three concentrations were tested in duplicate, the resulting IC_{50} values are approximate estimates; values reported as >50 µM indicate that the IC_{50} was not reached within the range tested.

ADMET prediction

Physicochemical descriptors and ADMET parameters were predicted with admetSAR 3.0[40]. SlogP was calculated by the RDKit algorithm; logP was obtained from the CLMGraph deep-learning model. Predicted parameters are distributed across seven supplementary tables: absorption and distribution in Supplementary Table S5; transporter interactions in Supplementary Table S6; metabolism and excretion in Supplementary Table S7; toxicity in Supplementary Table S8; endocrine disruption across the Tox21 panel in Supplementary Table S9; ecotoxicity in Supplementary Table S10; and full physicochemical descriptors in Supplementary Table S11. Predicted plasma protein binding values are unconstrained regression outputs and exceed 100% for several compounds, which reflects model extrapolation rather than a physical quantity. admetSAR 3.0 returns no hepatotoxicity model; the human liver microsomal (HLM) stability output is a metabolic stability endpoint and is reported as such, with drug-induced liver injury (DILI) the only hepatic-injury signal available.

Statistical analysis

All statistical computations were performed with SciPy under Python 3.12. Agreement between the two docking runs was assessed by Pearson correlation, Spearman rank correlation and Kendall's τ over all fifteen ligands, together with the mean signed difference, its standard deviation and the root-mean-square difference. Rank positions are assigned as tied minimum ranks, four pairs of ligands sharing a CB-Dock2 score and two a CB-Dock3 score.

The dependence of binding score on ligand size was evaluated by Pearson correlation of Pocket 1 score against heavy-atom count and against molecular weight, computed separately for each run. Heavy-atom counts were determined directly from the returned ligand coordinates rather than taken from descriptor software, and are listed in Supplementary Table S11. Ligand efficiency was calculated as the negative binding score divided by heavy-atom count[41], and is used solely as a diagnostic of size dependence; the metric retains a residual dependence on molecular size and is not a physical quantity, so its application requires care[42].

Structure–score relationships within the spiro series were examined descriptively. Pearson correlations between Pocket 1 score and individual predicted descriptors (SlogP, logP, TPSA, MW, HBA, HBD, BBB permeability, HIA) are reported both unadjusted and with Bonferroni adjustment for the eight descriptors tested (Supplementary Table S12). Predicted BBB permeability and HIA are themselves admetSAR outputs derived largely from TPSA, logP and molecular weight, so their association with docking score reflects descriptor collinearity rather than an independent structure–activity relationship. Compounds were assigned to substitution classes strictly according to the structures of Supplementary Table S34: unsubstituted at position 1 (KB-268); acetamide linker at position 1 (KB-269, KB-270, KB-271); carbothioamide at position 1 (KB-280); carboxamide at position 1 (KB-281, KB-282, KB-283); and acylation at position 6′ rather than position 1 (KB-272). Group differences were evaluated by Kruskal–Wallis test, which is appropriate given that two classes comprise single compounds, and are interpreted descriptively rather than inferentially for the same reason. Since the series contains no electron-donating aryl substituents, the electronic comparison was defined as compounds carrying an electron-withdrawing group on an *N*-aryl or *N*-benzyl substituent (KB-269, KB-270, KB-271, KB-281, KB-282) against those without (KB-268, KB-272, KB-280, KB-283), and was tested by Welch's *t*-test with Satterthwaite degrees of freedom and by the Mann–Whitney *U* test. That comparison is confounded, since three of the four compounds in the second group carry no aryl substituent at all, so the electronic contribution of the substituent cannot be separated from the steric and π-contact contribution of the ring itself.

Two Pocket-level runs returned positive scores on first submission and were resubmitted; the converged replacements are used throughout and the superseded values are retained, with their provenance, in Supplementary Tables S3 and S13. Where any statistic would be affected by that substitution it is reported on the converged values, and the alternative is given in the corresponding supplementary table.

Controlled receptor-variant study

The version comparison above confounds the metal term with every other change between CB-Dock2 and CB-Dock3, and this is recorded above as a limitation of the original design. A supplementary study was therefore carried out under CB-Dock3 alone, holding the engine, the template library and the scoring function fixed and varying only the receptor file, so that the metal term could be varied within a single version rather than across two.

Eight receptor variants were prepared from the same deposited 2V5X coordinates by deletion of records only, without rebuilding, renumbering or re-minimisation, so that every variant shares one coordinate frame: the A+B dimer with and without the catalytic Zn^{2+}; the chain A monomer with Zn^{2+} and K^+, with Zn^{2+} removed, with K^+ removed, and with all metals removed; and the chain B monomer with and without Zn^{2+}. Variant composition and atom counts are given in Supplementary Table S14. Nine ligands spanning the range of ligand-to-metal distance observed in the main study were submitted to every variant, and the p53-derived peptide was added subsequently as a tenth (Supplementary Tables S15 and S16). The pair of chain A arms differing only in the presence of the two K^+ ions, which lie 21.0 and 32.0 Å from the Pocket 1 centre and more than 8 Å outside the search box, pose an effectively identical search problem and are used to bound run-to-run variation of the sampling.

