Available Online: 27 November 2026 DOI: 10.54254/3029-0821/2026.39567

Exploratory docking of six flavonoids to an alphafold model of a human GPX4 sequence

Zihan Ni

School of Life Sciences, Nanjing Normal University, Nanjing, China

1539912192@qq.com

Abstract. Glutathione peroxidase 4 limits lipid peroxide accumulation and is a target of interest in ferroptosis research. Docking can help explore possible ligand poses, but its interpretation depends on the protein model and the method used. This study examined six flavonoids using an AlphaFold model associated with the unreviewed human GPX4 entry A0A9L9PXS3. A search region was defined from an experimental inhibitor complex, and each compound was docked with three random seeds. All eighteen runs completed successfully. Mean top-ranked Vina scores ranged from −5.495 to −5.108 kcal/mol; apigenin and luteolin had the most negative scores. For each compound, the maximum pairwise symmetry-corrected heavy-atom root-mean-square deviation among the three top-ranked poses was below 0.10 Å. Sequence checking nevertheless identified differences from the reviewed reference, including cysteine at the position corresponding to the catalytic selenocysteine. These findings show that the protocol returned consistent results for this model. They do not demonstrate binding, GPX4 inhibition, or ferroptosis induction. The receptor and docking method require further validation before the compounds can be prioritised on biological grounds.

Keywords: GPX4, ferroptosis, flavonoids, docking

1. Introduction

Ferroptosis is an iron-dependent form of regulated cell death associated with oxidative damage. Because it differs from apoptosis, it has attracted interest as a possible way to eliminate cancer cells susceptible to this process [1]. Glutathione Peroxidase 4 (GPX4) is an important part of that research. Experiments linking glutathione depletion and direct GPX4 inhibition to ferroptotic cell death established a basis for investigating the enzyme as a target [2].

Experimental GPX4 structures provide useful starting points for modelling. These include the covalent ML162 complex 6HKQ and the unliganded structure 6HN3 [3]. Their coordinates can help locate a search region, although differences in protein constructs and ligand chemistry need to be considered. A covalent complex alone cannot validate a conventional noncovalent docking protocol.

AlphaFold predictions, available through the AlphaFold Protein Structure Database, offer another source of structural models [4, 5]. Each prediction represents a particular amino acid sequence, however. Entries associated with the same gene may differ at residues relevant to the question being studied, so the sequence needs to be checked before docking results are assigned to a biological target.

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

This study examined six related flavonoids in one GPX4 sequence model, comparing docking scores, pose consistency, nearby residues, and calculated molecular properties. It also examined how the choice of an unreviewed sequence and a reference region derived from a covalent complex affected interpretation. The aim was to establish a reproducible pilot workflow and identify the evidence still needed for biological evaluation. The compounds were not assumed to have experimentally established GPX4 inhibitory activity. This work also has practical and methodological value. For researchers studying GPX4 and other ferroptosis-related targets, it provides a fully documented, reproducible pilot workflow that combines sequence verification, a reference region transferred from an experimental covalent complex, repeated docking across random seeds, and explicit assessment of pose consistency; the reported checks, especially the distinction between the unreviewed model and the selenocysteine-containing reviewed reference and the limits of a search region derived from a covalent inhibitor, offer a transferable example of how computational screening should be reported and where overinterpretation can arise. For natural-product and drug-discovery programmes, the six flavonoids supply an inexpensive and traceable set of structural hypotheses that can help direct later experimental resources toward better-justified candidates rather than unguided screening. The study makes no claim of therapeutic activity; rather, it defines the validation steps needed before such claims can be made, including enzyme inhibition, direct binding or target-engagement evidence, and cellular assays that distinguish ferroptosis from other forms of cell death, thereby supporting more rigorous and cost-effective early-stage decisions and providing a clear baseline for future work.

