RSC Advances

PAPER

Published on 12 November 2026
Licensed under CC-BY-NC 4.0

Check for updates

Cite this: DOI: 10.1039/d6ra07609b

AI-guided prioritization of potential FAM171A2 ligands through pharmacophore-aware and receptor-based screening

Baihui Yan,*[a] Yang Lu,[a] Yanni Tong,[a] Shaolong Lin *[b] and Wei Li*[a]

FAM171A2 has recently emerged as a potential therapeutic target in Parkinson's disease because it selectively recognizes α-synuclein fibrils and promotes their neuronal uptake. However, ligand prioritization for this receptor remains challenging because of scarce known binders, the absence of an experimentally resolved ligand-bound structure and the requirement for blood–brain barrier permeability. Here, we established an AI-guided computational screening framework that integrates pharmacophore-aware retrieval, hierarchical receptor-based screening, blood–brain barrier permeability prediction and molecular dynamics-based assessment to prioritize potential FAM171A2 ligands. Starting from 1.41 million compounds, the workflow progressively narrowed the search space to two structurally representative candidates, E634-1386 and F707-2364. Both candidates showed more favorable *in silico* profiles than the reference blocker Bemcentinib across complementary computational assessments: E634-1386 achieved higher KarmaDock and CarsiDock scores, whereas F707-2364 yielded the most favorable AutoDock Vina score and MM/GBSA estimate. Molecular dynamics simulations further suggested stable predicted FAM171A2-bound complexes for both candidates, with lower complex RMSD, reduced radius of gyration and lower protein solvent-accessible surface area relative to Bemcentinib, while maintaining broadly similar residue-level flexibility patterns. Representative interaction analysis suggested that E634-1386 may rely more on polar contacts, whereas F707-2364 may engage the predicted binding region through broader hydrophobic complementarity. Together, these results prioritize E634-1386 and F707-2364 as computationally selected, experimentally testable FAM171A2 ligand candidates and provide a practical strategy for early-stage ligand prioritization against structurally challenging central nervous system targets.

Received 15th August 2026
Accepted 4th October 2026

DOI: 10.1039/d6ra07609b

rsc.li/rsc-advances

1 Introduction

Parkinson's disease (PD) is a progressive neurodegenerative disorder characterized by dopaminergic neuronal loss and the pathological misfolding, aggregation and intercellular propagation of α-synuclein (α-syn) fibrils across interconnected brain regions.[1,2] Its global burden is increasing rapidly, with analyses from the Global Burden of Disease 2021 study estimating that 11.77 million people were living with PD worldwide in 2021 and projecting continued growth in the coming decades.[3,4] Current therapeutic strategies mainly focus on symptomatic dopamine replacement or disease-modifying approaches targeting α-syn, LRRK2, GBA1/GCase, proteostasis and neuroinflammatory pathways.[5] However, no approved therapy has definitively halted disease progression, and few strategies directly target the neuronal uptake and spread of pathological α-syn species. In this context, FAM171A2 has recently emerged as a potential therapeutic entry point for interfering with α-syn fibril uptake. Current evidence links FAM171A2 to PD through human genetic association, elevated expression in PD brain and cerebrospinal fluid, neuron-specific functional validation, and preferential recognition of α-syn fibrils over monomers.[6,7] Consistent with this mechanism, Bemcentinib has been identified as a proof-of-concept small-molecule blocker that interferes with the interaction between the positively charged extracellular domain 1 of FAM171A2 and the C-terminal region of α-syn, thereby reducing fibril uptake in cellular and mouse models.[6]

Despite this promise, FAM171A2 remains a challenging target for ligand discovery. Publicly available chemical starting points are extremely limited: in the initial report, virtual screening of 7173 approved drugs produced only a small set of candidates, with Bemcentinib emerging as the principal experimentally validated blocker.[6,8] As a result, there is little established structure–activity information to support conventional ligand-based pharmacophore modeling or similarity-guided optimization. Structure-based design is also complicated by the absence of an experimentally resolved ligand-bound FAM171A2 structure. The proposed binding interface has been inferred mainly from biochemical analysis and

[a]Department of Anesthesia, The Second Affiliated Hospital of Xi'an Jiaotong University, 710004, China. E-mail: 13891960816@163.com

[b]Department of Bioengineering, College of Life Science, Dalian Minzu University, 116600, China. E-mail: ybh258404025@126.com; shaolonglin2023@163.com

© 2026 The Author(s). Published by the Royal Society of Chemistry

AlphaFold-Multimer modeling of the competitive interaction between FAM171A2 extracellular domain 1 and the C-terminal region of α-syn fibrils. In addition, Bemcentinib itself may have limited translational value for PD because its reported brain penetration appears insufficient for central nervous system applications. These limitations create three linked challenges for FAM171A2 ligand discovery: scarce ligand information, uncertainty in receptor-level binding-site definition, and the need to prioritize compounds with CNS-compatible blood–brain barrier permeability early in the screening process.

Recent advances in AI-driven drug discovery provide complementary strategies for addressing these challenges.[9] Pharmacophore-aware molecular embeddings can retrieve compounds that remain functionally related to a known blocker even when conventional two-dimensional similarity is limited.[10] Deep learning-based docking and conventional docking can provide complementary assessments of receptor compatibility, reducing dependence on a single scoring model.[11] In parallel, graph neural network-based molecular property predictors can introduce developability-oriented prioritization, including prediction of blood–brain barrier permeability (BBBP).[12] The integration of ligand-space retrieval, receptor-based structural screening and property-based filtering is therefore particularly attractive for difficult CNS targets such as FAM171A2, where known ligands are scarce, receptor structures remain uncertain and brain exposure is essential for therapeutic translation.

Here, we developed a hierarchical computational screening framework for prioritizing potential FAM171A2 ligands. Starting from a ChemDiv library of 1 413 202 compounds, we first used Ouroboros to perform pharmacophore-aware embedding retrieval with Bemcentinib as the reference compound. Retrieved candidates were then subjected to sequential receptor-based screening using KarmaDock, CarsiDock and AutoDock Vina, followed by Chemprop-based BBBP prediction and diversity-aware clustering. Representative candidates were finally evaluated by molecular dynamics simulations to assess binding-mode stability and receptor–ligand interaction patterns. This integrated workflow was designed to reduce large chemical space efficiently while combining ligand-level similarity, receptor-level compatibility and CNS-oriented property prioritization for FAM171A2 ligand discovery.

2 Methods

2.1 Overall workflow

We established a hierarchical computational screening workflow to identify experimentally testable FAM171A2 ligand candidates. The workflow was motivated by three practical constraints of this target: only a small number of reported blockers are currently available, no ligand-bound FAM171A2 structure has been experimentally resolved, and any PD-oriented candidate must be compatible with central nervous system exposure. Rather than relying on a single ligand-based or structure-based method, we therefore used a staged strategy that first maintained broad chemical-space coverage and then progressively imposed receptor- and property-level constraints.

Bemcentinib, a reported blocker of FAM171A2-mediated α-syn fibril uptake, was used as the reference compound for the initial ligand retrieval step. Compounds from the ChemDiv library were embedded using Ouroboros,[13] and molecules with pharmacophore-aware embedding similarity to Bemcentinib were retained for receptor-based screening (Fig. 1A). These candidates were sequentially evaluated against an AlphaFold3 (ref. 14)-predicted FAM171A2 structure using KarmaDock,[15] CarsiDock[16] and AutoDock Vina[17] (Fig. 1B). At each docking stage, Bemcentinib served as the positive-control baseline, and compounds were retained only when they showed more favorable predicted receptor compatibility according to the scoring convention of the corresponding method. The resulting molecules were then filtered using a Chemprop[12] blood–brain barrier permeability model and predicted BBBP-positive compounds were clustered by K-means to preserve chemical diversity (Fig. 1C and D). Finally, representative candidates selected from the clusters were subjected to molecular dynamics simulations to examine the stability of the predicted complexes and the persistence of their binding modes (Fig. 1E). This staged workflow reduced the initial 1 413 202-compound ChemDiv library to two representative candidates for further computational evaluation.

2.2 Pharmacophore-aware embedding retrieval using Ouroboros

Ouroboros[13] was used as the initial ligand-retrieval module because its pharmacophore-aware molecular embeddings can capture feature-level similarity that may not be fully reflected by conventional two-dimensional fingerprints. Bemcentinib, a reported blocker of FAM171A2-mediated α-syn fibril uptake, was selected as the reference compound. Molecules from the ChemDiv library were encoded into the Ouroboros latent space, and their embedding similarity to Bemcentinib was calculated using the similarity metric implemented in the Ouroboros workflow.

Compounds with Ouroboros similarity scores ≥ 0.5 were retained for downstream screening. This cutoff has methodological support from a previous LRRK2 benchmark, in which 51.8% of active–active pairs and 1.2% of active–decoy pairs passed the Ouroboros similarity filter at 0.5.[18] The maximum observed similarity score was approximately 0.7321, and no upper similarity cutoff was imposed. For supplementary evaluation, 10 000 compounds were randomly sampled from the retained similarity range and 10 000 from compounds with similarity below 0.5. Both groups were processed using the same docking workflow, and their valid score distributions were compared.

This retrieval step reduced the ChemDiv library from 1 413 202 compounds to 201 091 candidates.