Returned complexes were analysed from the coordinate files by the procedure described above. Cavity identity was confirmed by matching the score in each returned filename to the CurPocket table, which agreed in every one of the 41 cases checked, and cavity position was expressed as the distance from the cavity centre to the catalytic Zn^{2+} in the deposited frame, together with a test of whether that ion falls inside the axis-aligned search box. One supplied run was identified as having been performed against the wrong receptor — it returned the zinc-free cavity volume and a complex carrying no zinc record — and was resubmitted; the superseded file is reported in Supplementary Table S17 and not used.

Ethics statement and use of large language models

No primary human tissue, animal experiments or patient data were involved in this study. The computational analyses were performed on publicly available structural data (PDB 2V5X), and the cytotoxicity measurements used established, commercially available cell lines. Ethical committee approval was therefore not required.

Claude Sonnet 5 and Claude Opus 5 (Anthropic) large language models were used for English-language editing and for writing the plotting and table-formatting scripts that render the authors' primary data as figures and supplementary tables. Every numerical value in the figures and tables was checked by the authors against the primary data, and all text was reviewed and edited by the authors, who take full responsibility for the content of the manuscript. The models do not meet authorship criteria and are not listed as authors.

Supplementary Information Supplementary Information file (Supplementary Figures S1–S4 and Supplementary Note 1) and Supplementary Data 1 (Supplementary Tables S1–S36) accompany this paper. Correspondence and requests for materials should be addressed to L.A.

Acknowledgements This manuscript was prepared in Zaporizhzhia, Ukraine, with support from the Armed Forces of Ukraine and the Territorial Defence Forces, which enabled safe working conditions during the ongoing conflict. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Author contributions Conceptualization and design of the study, methodology, project administration: L.A.; software: L.A. and O.A.; investigation: L.A., O.A. and Y.K.; resources: S.K.; original draft preparation: L.A.; data curation, formal analysis, visualization, writing—review and editing: L.A., O.A., V.S., M.M., H.L. and M.A.; supervision: S.K. and M.A. All authors read and approved the final manuscript.

Data availability All data supporting the conclusions of this article are included in the article and in its Supplementary Information. Supplementary Data 1 contains Supplementary Tables S1–S36, comprising the blind docking results for all five cavities under both server versions, the pose-accuracy validation, the per-chain contact analyses, the statistical analyses, the cytotoxicity measurements, the ADMET predictions and the complete controlled receptor-variant study; the Supplementary Information file contains Supplementary Figures S1–S4 and Supplementary Note 1. The crystal structure analysed is available from the Protein Data Bank under accession 2V5X. Any further information is available from the corresponding author upon reasonable request.

Declarations

Competing interests The authors declare no competing interests.

References

1. Chakrabarti, A. et al. HDAC8: a multifaceted target for therapeutic interventions. *Trends Pharmacol. Sci.* **36**, 481–492. https://doi.org/10.1016/j.tips.2015.04.013 (2015).

2. Somoza, J. R. et al. Structural snapshots of human HDAC8 provide insights into the class I histone deacetylases. *Structure* **12**, 1325–1334. https://doi.org/10.1016/j.str.2004.04.012 (2004).

3. Vannini, A. et al. Crystal structure of a eukaryotic zinc-dependent histone deacetylase, human HDAC8, complexed with a hydroxamic acid inhibitor. *Proc. Natl Acad. Sci. USA* **101**, 15064–15069. https://doi.org/10.1073/pnas.0404603101 (2004).

4. Deardorff, M. A. et al. HDAC8 mutations in Cornelia de Lange syndrome affect the cohesin acetylation cycle. *Nature* **489**, 313–317. https://doi.org/10.1038/nature11316 (2012).

5. Wilson, B. J., Tremblay, A. M., Deblois, G., Sylvain-Drolet, G. & Giguère, V. An acetylation switch modulates the transcriptional activity of estrogen-related receptor alpha. *Mol. Endocrinol.* **24**, 1349–1358. https://doi.org/10.1210/me.2009-0441 (2010).

6. Wu, R., Wang, S., Zhou, N., Cao, Z. & Zhang, Y. A proton-shuttle reaction mechanism for histone deacetylase 8 and the catalytic role of metal ions. *J. Am. Chem. Soc.* **132**, 9471–9479. https://doi.org/10.1021/ja103932d (2010).

7. Finnin, M. S. et al. Structures of a histone deacetylase homologue bound to the TSA and SAHA inhibitors. *Nature* **401**, 188–193. https://doi.org/10.1038/43710 (1999).