2. Materials and methods

2.1. Receptor identity and preparation

The structure AF-A0A9L9PXS3-F1-model_v6.pdb was obtained from the AlphaFold database. Its sequence was compared with the corresponding UniProt entry and the reviewed human GPX4 reference P36969 [6]. The model contained 195 residues, compared with 197 in the reviewed reference. Reference residues Tyr59 and Arg60 were absent, and model residue 71 was cysteine rather than the selenocysteine at reference position 73. These are observed sequence differences; their biological origin and the experimental validity of the unreviewed entry were not established in this study.

Per-residue predicted Local Distance Difference Test scores (pLDDT) were read from the confidence field in the Protein Data Bank (PDB) format file, using one alpha-carbon (Cα) record per residue. The complete model was retained for docking. Meeko 0.7.1 generated a rigid receptor PDBQT file and a parameterised JavaScript Object Notation (JSON) record. No receptor energy minimisation or molecular dynamics simulation was performed. The prepared receptor contained all 195 residues, including Cys71. Its file hash was recorded to identify the receptor used throughout the batch.

2.2. Reference region and compound selection

The experimental 6HKQ structure was used to locate a reference region rather than to supply a validated noncovalent binding pose. Matched Cα coordinates were fitted to the predicted model by a least-squares rigid transformation. The C66S construct mismatch in the experimental structure was excluded from the matched set. The resulting transformation was applied to the ML162 ligand coordinates using alternate location A where alternatives were present. The search box centre was 4.877, 17.340, and 3.442 Å along the three coordinate axes, with an edge length of 20 Å. No independent cavity prediction algorithm was applied.

Apigenin, luteolin, kaempferol, quercetin, myricetin, and isorhamnetin were selected before their docking scores were examined. Their shared flavonoid framework and different hydroxyl or methoxy substituents

provided a small set for structural comparison. This was a practical selection rather than a representative chemical library. Structures were obtained from PubChem, and the compound identifiers are given in Table 1 [7]. No member of the panel was treated as a validated GPX4 positive or negative control.

2.3. Ligand preparation and docking

RDKit 2026.09.1 was used to read the supplied structures, check molecular formulae and InChIKeys, and add hydrogens[1]. Ten conformers were generated per compound with experimental torsion knowledge distance geometry version 3 (ETKDGv3) and random seed 20261030. Merck Molecular Force Field (MMFF) optimisation allowed up to 2000 iterations; the lowest-energy converged conformer was selected. This energy was used only for conformer selection within each compound. Different molecular formulae were not compared by their MMFF energies. Meeko generated the ligand PDBQT files [8]. The supplied neutral structures were retained without enumeration of alternative protonation states or tautomers.

AutoDock Vina 1.2.7 was run with the Vina scoring function, exhaustiveness 32, eight CPU threads, and up to nine requested modes [9, 10]. Three separate runs used random seeds 20261030, 20261031, and 20261032 for each ligand, giving eighteen calculations. The receptor remained rigid. Logs, commands, ligand coordinates, software versions, and output hashes were retained. An earlier caffeine calculation served only as a technical workflow check and was excluded from the six-compound ranking.

2.4. Output assessment and visualisation

Analysis used the score and coordinates of the first-ranked pose from each run. Means and ranges summarised variation across the three random seeds; the repeats were not treated as biological observations, and no inferential statistical test was applied. Pose comparisons used the fixed receptor coordinate frame. Heavy atoms were mapped from the SMILES index records in the PDBQT output, with equivalent graph mappings considered to account for molecular symmetry. The ligand poses were not superimposed before comparison.

Nearby residues were defined by a protein–ligand heavy-atom separation of at most 4.0 Å. Minimum heavy-atom distances were also checked for obvious overlaps. These geometric criteria did not classify hydrogen bonds. RDKit calculated molecular weight, the octanol–water partition descriptor logP, hydrogen-bond donor and acceptor counts, and topological polar surface area. The calculations were descriptive and were not treated as measurements of absorption or toxicity. ChimeraX was used to inspect a representative apigenin pose and produce Figure 1 [11]. The exact ChimeraX version was not recorded. Low-confidence N-terminal residues were hidden in selected displays but remained present in the docking receptor.