2.3 Hierarchical receptor-based screening

The receptor-based screening stage consisted of three sequential docking modules: KarmaDock[15] for large-scale prescreening, CarsiDock[16] for geometry-aware reranking, and AutoDock Vina[17] for orthogonal classical docking assessment. The docking region was selected by integrating previous biochemical and

© 2026 The Author(s). Published by the Royal Society of Chemistry

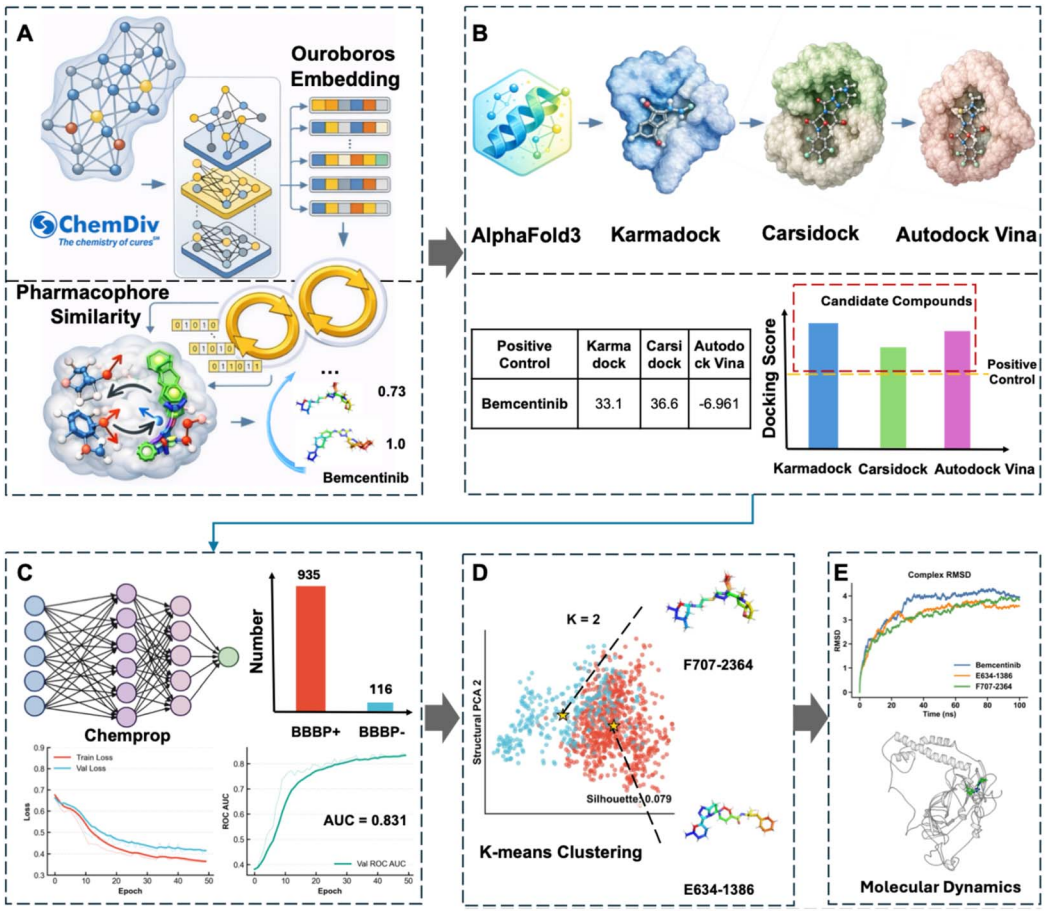


Fig. 1 Hierarchical AI-guided computational screening framework for prioritizing potential FAM171A2 ligands. (A) Bemcentinib was used as the reference blocker for pharmacophore-aware retrieval from the ChemDiv library using Ouroboros embeddings. (B) Retrieved compounds were sequentially screened against an AlphaFold3-predicted FAM171A2 structure using KarmaDock, CarsiDock and AutoDock Vina, with Bemcentinib serving as the positive-control baseline for comparative filtering. (C) Receptor-screened compounds were further evaluated using a Chemprop-based blood–brain barrier permeability (BBBP) prediction model, yielding predicted BBBP-positive and BBBP-negative candidates. (D) Predicted BBBP-positive candidates were clustered in fingerprint-derived chemical space using K-means, and two centroid-nearest representative candidates, F707-2364 and E634-1386, were selected. (E) The final representative candidates were advanced to molecular dynamics simulations to assess predicted complex stability and binding-mode persistence.

structural-modeling evidence concerning extracellular FAM171A2 recognition with the AlphaFold 3-predicted[14] FAM171A2–Bemcentinib model and complementary Proteins-Plus[19] pocket analysis. Local confidence was assessed using Cα pLDDT, with extracellular ligand-contacting residues defined by a minimum heavy-atom distance of ≤5 Å from Bemcentinib. Two additional receptor models were evaluated by D1 structural alignment and comparative docking of the three ligands.

At each docking stage, compounds were ranked according to the scoring convention of the corresponding method. For KarmaDock and CarsiDock, higher scores were interpreted as more favorable predicted receptor compatibility, whereas for AutoDock Vina, more negative binding energies were considered more favorable. Compounds were retained only when they performed more favorably than Bemcentinib according to the corresponding scoring criterion. This design was intended to reduce dependence on a single docking algorithm and to prioritize compounds supported by complementary receptor-based scoring schemes.

2.3.1 KarmaDock-based prescreening. KarmaDock[15] was used as the first docking stage because it provides efficient large-scale pose generation and scoring through an $E(n)$-equivariant graph neural network. Given a receptor pocket P and a ligand L, KarmaDock predicts a ligand pose $\hat{x}$ together with a docking score S_{Karma} that reflects the favorability of the protein–ligand interaction. In abstract form, the model can be written as:

$$(\hat{x}, S_{\mathrm{Karma}}) = f_\theta(P, L) \quad (1)$$

where f_θ denotes the trained KarmaDock model, $\hat{x}$ is the predicted ligand pose in the receptor binding site, and S_{Karma} is the corresponding docking score.

Because KarmaDock internally combines pose prediction and interaction scoring, ligand ranking in this stage was performed directly on the basis of the predicted score:

$$R_{\mathrm{Karma}}(L_i) = S_{\mathrm{Karma}}(L_i, P) \quad (2)$$

© 2026 The Author(s). Published by the Royal Society of Chemistry

where a higher score indicates a more favorable predicted interaction. Compounds were sorted in descending order of R_{Karma}, and the top-ranked subset was forwarded to the next screening stage.

2.3.2 CarsiDock-based reranking. CarsiDock was used as the second docking stage to further refine candidate ranking.[20] It predicts protein–ligand distance relationships and reconstructs ligand poses through geometry-guided optimization. For a given ligand L and receptor pocket P, the docking pose $\tilde{x}$ can be expressed as the solution of a geometry-consistency objective:

$$\tilde{x} = \arg\min_x L_{\mathrm{geo}}(x; P, L) \tag{3}$$

where L_{geo} represents the geometry-consistency objective induced by the learned protein–ligand distance constraints.

The corresponding CarsiDock score was then used for reranking

$$R_{\mathrm{Carsi}}(L_{\mathrm{i}}) = S_{\mathrm{Carsi}}(L_{\mathrm{i}}, P) \tag{4}$$

where S_{Carsi} denotes the final docking score reported by CarsiDock for ligand L_{i}. In this study, compounds were sorted according to R_{Carsi}, and only candidates with CarsiDock performance superior to the positive control were retained for the subsequent validation step.

2.3.3 AutoDock Vina validation. For each compound, the pose with the lowest Vina score was retained as the representative docking result. Ligands were ranked in ascending order of Vina score, with more negative values indicating more favorable predicted binding. Bemcentinib was docked using the same protocol and produced a Vina score of -6.961 kcal mol^{-1}. This value was used as the positive-control cutoff, and compounds with Vina scores lower than -6.961 kcal mol^{-1} were retained for downstream BBBP prediction and diversity-aware prioritization.

To reduce dependence on a single AI-based docking framework, shortlisted compounds were further evaluated using AutoDock Vina as a classical physics-based docking method. Docking was performed within a predefined binding pocket of FAM171A2 using a grid box of $30 \times 30 \times 30$ Å centered at $(x, y, z) = (82.474, 93.909, 85.977)$. The exhaustiveness parameter was set to 32, and multiple poses were generated for each ligand under the default Vina search protocol. For each compound, the pose with the lowest Vina score was retained as the representative result, and ligands were ranked in ascending order of docking score, with more negative values indicating more favorable predicted binding. Compounds were subsequently compared with the positive control Bemcentinib and used for downstream prioritization.

2.4 BBBP prediction using Chemprop

Because FAM171A2 is a CNS-related target, blood–brain barrier permeability was included as an explicit filtering criterion. The BBBP dataset was obtained from SSPBench.[21] After salt removal, charge neutralization, and deduplication, the dataset contained 2330 compounds, including 1741 BBBP-positive and 589 BBBP-negative compounds. A Chemprop v2 binary classifier was developed using scaffold-based training, validation, and test partitions at an approximately 8 : 1 : 1 ratio.[12] In Chemprop, molecules are represented as graphs and encoded by a directed message passing neural network. Let $h_{(u\to v)}^{(t)}$ denote the hidden representation on the directed bond from atom u to atom v at iteration t. The message passing and update steps are given by:

$$m_{(u\to v)}^{(t)} = \sum\nolimits_{-}(w \in N(u)\{v\})\, h_{(w\to u)}^{(t-1)} \tag{5}$$

$$h_{(u\to v)}^{(t)} = \mathrm{sigma}(W[x_{(u\to v)}; m_{(u\to v)}^{(t)}] + b) \tag{6}$$

After readout, the molecular representation h_{G} was mapped to a BBBP probability:

$$\hat{p} = \mathrm{sigmoid}(f(h_{\mathrm{G}})) \tag{7}$$

The model was optimized using binary cross-entropy loss:

$$L_{\mathrm{BCE}} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log\left(\hat{p}_i\right) + (1-y_i)\log\left(1-\hat{p}_i\right)\right] \tag{8}$$

where y_i is the ground-truth BBBP label and $\hat{p}_i$ is the predicted probability for compound i. The model used three message-passing steps, a hidden dimension of 300, one feed-forward hidden layer with 300 units, and a dropout rate of 0.1. Training used a batch size of 64 and class-balanced sampling for up to 100 epochs. The learning rate increased from 10^{-4} to 10^{-3} over two warmup epochs and subsequently decayed to 10^{-4}, with validation loss used for checkpoint monitoring. Binary predictions used a probability threshold of 0.5. Calibration was evaluated using a reliability diagram with ten equal-width bins, the Brier score, log loss, and expected calibration error (ECE); 95% confidence intervals for performance metrics were estimated using 2000 class-stratified bootstrap resamples of the held-out test set. The trained model was then used to predict BBBP for the receptor-screened compounds, and BBBP-positive molecules were retained for subsequent clustering and prioritization.