8. Yoshida, M., Kijima, M., Akita, M. & Beppu, T. Potent and specific inhibition of mammalian histone deacetylase both in vivo and in vitro by trichostatin A. *J. Biol. Chem.* **265**, 17174–17179. https://doi.org/10.1016/S0021-9258(17)44885-X (1990).

9. Atadja, P. Development of the pan-DAC inhibitor panobinostat (LBH589): successes and challenges. *Cancer Lett.* **280**, 233–241. https://doi.org/10.1016/j.canlet.2009.02.019 (2009).

10. Poole, R. M. Belinostat: first global approval. *Drugs* **74**, 1543–1554. https://doi.org/10.1007/s40265-014-0275-8 (2014).

11. Fournel, M. et al. MGCD0103, a novel isotype-selective histone deacetylase inhibitor, has broad spectrum antitumor activity in vitro and in vivo. *Mol. Cancer Ther.* **7**, 759–768. https://doi.org/10.1158/1535-7163.MCT-07-2026 (2008).

12. Furumai, R. et al. FK228 (depsipeptide) as a natural prodrug that inhibits class I histone deacetylases. *Cancer Res.* **62**, 4916–4921 (2002).

13. Han, J. W. et al. Apicidin, a histone deacetylase inhibitor, inhibits proliferation of tumor cells via induction of p21WAF1/Cip1 and gelsolin. *Cancer Res.* **60**, 6068–6074 (2000).

14. Balasubramanian, S. et al. A novel histone deacetylase 8 (HDAC8)-specific inhibitor PCI-34051 induces apoptosis in T-cell lymphomas. *Leukemia* **22**, 1026–1034. https://doi.org/10.1038/leu.2008.9 (2008).

15. Suzuki, T. et al. Rapid discovery of highly potent and selective inhibitors of histone deacetylase 8 using click chemistry to generate candidate libraries. *J. Med. Chem.* **55**, 9562–9575. https://doi.org/10.1021/jm300837y (2012).

16. Bandaru, N. V. M. R. et al. Targeted HDAC8 inhibition with non-hydroxamate[1,2,4]triazolo[4,3-a]quinoline compounds. *Sci. Rep.* **16**, 7472. https://doi.org/10.1038/s41598-026-38490-y (2026).

17. Vaidya, A. S. et al. Novel histone deacetylase 8 ligands without a zinc chelating group: exploring an 'upside-down' binding pose. *Bioorg. Med. Chem. Lett.* **22**, 6621–6627. https://doi.org/10.1016/j.bmcl.2012.08.104 (2012).

18. Antypenko, L. et al. In silico identification and characterization of spiro[1,2,4]triazolo[1,5-c]quinazolines as diacylglycerol kinase α modulators. *Molecules* **30**, 2324. https://doi.org/10.3390/molecules30112324 (2025).

19. Antypenko, O. M., Kovalenko, S. I. & Karpenko, O. V. Synthesis and hydrolytic cleavage of tetrazolo[1,5-c]quinazolines. *Synth. Commun.* **46**, 551–555. https://doi.org/10.1080/00397911.2016.1156131 (2016).

20. Antypenko, L. M., Kovalenko, S. I., Los', T. S. & Rebets', O. L. Synthesis and characterization of novel N-(phenyl, benzyl, hetaryl)-2-([1,2,4]triazolo[1,5-c]quinazolin-2-ylthio)acetamides by spectral data, antimicrobial activity, molecular docking and QSAR studies. *J. Heterocycl. Chem.* **54**, 1267–1278. https://doi.org/10.1002/jhet.2702 (2016).

21. Stavytskyi, V. et al. Substituted 3-R-2,8-dioxo-7,8-dihydro-2H-pyrrolo[1,2-a][1,2,4]triazino[2,3-c]quinazoline-5a(6H)carboxylic acids and their salts – a promising class of anti-inflammatory agents. *Anti-Inflamm. Anti-Allergy Agents Med. Chem.* **20**, 75–88. https://doi.org/10.2174/1871523019666200505073232 (2021).

22. Antypenko, L. et al. 5,6-Dihydrotetrazolo[1,5-c]quinazolines: toxicity prediction, synthesis, antimicrobial activity, molecular docking, and perspectives. *Arch. Pharm. (Weinheim)* **356**, e2300029. https://doi.org/10.1002/ardp.202300029 (2023).

23. Bilai, I. M., Dariy, V. I., Khilkovets, A. V. & Bilai, A. I. Current research trends of 1,2,4-triazole derivatives biological activity (literature review). *Curr. Issues Pharm. Med. Sci. Pract.* **18**, 197–205. https://doi.org/10.14739/2409-2932.2025.2.322344 (2025).