3. Results

3.1. Receptor confidence and sequence limitations

The mean confidence score across the model was 88.57. Confidence was uneven: residues 1–30 had a mean score of 43.98, while residues 31–195 had a mean of 96.68. The overall average therefore gave limited information about the N-terminal region. Fitting the model to 6HKQ gave a root-mean-square deviation of approximately 0.534 Å across 162 Cα pairs. The backbone agreement allowed a reference region to be

[1] RDKit contributors. (2026). The RDKit book and getting started with the RDKit in Python. RDKit 2026.09.1 documentation. https: //www.rdkit.org/docs/RDKit_Book.html and https: //www.rdkit.org/docs/GettingStartedInPython.html

transferred between structures, but did not establish that their side-chain chemistry or ligand recognition was equivalent.

Sequence comparison identified a limitation that the confidence scores could not address. In addition to the two missing reference residues, the model had cysteine at the position corresponding to the reference selenocysteine. The results were therefore assigned specifically to A0A9L9PXS3. A close backbone fit and a GPX4 annotation were insufficient grounds for treating this model as the reviewed wild-type target. Unless stated otherwise, all residue numbers below refer to the 195-residue model.

3.2. Docking scores and search consistency

All eighteen docking runs completed successfully. Scores in the output coordinate files matched the numerical summary, and output file hashes matched the run records. Table 1 reports the three scores for each compound. Apigenin had the most negative mean score at −5.495 kcal/mol, followed by luteolin at −5.446 kcal/mol. Mean scores for quercetin, kaempferol, myricetin, and isorhamnetin were −5.292, −5.255, −5.190, and −5.108 kcal/mol, respectively.

Score ranges within compounds were 0.002–0.019 kcal/mol. Apigenin and luteolin differed by only 0.049 kcal/mol in their mean scores; the difference between the most and least negative means was 0.387 kcal/mol. This ordering describes the calculations and was not interpreted as a ranking of experimental potency. Coordinate comparisons also showed little variation: the maximum pairwise symmetry-corrected heavy-atom deviation among the three first-ranked poses ranged from 0.023 to 0.094 Å. This analysis was limited to the top-ranked poses and did not assess convergence of all returned modes.

Table 1. Docking scores and pose consistency across three random seeds

Compound	PubChem CID	Seed 1030	Seed 1031	Seed 1032	Mean score	Max RMSD
Apigenin	5280443	-5.501	-5.491	-5.492	-5.495	0.063
Luteolin	5280445	-5.437	-5.450	-5.450	-5.446	0.024
Quercetin	5280343	-5.291	-5.293	-5.292	-5.292	0.023
Kaempferol	5280863	-5.262	-5.258	-5.245	-5.255	0.050
Myricetin	5281672	-5.195	-5.190	-5.185	-5.190	0.089
Isorhamnetin	5281654	-5.105	-5.100	-5.119	-5.108	0.094

Scores are in kcal/mol. Seed labels abbreviate 20261030, 20261031, and 20261032. Max RMSD is the maximum pairwise symmetry-corrected heavy-atom RMSD in Å, calculated without ligand superposition. All repeats are technical.

3.3. Local geometry and molecular properties

Gly72, Gln106, Trp161, and Met181 lay within 4.0 Å of ligand heavy atoms in all eighteen first-ranked poses. Lys73 and Lys160 also fell within this distance for each apigenin pose. Figure 1 shows the apigenin pose from seed 20261030. The displayed residues are spatial neighbours; the analysis did not establish a hydrogen-bond network or a covalent linkage. Labels use model residue numbers, which should not be read as crystallographic residue numbers.