2.5 Diversity-aware clustering by K-means

To preserve chemical diversity among BBBP-positive candidates, clustering was performed on ECFP4 fingerprints generated by RDKit (radius = 2, 2048 bits).[22] Invalid SMILES were removed before analysis. The fingerprints were reduced to 50 dimensions by PCA for clustering, and a separate two-dimensional PCA projection was used for visualization. MiniBatchKMeans[23] was applied with $K = 2$ to 6, and clustering quality was assessed using the silhouette score:

$$s(i) = (b(i) - a(i))/\max a(i), b(i) \tag{9}$$

Based on the silhouette score trend, $K = 2$ was selected. For each cluster, the representative compound was defined as the molecule closest to the cluster centroid in the reduced feature space according to Euclidean distance:

$$d(x_i, \mu_k) = \|x_i - \mu_k\|_2 \tag{10}$$

The two centroid-nearest molecules were retained as final representatives for downstream molecular dynamics

© 2026 The Author(s). Published by the Royal Society of Chemistry

simulations. As a supplementary post-selection assessment, admetSAR3.0 (ref. 24) was used to compare Bemcentinib and the two candidates in terms of physicochemical properties, drug-likeness, ADME properties, microsomal stability, and toxicity-related endpoints, including CYP, and transporter.

2.6 Molecular dynamics simulations

Three independent MD simulations per complex were used to assess repeatability using the same initial receptor model. MM/GBSA results were summarized as the mean ± SD of the three trajectory-averaged estimates. All molecular dynamics simulations and MM/GBSA analyses were performed using the AMBER24 software suite and its associated tools, including cpptraj and MMPBSA.py, to evaluate the stability of the final FAM171A2–ligand complexes.[25] The protein was parameterized with ff19SB, and ligands were described using GAFF2 with AM1-BCC charges. Complexes were solvated in a TIP3P water box with 10 Å padding and neutralized with counterions. After minimization, heating and equilibration, 100 ns production simulations were carried out under the NPT ensemble at 300 K and 1 bar with a 2 fs time step.

Trajectory analysis was performed using cpptraj after alignment to the protein backbone. Structural stability was assessed by RMSD:

$$\mathrm{RMSD}(t) = \mathrm{sqrt}\left((1/N)\mathrm{sum}_{(i=1\ \mathrm{to}\ N)}\left|\left|r_i(t) - r_i^{(\mathrm{ref})}\right|\right|^2\right) \quad (11)$$

Residue flexibility was assessed by RMSF:

$$\mathrm{RMSF}_i = \mathrm{sqrt}\left((1/T)\mathrm{sum}_{(t=1\mathrm{to}T)}\left|\left|r_i(t) - r_i^{(\mathrm{mean})}\right|\right|^2\right) \quad (12)$$

Binding free energies were estimated using MM/GBSA on MD snapshots. The binding free energy was computed as:

$$\Delta G_{\mathrm{bind}} = G_{\mathrm{complex}} - G_{\mathrm{protein}} - G_{\mathrm{ligand}} \approx \Delta E_{\mathrm{MM}} + \Delta G_{\mathrm{solv}} - T\Delta S \quad (13)$$

where ΔE_{MM} includes van der Waals and electrostatic interactions, and ΔG_{solv} is decomposed into polar and nonpolar components:

$$G = E_{\mathrm{MM}} + G_{\mathrm{GB}} + G_{\mathrm{SA}} \quad (14)$$

In this work, the entropic term $(-T\Delta S)$ was not explicitly calculated due to computational cost and because relative comparisons among similar candidates are often dominated by enthalpic and solvation contributions. Accordingly, we report $\Delta G_{\mathrm{MM/GBSA}}$ as an approximate, comparative binding metric.

3 Results and discussion

3.1 Pharmacophore-aware embedding retrieval reduced the initial compound library while maintaining candidate diversity

We first applied Ouroboros-based pharmacophore-aware embedding retrieval to the ChemDiv library using Bemcentinib as the reference compound. Retaining compounds with embedding similarity scores ≥0.5 reduced the initial search space from 1 413 202 compounds to 201 091 candidates, corresponding to an approximately sevenfold reduction of the original library.

The retained similarity scores were concentrated near the lower cutoff of 0.5, with fewer compounds observed at higher scores and a maximum of approximately 0.7321 (Fig. 2A). This pattern suggests that the retrieval step preserved a large pool of moderately related compounds rather than concentrating only on close analogs of Bemcentinib. For an early-stage screening workflow, such a candidate pool is useful because it reduces the computational burden of downstream docking while maintaining sufficient diversity for receptor-based prioritization. The supplementary comparison showed higher mean KarmaDock scores (27.678 *versus* 23.391) and CarsiDock scores (28.601 *versus* 23.666) for the similarity ≥0.5 group than for the similarity <0.5 group (Fig. S1). Higher median scores and consistent rightward shifts in both distributions further supported the enrichment of higher-scoring compounds through similarity-based prescreening.

To further examine the embedding-space coverage of the retrieved candidates, we projected the selected compounds into a two-dimensional t-SNE space and visualized their distribution relative to the broader embedding landscape (Fig. 2B). The retained candidates were broadly dispersed across the mapped space rather than concentrated in a single local region, indicating that Ouroboros retrieval captured multiple regions of the pharmacophore-aware embedding space. In other words, the retrieval step reduced the library size without collapsing the candidate set into a narrow cluster of close Bemcentinib analogs. Although t-SNE provides only a qualitative visualization of high-dimensional relationships, this distribution supports the use of the retrieved compounds as a diverse starting pool for subsequent receptor-based screening.[26]

3.2 Hierarchical receptor-based screening narrowed the candidate space using Bemcentinib as a comparative baseline

The extracellular region and D1 domain had mean Cα pLDDT values of 90.87 and 95.12, respectively, while the 16 extracellular ligand-contacting residues had a mean of 90.64. ProteinsPlus pockets P_0 and P_1 had druggability scores of 0.809 and 0.740. Trp109, Arg110 and Ala118 were within 5 Å of Bemcentinib, supporting partial overlap with the literature-proposed recognition interface. Two additional models showed D1 Cα RMSDs of 0.777 and 0.782 Å relative to the original model (Fig. S3).

Using the 201 091 compounds retained by Ouroboros as input, we next performed hierarchical receptor-based screening with KarmaDock and CarsiDock. Bemcentinib was included as a positive-control reference compound and achieved docking scores of 33.1 in KarmaDock and 36.6 in CarsiDock against the AlphaFold3-predicted FAM171A2 structure. These values were used as practical baseline thresholds for comparative prioritization (Fig. 3C). Using a literature-supported blocker as the reference provided a more interpretable filtering criterion than absolute score ranking alone, although the resulting docking

© 2026 The Author(s). Published by the Royal Society of Chemistry

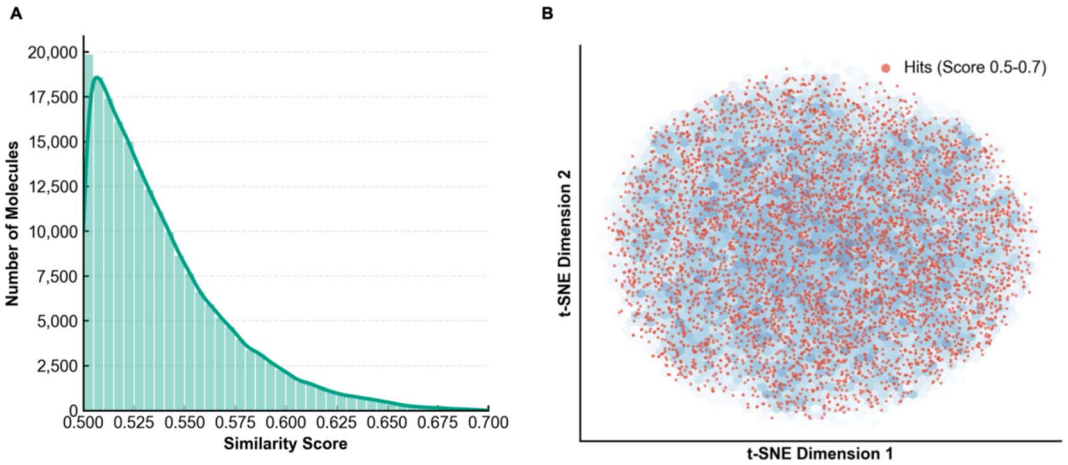


Fig. 2 Ouroboros-based pharmacophore-aware retrieval of candidate molecules. (A) Distribution of Ouroboros embedding similarity scores relative to Bemcentinib among retained compounds (similarity ≥0.5). (B) t-SNE projection of the retrieved candidate set, showing its distribution across the pharmacophore-aware embedding space.