24. Liu, Y. et al. CB-Dock2: improved protein-ligand blind docking by integrating cavity detection, docking and homologous template fitting. *Nucleic Acids Res.* **50**, W159–W164. https://doi.org/10.1093/nar/gkac394 (2022).

25. CB-Dock2: cavity detection guided blind docking. https://cadd.labshare.cn/cb-dock2/index.php (accessed 22 April 2025).

26. Liu, Y. et al. CB-Dock3: an enhanced web server for protein–ligand blind docking. *Nucleic Acids Res.* **54**, W238–W245. https://doi.org/10.1093/nar/gkag417 (2026).

27. Agboola, O. E. et al. Ranking metrics conceal pose errors in blind docking with AutoDock Vina, a cross-target benchmark. *Results Eng.* **32**, 112651. https://doi.org/10.1016/j.rineng.2026.112651 (2026).

28. Giap, V., Wang, E. & Nguyen, C. Measuring the re-executability of published molecular docking claims. Preprint at https://arxiv.org/abs/2607.12117 (2026).

29. Vannini, A. et al. Substrate binding to histone deacetylases as shown by the crystal structure of the HDAC8-substrate complex. *EMBO Rep.* **8**, 879–884. https://doi.org/10.1038/sj.embor.7401047 (2007).

30. Antypenko, L. et al. Spiro[indoline-tetrazolo[1,5-c]quinazoline] derivatives as candidate HDAC8 inhibitors: a molecular docking and ADMET analysis using the HDAC8–inhibitor complex (PDB 2V5X). Preprint at https://doi.org/10.21203/rs.3.rs-10243724/v1 (2026).

31. Bressi, J. C. et al. Exploration of the HDAC2 foot pocket: synthesis and SAR of substituted N-(2-aminophenyl)benzamides. *Bioorg. Med. Chem. Lett.* **20**, 3142–3145. https://doi.org/10.1016/j.bmcl.2010.03.091 (2010).

32. Xu, M., Shen, C., Yang, J., Wang, Q. & Huang, N. Systematic investigation of docking failures in large-scale structure-based virtual screening. *ACS Omega* **7**, 39417–39428. https://doi.org/10.1021/acsomega.2c05826 (2022).

33. Dowling, D. P., Gantt, S. L., Gattis, S. G., Fierke, C. A. & Christianson, D. W. Structural studies of human histone deacetylase 8 and its site-specific variants complexed with substrate and inhibitors. *Biochemistry* **47**, 13554–13563. https://doi.org/10.1021/bi801610c (2008).

34. Eberhardt, J., Santos-Martins, D., Tillack, A. F. & Forli, S. 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 (2021).

35. Santos-Martins, D., Forli, S., Ramos, M. J. & Olson, A. J. AutoDock4Zn: an improved AutoDock force field for small-molecule docking to zinc metalloproteins. *J. Chem. Inf. Model.* **54**, 2371–2379. https://doi.org/10.1021/ci500209e (2014).

874 36. Hu, E. et al. Identification of novel isoform-selective inhibitors within class I histone deacetylases.
875 *J. Pharmacol. Exp. Ther.* **307**, 720–728. https://doi.org/10.1124/jpet.103.055541 (2003).

876 37. Whitehead, L. et al. Human HDAC isoform selectivity achieved via exploitation of the acetate
877 release channel with structurally unique small molecule inhibitors. *Bioorg. Med. Chem.* **19**,
878 4626–4634. https://doi.org/10.1016/j.bmc.2011.06.030 (2011).

879 38. Kunze, M. B. A. et al. Loop interactions and dynamics tune the enzymatic activity of the human
880 histone deacetylase 8. *J. Am. Chem. Soc.* **135**, 17862–17868. https://doi.org/10.1021/ja408184x
881 (2013).

882 39. Yang, X., Liu, Y., Gan, J., Xiao, Z.-X. & Cao, Y. FitDock: protein–ligand docking by template
883 fitting. *Brief. Bioinform.* **23**, bbac087. https://doi.org/10.1093/bib/bbac087 (2022).

884 40. Laboratory of Molecular Modeling and Design. admetSAR 3.0.
885 https://lmmd.ecust.edu.cn/admetsar3/predict_submit.php (accessed 26 April 2025).

886 41. Hopkins, A. L., Groom, C. R. & Alex, A. Ligand efficiency: a useful metric for lead selection.
887 *Drug Discov. Today* **9**, 430–431. https://doi.org/10.1016/S1359-6446(04)03069-7 (2004).

888 42. Hopkins, A. L., Keserü, G. M., Leeson, P. D., Rees, D. C. & Reynolds, C. H. The role of ligand
889 efficiency metrics in drug discovery. *Nat. Rev. Drug Discov.* **13**, 105–121.
890 https://doi.org/10.1038/nrd4163 (2014).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryData1.xlsx
- SupplementaryInformation.pdf