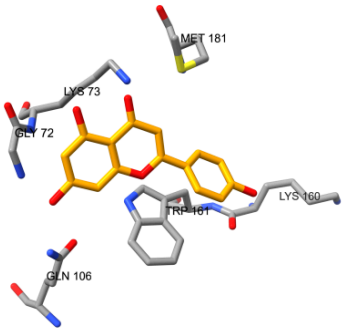

Figure 1. Predicted apigenin pose in the A0A9L9PXS3 model

Ligand carbon atoms are orange; protein carbon atoms are grey; oxygen, nitrogen, and sulfur are red, blue, and yellow. The displayed residues have a heavy atom within 4.0 Å of the ligand. Model numbering is used. The pose is from seed 20261030 with a Vina score of −5.501 kcal/mol. The figure depicts proximity, not validated hydrogen bonds or inhibition.

The shortest protein–ligand heavy-atom distances across the eighteen first-ranked poses were approximately 2.85–3.10 Å. None was below 2.0 Å. The shortest distance between a ligand heavy atom and the sulfur atom of Cys71 was 4.18–4.46 Å. This check found no obvious heavy-atom overlap under the stated criterion, although it did not establish the physical accuracy of the predicted poses.

Table 2 summarises the calculated molecular properties. Molecular weights were 270.24–318.24 g/mol and calculated logP values were 1.69–2.58. Myricetin had the highest donor count and polar surface area, with six donors and 151.59 Å². Isorhamnetin, although among the heavier molecules, had four donors. These values describe structural differences within the panel. Absorption, distribution, metabolism, excretion, and toxicity were neither measured nor independently predicted.

Table 2. Calculated physicochemical descriptors of the six flavonoids

Compound	MW g/mol	logP	HBD	HBA	TPSA Å²
Apigenin	270.24	2.58	3	5	90.90
Luteolin	286.24	2.28	4	6	111.13
Quercetin	302.24	1.99	5	7	131.36
Kaempferol	286.24	2.28	4	6	111.13
Myricetin	318.24	1.69	6	8	151.59
Isorhamnetin	316.26	2.29	4	7	120.36

Note. MW, molecular weight; HBD, hydrogen-bond donors; HBA, hydrogen-bond acceptors; TPSA, topological polar surface area. Descriptors were calculated with RDKit and do not establish pharmacokinetic behaviour.

4. Discussion

4.1. Meaning of the observed consistency

The three searches for each compound returned similar top-ranked poses as well as similar scores. Examining both was useful: agreement between scores alone would not show whether the ligand occupied the same

region. Accounting for symmetry also avoided interpreting exchanges between chemically equivalent atoms as differences in the pose. In this dataset, the coordinate comparisons supported consistency across random seeds for each compound under the chosen preparation, search box, and scoring settings.

That consistency does not establish predictive accuracy. Changing the random seed changes the search, while the receptor, scoring function, ligand state, and box remain the same. An error in any of those choices can therefore affect all three runs. The narrow score ranges should not be read as confidence intervals for experimental affinity. The seeds also do not provide independent biological observations from which to claim statistical significance. What the repeats demonstrate is limited variation within this calculation; independent reference data would be needed to test whether the method predicts experimentally supported poses or activity.

4.2. Receptor choice and reference chemistry

Receptor identity is the main limitation. A confident prediction of the fold cannot resolve a difference between the model sequence and the intended target. Here, the model contains cysteine at the position corresponding to the reference selenocysteine. This matters particularly because the reference ligand was obtained from a covalent inhibitor complex. The findings must consequently be interpreted for the actual sequence and chemical representation used. They cannot be transferred to the reviewed human GPX4 reference solely on the basis of a shared annotation or similar backbone coordinates.

The 6HKQ structure helped define the search region, but its ML162 ligand is covalently bound [3]. The current protocol did not model bond formation, a reaction pathway, or the associated energetics. It therefore did not test whether the flavonoids reproduce that inhibitory mechanism. Nor does the Cα fit establish that side chains adopt equivalent local arrangements. Before screening is extended, the reference sequence, construct, residue mapping, and treatment of selenium need to be settled. Using an experimental structure may help, provided that its mutations and ligand-bound state are also considered in receptor preparation and interpretation.