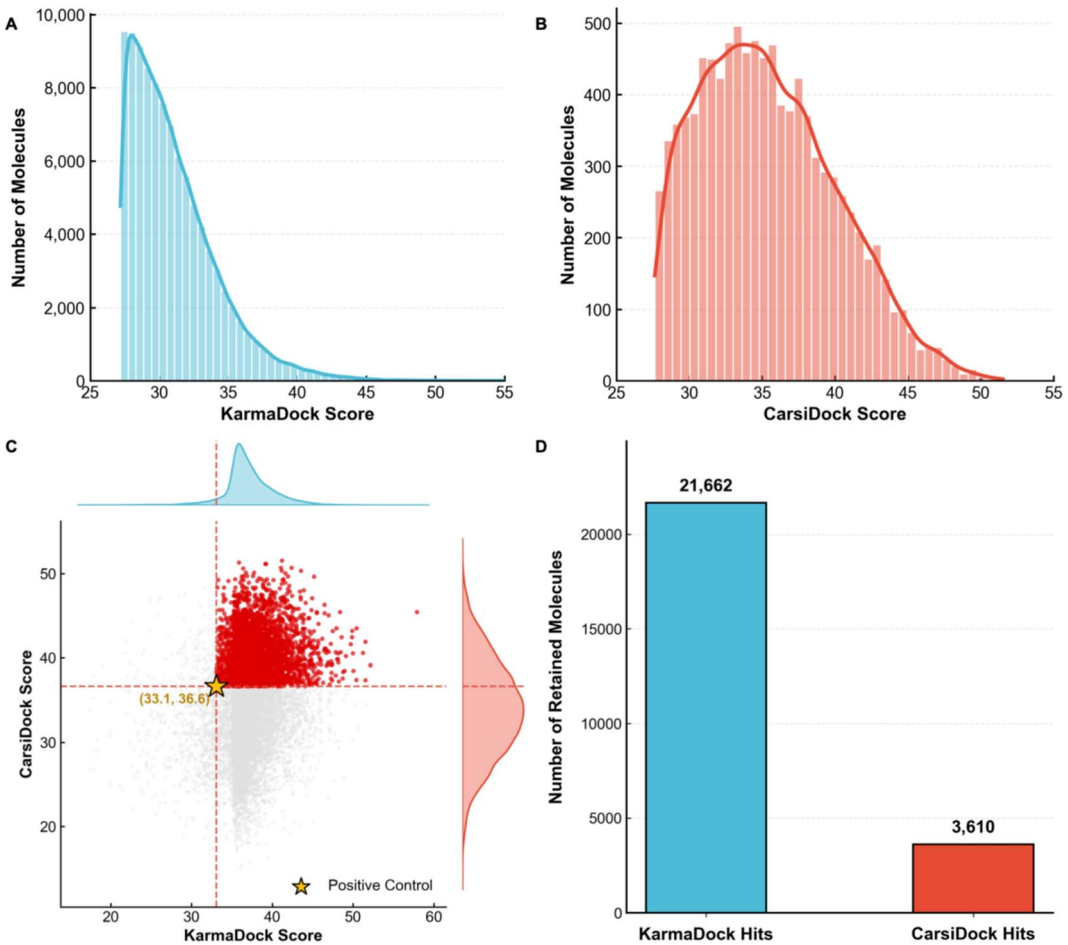


Fig. 3 Hierarchical receptor-based screening relative to the Bemcentinib baseline. (A) Distribution of KarmaDock scores for the Ouroboros-retrieved candidates. (B) Distribution of CarsiDock scores for candidates retained after KarmaDock screening. (C) Joint distribution of KarmaDock and CarsiDock scores. The yellow star marks Bemcentinib, with reference scores of 33.1 in KarmaDock and 36.6 in CarsiDock. Red dashed lines indicate the corresponding Bemcentinib-based cutoffs. (D) Numbers of compounds retained after applying the KarmaDock and CarsiDock cutoffs, yielding 21 662 and 3610 candidates, respectively.

© 2026 The Author(s). Published by the Royal Society of Chemistry

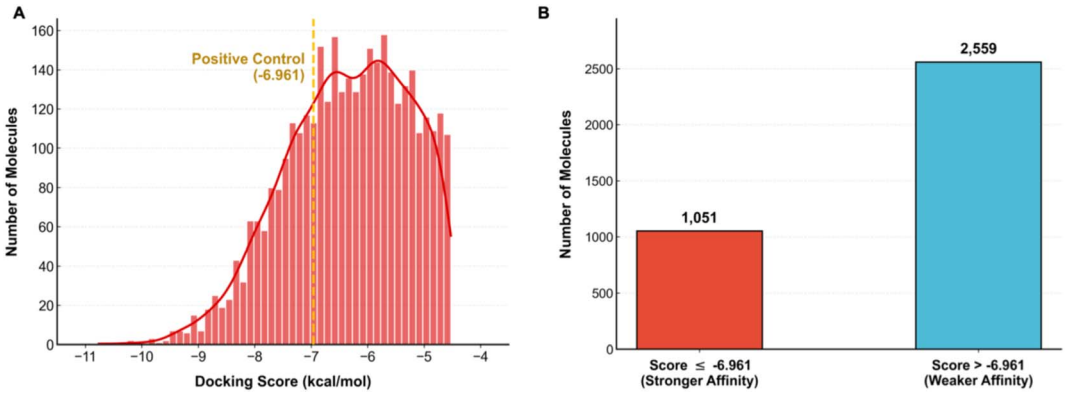


Fig. 4 AutoDock Vina-based docking assessment relative to the Bemcentinib baseline. (A) Distribution of Vina docking scores for compounds retained after KarmaDock and CarsiDock screening. The dashed line marks the Bemcentinib reference score of −6.961 kcal mol^{-1}. (B) Numbers of compounds with Vina scores more favorable or less favorable than the Bemcentinib-based cutoff.

scores should be interpreted as computational prioritization metrics rather than direct measures of binding activity.

KarmaDock was used as the first receptor-based screening stage to evaluate the Ouroboros-retrieved candidates at large scale. Its score distribution spanned roughly 27–55, with most compounds concentrated in the lower-to-middle range and a long tail extending toward higher values (Fig. 3A). Applying the Bemcentinib-based cutoff retained 21 662 compounds with KarmaDock scores above 33.1. This first docking stage therefore removed a large fraction of molecules with lower predicted receptor compatibility while retaining a broad candidate pool for subsequent refinement. At this stage, the purpose was not to define final hits, but to preserve a sufficiently large set of compounds with plausible predicted compatibility with the FAM171A2 binding region.

CarsiDock was then applied as a second-stage reranking method to introduce an additional geometry-aware docking criterion. Compared with KarmaDock, the CarsiDock score distribution was narrower and shifted toward a more restricted range (Fig. 3B), and 3610 compounds exceeded the Bemcentinib-based threshold of 36.6 (Fig. 3D). This reduction reflects the different scoring assumptions of the two docking methods. Whereas KarmaDock was used for high-throughput pocket-level prescreening, CarsiDock places greater emphasis on protein–ligand distance relationships and geometry-guided pose evaluation. Consequently, compounds prioritized by KarmaDock were not uniformly retained by CarsiDock, suggesting that the two methods provided partially non-overlapping receptor-based prioritization signals.

This partial complementarity is further illustrated by the joint score distribution (Fig. 3C). Compounds surpassing the Bemcentinib thresholds in both docking schemes were located in the upper-right quadrant, whereas many molecules with favorable scores in one method failed to meet the cutoff in the other.[20] In practical terms, KarmaDock served as a high-throughput structural prescreen, while CarsiDock acted as a geometry-aware refinement step that reduced candidates favored only under a single scoring framework. Such a hierarchical design is useful for FAM171A2 because experimentally validated ligands remain scarce and the receptor model is prediction-derived rather than obtained from a ligand-bound experimental structure. Under these conditions, requiring agreement between complementary docking methods provides a more conservative computational prioritization strategy than reliance on either method alone. Uncertainty in receptor conformation and pocket definition may alter predicted poses, scores, and ligand rankings; using multiple docking methods with the same receptor does not eliminate this shared source of uncertainty.

Taken together, the receptor-based stage reduced the candidate space from 201 091 to 21 662 and then to 3610 compounds, while enriching for molecules that performed favorably relative to Bemcentinib in both docking frameworks. These results support the use of a coarse-to-fine docking strategy for FAM171A2 candidate prioritization, in which large-scale receptor-compatibility screening and geometry-aware reranking are combined before subsequent AutoDock Vina reassessment, BBBP prediction and diversity-aware clustering.

3.3 AutoDock Vina provided an orthogonal docking filter relative to the Bemcentinib baseline

We next evaluated the 3610 compounds retained after the KarmaDock–CarsiDock stage using AutoDock Vina as an orthogonal conventional docking method with an empirical scoring function. Bemcentinib yielded a reference score of

Table 1 Comparison of molecular representation models for BBB permeability prediction

	Chemprop	Uni-mol	AttentionFP	MolFormer	KPGT	MolCLR
ROC-AUC	0.831	0.818	0.790	0.815	0.822	0.787

© 2026 The Author(s). Published by the Royal Society of Chemistry

−6.961 kcal mol^{-1} under the same docking protocol, which was used as the Bemcentinib-based cutoff for comparative prioritization. An additional assessment of the second-ranked ProteinsPlus pocket yielded a Bemcentinib docking score of −5.681. Using the primary-pocket cutoff of −6.961 kcal mol^{-1}, 1051 compounds achieved more favorable Vina scores, whereas 2559 compounds did not (Fig. 4A and B). This step further reduced the candidate set by approximately 3.4-fold before BBBP prediction and diversity-aware clustering.

The contribution of this stage was not limited to numerical reduction. In contrast to KarmaDock and CarsiDock, which rely on learned representations for pose evaluation and scoring, AutoDock Vina applies a conventional empirical scoring function. The observation that only a subset of compounds remained favorable under Vina suggests that compounds prioritized by the preceding AI-based docking models were not automatically supported by the empirical scoring scheme. The Vina-retained subset therefore represents compounds that remained favorable under both learned docking models and an empirical docking assessment.

The Vina score distribution also provides context for the use of Bemcentinib as the comparative cutoff (Fig. 4A). Bemcentinib was positioned near the transition region of the distribution rather than at an extreme tail, suggesting that the cutoff was neither trivially permissive nor overly restrictive for this screened set. This comparative strategy is useful for FAM171A2, where available experimentally supported ligands remain sparse and the receptor structure is prediction-derived rather than resolved in a ligand-bound experimental complex. Under these conditions, prioritization relative to a literature-supported reference blocker is more interpretable than ranking by absolute docking score alone.

Overall, the Vina step served as a useful orthogonal checkpoint in the receptor-based screening cascade. It did not experimentally validate binding, but it helped remove compounds favored only by the preceding AI-based docking models and retained a smaller subset that remained favorable under a conventional empirical scoring function. This subset was therefore used for the subsequent BBBP prediction and diversity-aware prioritization steps.

3.4 Chemprop-based BBBP prediction introduced a CNS-oriented property filter

Given the CNS-related therapeutic context of FAM171A2 in Parkinson's disease, predicted blood–brain barrier permeability (BBBP) was incorporated as an additional prioritization

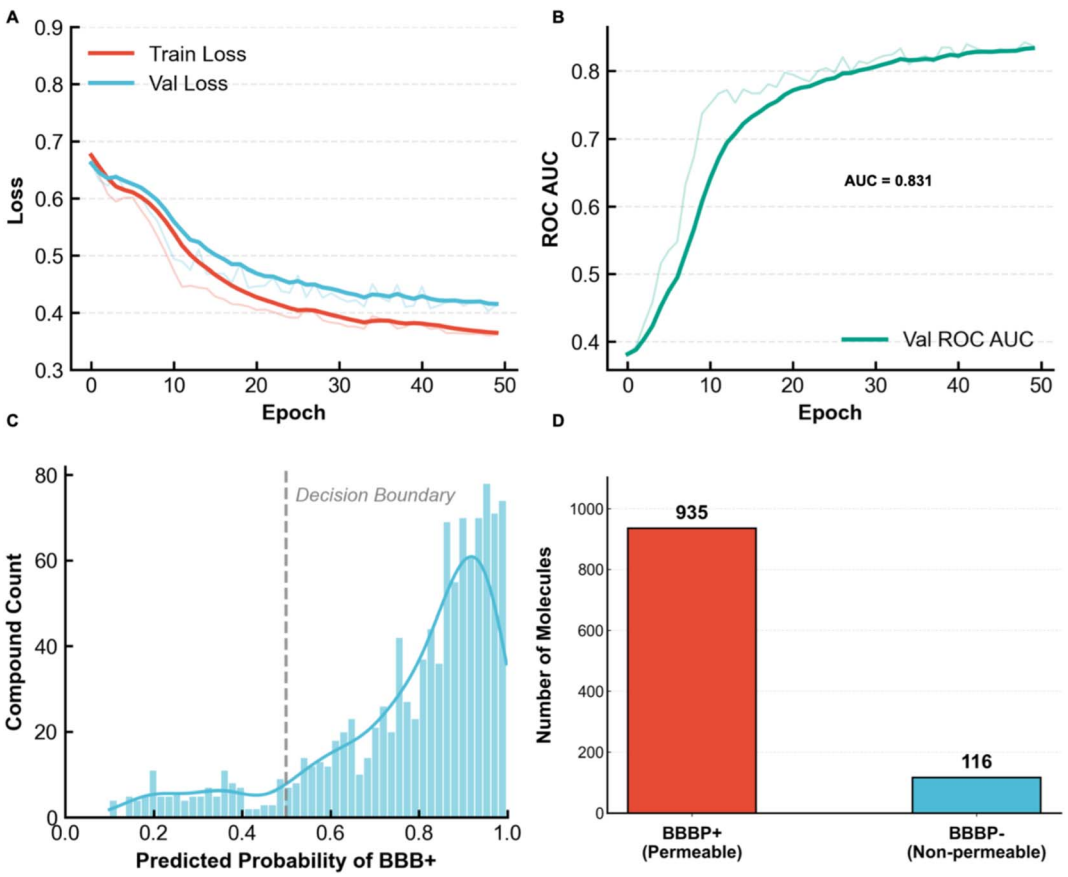


Fig. 5 Chemprop-based BBBP prediction for compounds retained after Vina reassessment. (A) Training and validation loss curves of the Chemprop BBBP classifier. (B) Validation ROC AUC across training epochs, reaching 0.831. (C) Distribution of predicted BBBP probabilities for the 1051 compounds retained after AutoDock Vina-based filtering. The dashed line marks the classification threshold. (D) Numbers of compounds classified as predicted BBBP-positive and predicted BBBP-negative, yielding 935 and 116 molecules, respectively.

© 2026 The Author(s). Published by the Royal Society of Chemistry

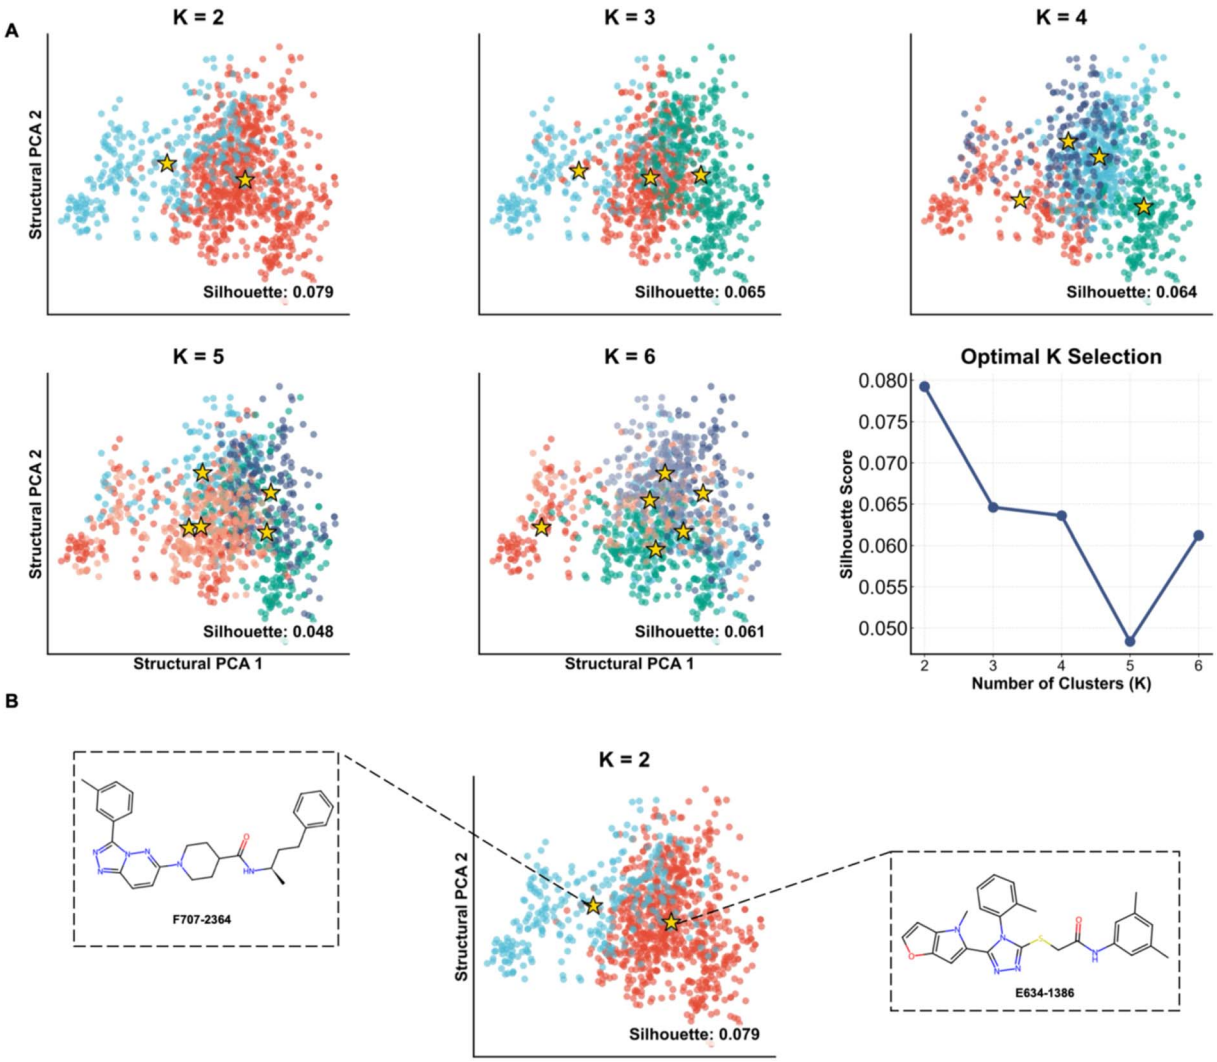


Fig. 6 Diversity-aware clustering of BBBP-positive candidates. (A) *K*-means clustering of the BBBP-positive compounds with $K = 2$ to 6, together with the corresponding silhouette scores and the silhouette-score trend used for selecting the optimal number of clusters. (B) Final $K = 2$ solution and centroid-nearest representative compounds selected for downstream molecular dynamics simulations.

criterion after receptor-based screening.[12,27] A Chemprop-based BBBP classifier was trained using 2330 annotated compounds and compared with several representative molecular representation models, including Uni-Mol,[28] AttentionFP,[29] MolFormer,[30] KPGT[31] and MolCLR[32] (Table 1). Among the evaluated models, Chemprop achieved the highest validation ROC-AUC value (0.831), indicating its competitive performance for BBBP classification under the benchmark setting used in this study (Fig. 5).

On the scaffold-separated held-out test set of 234 compounds, including 183 BBBP-positive and 51 BBBP-negative compounds, Chemprop achieved a ROC-AUC of 0.888 (95% CI: 0.833–0.931). At a classification threshold of 0.5, sensitivity (recall), specificity, and precision were 0.825, 0.725, and 0.915, respectively. The Brier score, log loss, and ECE were 0.131, 0.406, and 0.140, respectively (Fig. S4). These results support the model's ability to generalize to compounds with scaffolds absent from the training set and its use for BBBP-based prioritization in subsequent virtual screening.

The Chemprop model was subsequently applied to the 1051 compounds retained after AutoDock Vina-based reassessment. Based on the predefined classification threshold, 935 compounds were assigned as predicted BBBP-positive candidates, whereas 116 compounds were classified as predicted BBBP-negative candidates (Fig. 5C and D). These predictions provided an additional property-oriented criterion to prioritize compounds with potentially favorable CNS-related characteristics, while reducing candidates with less favorable predicted permeability profiles.