4.3. Compound comparison and biological interpretation

The six compounds form a small structural comparison set. Without a validated activity assay or a benchmark ligand set, their scores cannot establish which compound should bind more strongly. The difference between apigenin and luteolin is particularly small, despite their positions at the top of the numerical ranking. Selectivity was not tested either: the workflow did not compare these compounds across other protein targets. The ordering is useful as a record of this model's output, but further evidence is required before it can guide biological prioritisation.

The molecular descriptors provide a different comparison. Within this panel, structures with additional hydroxyl groups had higher donor counts and polar surface areas, while methoxy substitution changed that balance. These differences do not establish whether a compound can reach GPX4 in a cell. Solubility, stability, exposure, and target engagement could all affect the relationship between a docking result and an observed cellular response. A cellular effect, if one were later detected, could also involve another target or pathway. The present data therefore provide no evidence that these flavonoids induce ferroptosis or have anticancer activity.

4.4. Validation priorities and further work

Further work should start by resolving the receptor representation. The reviewed sequence and candidate experimental constructs need to be compared, and the relevant residue chemistry must be supported by the preparation and docking tools. Validation should then use experimental ligand evidence appropriate to that

method. A covalent reference may require a protocol designed for covalent interactions; a conventional noncovalent redocking score would not test the same process. Where suitable noncovalent reference data exist, pose recovery and discrimination between active and appropriately chosen inactive or decoy compounds could be assessed before screening additional molecules.

Once these choices are justified, calculations could test sensitivity to receptor conformation, protonation, and alternative ligand states. AlphaFold 3 offers an additional approach to predicting protein–small molecule complexes, but another structural prediction would still need independent validation [12]. It was not used in this study, which was limited to a database receptor and a conventional docking workflow. A meaningful comparison would require consistent sequence definitions and chemical states, so differences between methods could be separated from differences in their inputs.

A claim of GPX4 inhibition would ultimately require an experimental test of enzyme activity, supported where appropriate by direct binding or target-engagement evidence. Cellular studies would need further controls to distinguish ferroptosis from other causes of cell death. None of these experiments were performed here. They are necessary follow-up questions rather than conclusions from the docking scores. The value of the current dataset is as a documented pilot: it establishes what was calculated, shows the limited variation between searches, and identifies assumptions that need attention before the project can support stronger biological claims.

5. Conclusion

Six flavonoids were docked to an AlphaFold model associated with the human GPX4 entry A0A9L9PXS3. All eighteen runs completed successfully, and the first-ranked poses changed little across three random seeds per compound. Apigenin and luteolin had the most negative mean scores, separated by a small numerical difference. Several nearby residues were common to all six compounds, and the calculated molecular properties described variation within the panel.

These results show that the selected workflow returned consistent outputs for one receptor model. They do not establish binding affinity, GPX4 inhibition, selectivity, or ferroptosis induction. The model differs from the reviewed reference sequence at the position corresponding to the catalytic selenocysteine and at two additional reference residues. The search region was also derived from a covalent inhibitor complex, although the docking calculation did not model covalent chemistry. These factors limit the biological interpretation of the score ordering.

The next stage should address receptor identity and validate the method against suitable independent evidence. Testing justified molecular states and, where possible, obtaining experimental evidence would then allow a more meaningful assessment of the compounds. Until those steps are completed, the findings remain preliminary and specific to the model used.

References

[1] Dixon, S. J., Lemberg, K. M., Lamprecht, M. R., Skouta, R., Zaitsev, E. M., Gleason, C. E., Patel, D. N., Bauer, A. J., Cantley, A. M., Yang, W. S., Morrison, B., III, & Stockwell, B. R. (2012). Ferroptosis: An iron-dependent form of nonapoptotic cell death. *Cell, 149*(5), 1060–1072. https://doi.org/10.1016/j.cell.2012.03.042

[2] Yang, W. S., SriRamaratnam, R., Welsch, M. E., Shimada, K., Skouta, R., Viswanathan, V. S., Cheah, J. H., Clemons, P. A., Shamji, A. F., Clish, C. B., Brown, L. M., Girotti, A. W., Cornish, V. W., Schreiber, S. L., &