The incorporation of BBBP prediction complements receptor-based screening because favorable docking scores alone do not necessarily indicate suitable developability for CNS applications. By integrating predicted permeability with receptor compatibility, the workflow considered both target-related computational prioritization and a key pharmacokinetic property relevant to brain-directed drug discovery. This consideration is particularly relevant for FAM171A2, as the reported reference blocker Bemcentinib has limited brain

© 2026 The Author(s). Published by the Royal Society of Chemistry

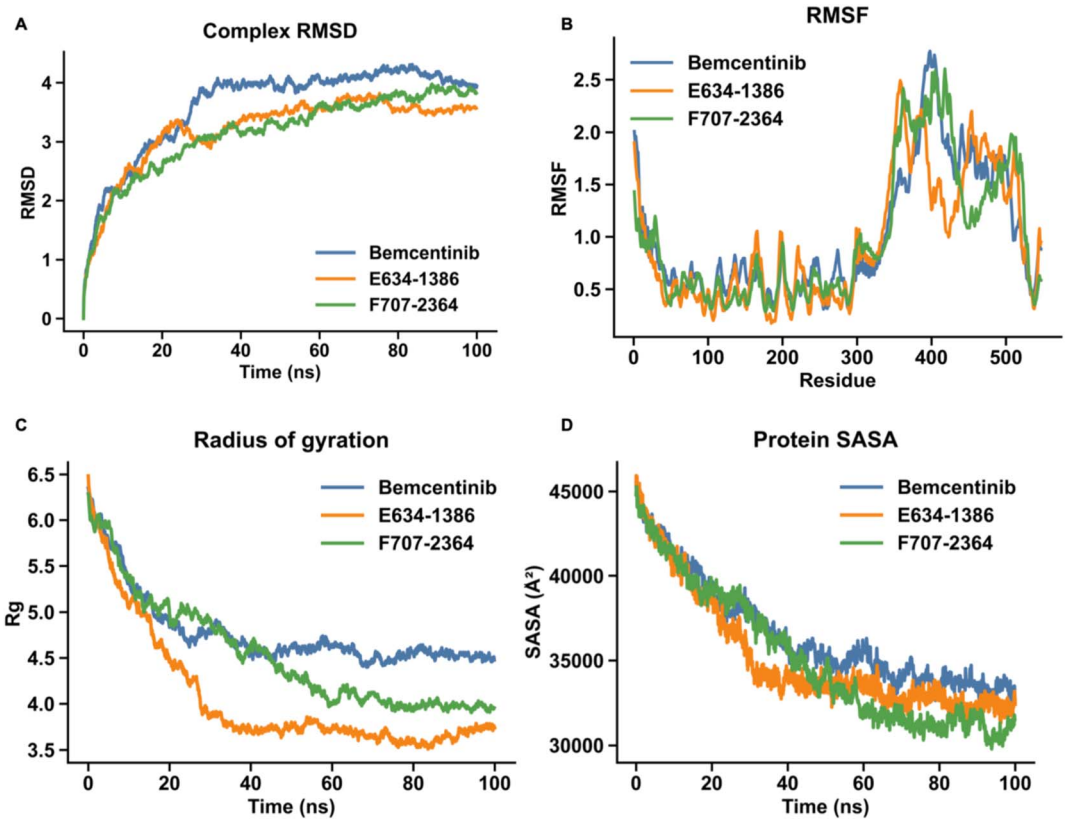


Fig. 7 Molecular dynamics analyses of the three protein–ligand complexes. (A) Complex RMSD. (B) RMSF. (C) Radius of gyration (R_g). (D) Protein solvent-accessible surface area (SASA).

penetration, highlighting the importance of early CNS-oriented property assessment.

The distribution of predicted BBBP probabilities further provided a continuous measure of model preference beyond binary classification. Although most Vina-retained compounds were assigned to the predicted BBBP-positive class, the prediction scores covered a broad range, reflecting different confidence levels within the retained chemical space (Fig. 5C). Following BBBP-based filtering, 935 candidates were retained for subsequent diversity-oriented clustering and representative compound selection.

3.5 Diversity-aware clustering selected two representative candidates from the predicted BBBP-positive set

After Chemprop-based filtering, 935 predicted BBBP-positive compounds remained. Rather than selecting candidates solely by docking score or BBBP probability, we introduced a diversity-aware clustering step to obtain a small number of structurally representative molecules for molecular dynamics simulations. ECFP4 fingerprints were used to describe the retained compounds, followed by dimensionality reduction and K-means clustering.[33,34] This step was intended to reduce redundancy among closely related analogs while preserving representative regions of the final candidate space.

We evaluated K values from 2 to 6 and visualized the resulting partitions in a two-dimensional PCA projection (Fig. 6A). Across these settings, the compounds formed a largely continuous chemical landscape, with partial overlap between clusters rather than clearly separated scaffold groups. Consistent with this observation, the silhouette scores were low for all tested K values.[35,36] Among the tested options, $K = 2$ produced the highest silhouette score, although the absolute value remained modest at 0.079. We therefore used $K = 2$ not as evidence of two sharply separated chemical classes, but as a simple and practical partition for diversity-aware representative selection.

Under the $K = 2$ solution, the two cluster centroids were located in different regions of the projected chemical space (Fig. 6B). The molecules closest to these centroids were selected

Table 2 Comparative computational evaluation of Bemcentinib, E634-1386 and F707-2364 using docking scores and MM/GBSA estimates[a]

	KarmaDock	CarsiDock	AutoDock Vina	MD (MMGBSA)
Bemcentinib	33.1	36.6	−6.961	−50.36
E634-1386	39.5	38.4	−7.112	−52.65
F707-2364	35.7	37.1	−7.592	−56.11

[a] Higher scores indicate better predicted binding in KarmaDock and CarsiDock, whereas lower (more negative) values indicate more favorable binding in AutoDock Vina and MM/GBSA.

© 2026 The Author(s). Published by the Royal Society of Chemistry

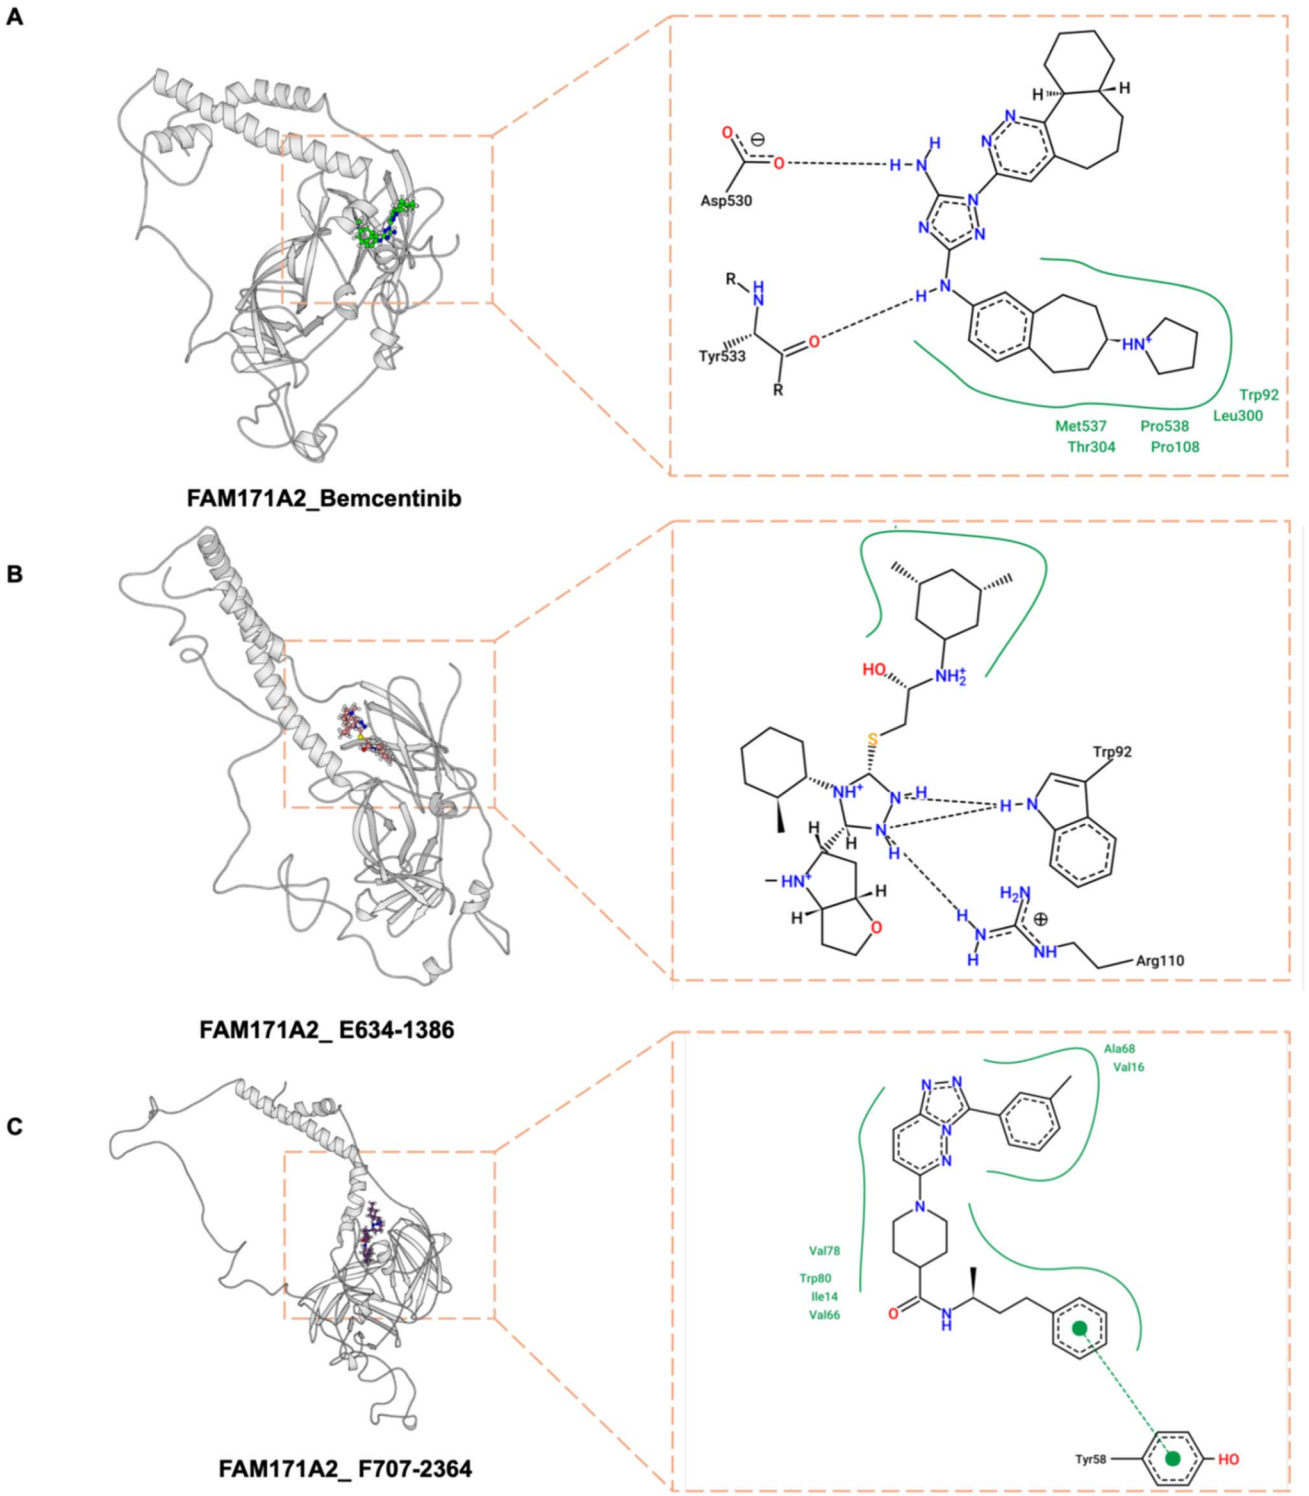


Fig. 8 Representative binding modes and interaction patterns of Bemcentinib, E634-1386 and F707-2364 in the predicted FAM171A2 binding region. (A) Representative conformation of the FAM171A2–Bemcentinib complex and corresponding two-dimensional interaction diagram. (B) Representative conformation of the FAM171A2–E634-1386 complex, showing polar contacts involving residues such as Lys92 and Arg110 and local hydrophobic interactions. (C) Representative conformation of the FAM171A2–F707-2364 complex, showing a hydrophobic interaction pattern involving residues such as Ala64, Val66, Val68 and Val73, together with local polar contacts. Representative structures were extracted from relatively stable regions of the molecular dynamics trajectories.

as representative candidates, yielding F707-2364 and E634-1386. Their ECFP4 Tanimoto similarity was 0.184, supporting their structural dissimilarity (Fig. S2). Centroid proximity was used to select representatives from the prefiltered pool; it does not establish that these compounds are the highest-ranked candidates under a comprehensive multi-criteria strategy.

© 2026 The Author(s). Published by the Royal Society of Chemistry

This clustering step served as a transition from large-scale screening to detailed simulation-based assessment. Earlier stages enriched compounds according to pharmacophore-aware similarity, predicted receptor compatibility and predicted BBBP, but those criteria alone could still favor structurally redundant analogs. The *K*-means step therefore provided a practical way to select a compact, structurally representative subset for molecular dynamics simulations. This procedure provided two structurally distinct representatives for detailed computational assessment.

3.6 Molecular dynamics simulations supported stable predicted complexes for E634-1386 and F707-2364

Molecular dynamics simulations were performed to further examine the predicted FAM171A2–ligand complexes obtained from the docking workflow. We first monitored complex RMSD to evaluate the overall structural deviation from the initial docked conformations. All three systems reached relatively stable RMSD ranges during the 100 ns simulations (Fig. 7A). The mean complex RMSD values were 3.61 ± 0.78 for Bemcentinib, 3.20 ± 0.65 for E634-1386 and 3.15 ± 0.71 for F707-2364. Thus, the two representative candidate complexes showed slightly lower overall deviations than the Bemcentinib complex in these trajectories.

Residue-level flexibility was then assessed using RMSF. The three systems displayed broadly similar RMSF profiles, with mean values of 1.01 ± 0.60 for Bemcentinib, 0.96 ± 0.61 for E634-1386 and 1.00 ± 0.64 for F707-2364 (Fig. 7B). The most flexible region was located around residues 300–500, consistent with the presence of more mobile loop-rich segments in this part of the predicted FAM171A2 model. Overall, the candidate ligands did not markedly reshape the global flexibility pattern of the protein, suggesting that the three complexes maintained comparable residue-level dynamic behavior.

We next analyzed the radius of gyration and protein solvent-accessible surface area to compare compactness and solvent exposure across the simulations. E634-1386 and F707-2364 showed lower mean R_g values than Bemcentinib, with values of 4.06 ± 0.66 and 4.53 ± 0.59 compared with 4.77 ± 0.40 for Bemcentinib (Fig. 7C). The mean protein SASA values were also lower for E634-1386 and F707-2364 than for Bemcentinib, with values of 35 104 ± 3421 Å^2 and 35 067 ± 3997 Å^2 compared with 36 442 ± 3081 Å^2 (Fig. 7D). These trends are consistent with slightly more compact and less solvent-exposed protein conformations in the two candidate-bound simulations. Together, the RMSD, RMSF, R_g and SASA analyses suggest that E634-1386 and F707-2364 formed stable predicted complexes with FAM171A2 under the simulation conditions used here.

The docking and MD-derived scoring results are summarized in Table 2. Table 2 presents the original calculations, whereas Table S1 summarizes docking scores across three receptor models and MM/GBSA estimates from three independent simulations per complex. Both candidates showed more favorable mean docking scores than Bemcentinib, supporting their prioritization within the tested models. These metrics should be interpreted as complementary computational prioritization indicators rather than as a single unified activity score. E634-1386 showed higher KarmaDock and CarsiDock scores than Bemcentinib, whereas F707-2364 showed the most favorable AutoDock Vina score and MM/GBSA estimate. Although the magnitude of these differences should be interpreted cautiously, particularly for MM/GBSA, both representative candidates showed more favorable values than Bemcentinib across the reported computational metrics. This pattern supports their continued prioritization for further experimental testing.

To inspect the predicted binding modes in more detail, representative low-energy conformations from relatively stable portions of the trajectories were selected for interaction analysis. The analyzed structures were extracted at 85.292 ns for Bemcentinib, 42.99 ns for E634-1386 and 18.63 ns for F707-2364. In the selected Bemcentinib conformation, polar contacts were observed with Asp530 and Tyr533, together with hydrophobic contacts involving Tyr92, Met537, Pro538 and Leu300. E634-1386 showed a distinct polar-contact pattern involving residues such as Trp92 and Arg110, while its hydrophobic substituent extended into a nearby nonpolar region. F707-2364 displayed a more hydrophobic interaction pattern involving residues including Ala68, Val78, Val66 and Ala68, together with local polar contacts near Tyr74 (Fig. 8).

These representative interaction patterns suggest that the two candidates may engage the predicted FAM171A2 binding region through different interaction modes: E634-1386 with a stronger contribution from polar contacts, and F707-2364 with broader hydrophobic complementarity. Collectively, the MD trajectories, MM/GBSA estimates and representative interaction analyses provide convergent computational support for prioritizing E634-1386 and F707-2364 as experimentally testable FAM171A2 ligand candidates.

Additional admetSAR3.0 profiling showed lower molecular weights and TPSA values, more favorable predicted aqueous solubility, and higher predicted BBB permeability scores for both candidates than for Bemcentinib (Table S2). The BBB scores were 0.833 and 0.921 for E634-1386 and F707-2364, respectively, compared with 0.332 for Bemcentinib. Both candidates were classified as acceptable under the platform's Lipinski criteria.

4 Conclusion

In this study, we developed a hierarchical AI-guided computational screening framework for prioritizing potential FAM171A2 ligands. The workflow integrated pharmacophore-aware retrieval, multi-stage receptor-based screening, Chemprop-based BBBP prediction, diversity-aware clustering and molecular dynamics-based assessment. Starting from a ChemDiv library of 1 413 202 compounds, this staged strategy reduced the search space to two structurally representative candidates, E634-1386 and F707-2364, while incorporating both receptor-level compatibility and predicted CNS-oriented permeability into the prioritization process.

Relative to the reference blocker Bemcentinib, E634-1386 and F707-2364 showed favorable *in silico* profiles across

 © 2026 The Author(s). Published by the Royal Society of Chemistry

complementary computational metrics. E634-1386 achieved higher KarmaDock and CarsiDock scores, whereas F707-2364 showed the most favorable AutoDock Vina score and MM/GBSA estimate. Molecular dynamics analyses suggested that both candidate complexes maintained slightly lower complex RMSD, R_g and protein SASA than the Bemcentinib complex, while showing broadly similar residue-level flexibility patterns. Representative interaction analysis further suggested that E634-1386 may rely more on polar contacts, whereas F707-2364 may engage the predicted binding region through broader hydrophobic complementarity.

Taken together, these results prioritize E634-1386 and F707-2364 as experimentally testable FAM171A2 ligand candidates rather than confirmed binders. More broadly, this study illustrates that combining pharmacophore-aware ligand retrieval, receptor-based structural filtering and predicted BBBP prioritization can provide a practical computational strategy for difficult CNS targets with limited ligand information and uncertain receptor structures. Future experimental studies will be required to validate direct FAM171A2 binding, inhibition of α-synuclein fibril uptake, cellular safety and blood–brain barrier-related properties.