Stockwell, B. R. (2014). Regulation of ferroptotic cancer cell death by GPX4. *Cell, 156*(1–2), 317–331. https: //doi.org/10.1016/j.cell.2013.12.010

[3] Moosmayer, D., Hilpmann, A., Hoffmann, J., Schnirch, L., Zimmermann, K., Badock, V., Furst, L., Eaton, J. K., Viswanathan, V. S., Schreiber, S. L., Gradl, S., & Hillig, R. C. (2021). Crystal structures of the selenoprotein glutathione peroxidase 4 in its apo form and in complex with the covalently bound inhibitor ML162. *Acta Crystallographica Section D, 77*(2), 237–248. https: //doi.org/10.1107/S2059798320016125

[4] Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., ... Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature, 596*(7873), 583–589. https: //doi.org/10.1038/s41586-021-03819-2

[5] Varadi, M., Bertoni, D., Magana, P., Paramval, U., Pidruchna, I., Radhakrishnan, M., Tsenkov, M., Nair, S., Mirdita, M., Yeo, J., Kovalevskiy, O., Tunyasuvunakool, K., Laydon, A., Žídek, A., Tomlinson, H., Hariharan, D., Abrahamson, J., Green, T., Jumper, J., ... Velankar, S. (2024). AlphaFold Protein Structure Database in 2024: Providing structure coverage for over 214 million protein sequences. *Nucleic Acids Research, 52*(D1), D368–D375. https: //doi.org/10.1093/nar/gkad1011

[6] The UniProt Consortium. (2025). UniProt: The universal protein knowledgebase in 2025. *Nucleic Acids Research, 53*(D1), D609–D617. https: //doi.org/10.1093/nar/gkae1010

[7] Kim, S., Chen, J., Cheng, T., Gindulyte, A., He, J., He, S., Li, Q., Shoemaker, B. A., Thiessen, P. A., Yu, B., Zaslavsky, L., Zhang, J., & Bolton, E. E. (2023). PubChem 2023 update. *Nucleic Acids Research, 51*(D1), D1373–D1380. https: //doi.org/10.1093/nar/gkac956

[8] Santos-Martins, D., He, Y., Eberhardt, J., Sharma, P., Bruciaferri, N., Holcomb, M., Llanos, M. A., Hansel-Harris, A., Barkdull, A. P., Tillack, A. F., Bianco, G., Paulsen, M.-L., Mato, J., Taneja, I., & Forli, S. (2025). Meeko: Molecule parametrization and software interoperability for docking and beyond. *Journal of Chemical Information and Modeling, 65*(24), 13045–13050. https: //doi.org/10.1021/acs.jcim.5c02271

[9] Trott, O., & Olson, A. J. (2010). AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry, 31*(2), 455–461. https: //doi.org/10.1002/jcc.21334

[10] Eberhardt, J., Santos-Martins, D., Tillack, A. F., & Forli, S. (2021). AutoDock Vina 1.2.0: New docking methods, expanded force field, and Python bindings. *Journal of Chemical Information and Modeling, 61*(8), 3891–3898. https: //doi.org/10.1021/acs.jcim.1c00203

[11] Pettersen, E. F., Goddard, T. D., Huang, C. C., Meng, E. C., Couch, G. S., Croll, T. I., Morris, J. H., & Ferrin, T. E. (2021). UCSF ChimeraX: Structure visualization for researchers, educators, and developers. *Protein Science, 30*(1), 70–82. https: //doi.org/10.1002/pro.3943

[12] Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., Bodenstein, S. W., Evans, D. A., Hung, C.-C., O'Neill, M., Reiman, D., Tunyasuvunakool, K., Wu, Z., Žemgulytė, A., Arvaniti, E., ... Jumper, J. M. (2024). Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature, 630*(8016), 493–500. https: //doi.org/10.1038/s41586-024-07487-w