Future FAM171A2 research may benefit from broader investigations of neuroprotective mechanisms, including studies of natural polysaccharides targeting protein aggregation and neuroinflammation in Alzheimer's disease.[37] Broader AI applications, exemplified by deep learning-based skin-lesion classification, also motivate continued refinement of biomedical feature learning and quantitative validation to strengthen computational candidate prioritization.[38]

Conflicts of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All necessary data supporting the findings of this study are provided in the article and the supplementary information (SI). Additional data and computational resources are available at https://github.com/linshaolong11/FAM171A2 and from the corresponding author upon reasonable request. Supplementary information is available. See DOI: https://doi.org/10.1039/d6ra07609b.

References

1 H. R. Morris, M. G. Spillantini, C. M. Sue and C. H. Williams-Gray, *Lancet*, 2024, **403**, 293–304.

2 I. Frigerio, M. L. Morella, A. B. Winkelaar, D. T. Vorgia, J. J. Brevé, W. A. Boiten, B. L. van der Gaag, I. L. Calandri, A. J. Jonker, H. W. Berendse, D. Mona, M. Britschgi, L. E. Jonkman and W. D. van de Berg, *Nat. Commun.*, 2025, **16**, 6630.

3 Y. Luo, L. Qiao, M. Li, X. Wen, W. Zhang and X. Li, *Front. Aging Neurosci.*, 2025, **16**, 1498756.

4 D. Su, Y. Cui, C. He, P. Yin, R. Bai, J. Zhu, J. S. T. Lam, J. Zhang, R. Yan, X. Zheng, J. Wu, D. Zhao, A. Wang, M. Zhou and T. Feng, *BMJ*, 2025, **388**, e080952.

5 E. Menozzi and A. H. Schapira, *Annu. Rev. Pharmacol. Toxicol.*, 2025, **65**, 237–258.

6 K.-M. Wu, Q.-H. Xu, Y.-Q. Liu, Y.-W. Feng, S.-D. Han, Y.-R. Zhang, S.-D. Chen, Y. Guo, B.-S. Wu and L.-Z. Ma, *Science*, 2025, **387**, 892–900.

7 K.-M. Wu, Y. Zhang, T.-T. Wang, J.-Y. Wu, Y.-Y. Huang, P. Yuan and J.-T. Yu, *Innovation*, 2026, **7**, 101112.

8 K.-M. Wu, Y. Zhang, T.-T. Wang, J.-Y. Wu, Y.-Y. Huang, P. Yuan and J.-T. Yu, *Expert Opin. Ther. Targets*, 2025, **29**, 667–670.

9 S. Yang, S. Li and J. Chang, *RSC Adv.*, 2022, **12**, 13500–13510.

10 S. Wang, Y. Zeng, H. Yang, S. Tian, Y. Zhou, L. Wang, X. Chen, H. Wang, Z. Gao and F. Bai, *Acta Pharmacol. Sin.*, 2025, **46**, 3108–3115.

11 S.-L. Lin, Y.-S. Chen, R.-Y. Liu, M.-Y. Zhu, T. Zhu, M.-Q. Wang and B.-Q. Liu, *RSC Adv.*, 2024, **14**, 8240–8250.

12 D. E. Graff, N. K. Morgan, J. W. Burns, A. C. Doner, B. Li, S.-C. Li, J. Manu, A. Menon, H.-W. Pang, H. Wu, A. S. Zalte, J. W. Zheng, C. W. Coley, W. H. Green and K. P. Greenman, *J. Chem. Inf. Model.*, 2026, **66**, 28–33.

13 L. Wang, Y. Wu, H. Luo, M. Liang, Y. Zhou, C. Chen, C. Liu, J. Zhang and Y. Zhang, *Adv. Sci.*, 2026, **13**, e13556.

14 J. Abramson, J. Adler, J. Dunger, R. Evans, T. Green, A. Pritzel, O. Ronneberger, L. Willmore, A. J. Ballard, J. Bambrick, S. W. Bodenstein, D. A. Evans, C.-C. Hung, M. O'Neill, D. Reiman, K. Tunyasuvunakool, Z. Wu, A. Žemgulytė, E. Arvaniti, C. Beattie, O. Bertolli, A. Bridgland, A. Cherepanov, M. Congreve, A. I. Cowen-Rivers, A. Cowie, M. Figurnov, F. B. Fuchs, H. Gladman, R. Jain, Y. A. Khan, C. M. R. Low, K. Perlin, A. Potapenko, P. Savy, S. Singh, A. Stecula, A. Thillaisundaram, C. Tong, S. Yakneen, E. D. Zhong, M. Zielinski, A. Žídek, V. Bapst, P. Kohli, M. Jaderberg, D. Hassabis and J. M. Jumper, *Nature*, 2024, **630**, 493–500.

15 X. Zhang, O. Zhang, C. Shen, W. Qu, S. Chen, H. Cao, Y. Kang, Z. Wang, E. Wang and J. Zhang, *Nat. Comput. Sci.*, 2023, **3**, 789–804.

16 H. Cai, C. Shen, T. Jian, X. Zhang, T. Chen, X. Han, Z. Yang, W. Dang, C.-Y. Hsieh and Y. Kang, *Chem. Sci.*, 2024, **15**, 1449–1471.

17 O. Trott and A. J. Olson, *J. Comput. Chem.*, 2010, **31**, 455–461.

18 X. Gong, J. Huang, C. Fu, S. Guo, Y. Tan, R. Lu, B. Liu, L. Zhao, S. Wang, H. H. Y. Tong, X. Yao, C. Chen and H. Liu, *J. Chem. Inf. Model.*, 2026, **66**, 6706–6718.

19 K. Schöning-Stierand, K. Diedrich, R. Fährrolfes, F. Flachsenberg, A. Meyder, E. Nittinger, R. Steinegger and M. Rarey, *Nucleic Acids Res.*, 2020, **48**, W48–W53.

20 S. Gu, C. Shen, X. Zhang, H. Sun, H. Cai, H. Luo, H. Zhao, B. Liu, H. Du, Y. Zhao, C. Fu, S. Zhai, Y. Deng, H. Liu, T. Hou and Y. Kang, *Nat. Mach. Intell.*, 2025, **7**, 509–520.

© 2026 The Author(s). Published by the Royal Society of Chemistry

21 S. Lin, S. Zhai, S. Wang, X. Zhan, Y. Li, W. Li, L. Qin, L. Shi, Y. Tian, K. Xu, K. Zhou, C. Gu, H. Liu and X. Yao, *Chem. Sci.*, 2026, **17**, 19964–19977.
22 G. Landrum, P. Tosco, R. Rodriguez, B. Kelley, D. Cosgrove, R. Vianello, S. Riniker, P. Gedeck, G. Jones, D. Nealschneider, E. Kawashima, N. Schneider, T. Hurst, A. Dalke, N. Maeder, M. Swain, Y. Pechersky, B. Cole, K. Boyd, S. Turk, A. Savelev, R. Walker, A. Vaucher, M. Wójcikowski, H. Faara, I. Take, V. F. Scalfani, S. Kearnes, K. Ujihara and D. Probst, *Zenodo*, 2026.
23 K. P. Sinaga and M.-S. Yang, *IEEE Access*, 2020, **8**, 80716–80727.
24 Y. Gu, Z. Yu, Y. Wang, L. Chen, C. Lou, C. Yang, W. Li, G. Liu and Y. Tang, *Nucleic Acids Res.*, 2024, **52**, W432–W438.
25 R. Salomon-Ferrer, D. A. Case and R. C. Walker, *Wiley Interdiscip. Rev.:Comput. Mol. Sci.*, 2013, **3**, 198–210.
26 L. Van der Maaten and G. Hinton, *J. Mach. Learn. Res.*, 2008, **9**, 2579–2605.
27 L. Fu, S. Shi, J. Yi, N. Wang, Y. He, Z. Wu, J. Peng, Y. Deng, W. Wang, C. Wu, A. Lyu, X. Zeng, W. Zhao, T. Hou and D. Cao, *Nucleic Acids Res.*, 2024, **52**, W422–W431.
28 G. Zhou, Z. Gao, Q. Ding, H. Zheng, H. Xu, Z. Wei, L. Zhang and G. Ke, in *The Eleventh International Conference on Learning Representations*, 2023.
29 Z. Xiong, D. Wang, X. Liu, F. Zhong, X. Wan, X. Li, Z. Li, X. Luo, K. Chen, H. Jiang and M. Zheng, *J. Med. Chem.*, 2020, **63**, 8749–8760.
30 J. Ross, B. Belgodere, V. Chenthamarakshan, I. Padhi, Y. Mroueh and P. Das, *Nat. Mach. Intell.*, 2022, **4**, 1256–1264.
31 H. Li, R. Zhang, Y. Min, D. Ma, D. Zhao and J. Zeng, *Nat. Commun.*, 2023, **14**, 7568.
32 Y. Wang, J. Wang, Z. Cao and A. B. Farimani, *Nat. Mach. Intell.*, 2022, **4**, 279–287.
33 J. Lyu, S. Wang, T. E. Balius, I. Singh, A. Levit, Y. S. Moroz, M. J. O'Meara, T. Che, E. Algaa and K. Tolmachova, *Nature*, 2019, **566**, 224–229.
34 J. A. Hartigan and M. A. Wong, *J. Roy. Stat. Soc. C Appl. Stat.*, 1979, **28**, 100–108.
35 C. Croux, P. Filzmoser and M. R. Oliveira, *Chemom. Intell. Lab. Syst.*, 2007, **87**, 218–225.
36 Y. Januzaj, E. Beqiri and A. Luma, *Int. J. Online Biomed. Eng.*, 2023, **19**, 174–182.
37 N. Muhammad, R. Ai, L. Yang, Y. Cui, L. Wang, S. Weng, Y. Wang, M. U. Kakar, S. Khan, S. Ahmad and I. Iqbal, *Front. Aging Neurosci.*, 2026, **18**, 1877358.
38 I. Iqbal, M. Younus, K. Walayat, M. U. Kakar and J. Ma, *CMIG*, 2021, **88**, 101843.

© 2026 The Author(s). Published by the Royal Society of Chemistry