

Assembly of Protein Complexes in and on the Membrane with Predicted Spatial Arrangement Constraints

Charles Christoffer ¹, Kannan Harini ^{2,3}, Gupta Archit ^{3,4} and Daisuke Kihara ^{1,3,5*}

1 - Department of Computer Science, Purdue University, West Lafayette, IN 47907, USA

2 - Department of Biotechnology, Bhupat and Jyoti Mehta School of Biosciences, Indian Institute of Technology Madras, Chennai 600036, India

3 - Department of Biological Sciences, Purdue University, West Lafayette, IN 47907, USA

4 - Department of Genetic Engineering, SRM Institute of Science and Technology, Kattankulathur 603203, India

5 - Purdue University Center for Cancer Research, Purdue University, West Lafayette, IN 47907, USA

Correspondence to Daisuke Kihara: dkihara@purdue.edu (D. Kihara)

<https://doi.org/10.1016/j.jmb.2024.168486>

Edited by Michael Sternberg

Abstract

Membrane proteins play crucial roles in various cellular processes, and their interactions with other proteins in and on the membrane are essential for their proper functioning. While an increasing number of structures of more membrane proteins are being determined, the available structure data is still sparse. To gain insights into the mechanisms of membrane protein complexes, computational docking methods are necessary due to the challenge of experimental determination. Here, we introduce Mem-LZerD, a rigid-body membrane docking algorithm designed to take advantage of modern membrane modeling and protein docking techniques to facilitate the docking of membrane protein complexes. Mem-LZerD is based on the LZerD protein docking algorithm, which has been constantly among the top servers in many rounds of CAPRI protein docking assessment. By employing a combination of geometric hashing, newly constrained by the predicted membrane height and tilt angle, and model scoring accounting for the energy of membrane insertion, we demonstrate the capability of Mem-LZerD to model diverse membrane protein–protein complexes. Mem-LZerD successfully performed unbound docking on 13 of 21 (61.9%) transmembrane complexes in an established benchmark, more than shown by previous approaches. It was additionally tested on new datasets of 44 transmembrane complexes and 92 peripheral membrane protein complexes, of which it successfully modeled 35 (79.5%) and 15 (16.3%) complexes respectively. When non-blind orientations of peripheral targets were included, the number of successes increased to 54 (58.7%). We further demonstrate that Mem-LZerD produces complex models which are suitable for molecular dynamics simulation. Mem-LZerD is made available at <https://lzerd.kiharalab.org>.

© 2024 Elsevier Ltd. All rights reserved.

Introduction

Protein-protein interactions are fundamental to many biological processes in living cells. Membrane proteins play an essential role in many of these processes, where they act as gateways of cellular signaling pathways, pumps, and more,

facilitating selective transport processes across membranes. To understand the detailed mechanisms of these processes, modeling the 3D structures of their associated protein complexes is a critical step. While protein complex structures are steadily being determined by experiment and deposited in the Protein Data Bank (PDB),^{1,2}

experiments are still expensive and a substantial human and instrument time investment, assuming the proteins at hand are compatible with the methods available.³ Moreover, structures of protein complexes are often extremely difficult to determine by experiments, even without the involvement of a membrane. Thus, when a protein complex structure has not yet been experimentally determined, computational tools can be used to construct atomic models.⁴ A protein docking program can take component proteins, called subunits, as input and assemble them into models of the protein complex. Many general protein docking methods and specialized versions thereof have been publicly released, such as ZDOCK,⁵ HADDOCK,⁶ ClusPro,⁷ RosettaDock,⁸ HEX,⁹ SwarmDock,¹⁰ and ATTRACT.¹¹ Even protein structure prediction methods like AlphaFold¹² have been tweaked to be able to output multimeric structures.¹³ The rigid-body docking method LZerD^{14–17} in particular has been consistently ranked highly in the server category in CAPRI,^{18,19} the blind communitywide assessment of protein docking methods.

The fact that proteins *in vivo* are not generally interacting in isolation in a uniform environment is often confounding to computational modeling of complexes. Even with state-of-the-art modeling techniques, existing docking methods struggle to rigorously handle environments other than a uniform aqueous environment.^{20,21} Membranes create an environment where hydrophobic surfaces are not as energetically incentivized to be buried in the protein–protein interface, since such surfaces may competitively interact with the membrane itself. For example, in the transmembrane halorhodopsin protein family as found in halobacteria, which is included in the benchmark of Mem-LZerD, the protein–protein interface is not especially rich in hydrophobic amino acids relative to the remainder of the molecular surface, much of which participates in the protein–lipid interface.²² Highlighting the utility of docking methods, a halorhodopsin structure for *Halobacterium salinarium* was available in 2000, but no structure for *Natronomonas pharaonic* was available until 2009.²³ Homology-based methods can be used to model the subunits, and a docking method can then be used to explore the space of interaction poses. Other proteins do not wholly embed in the membrane, but instead pass only part-way through or interact only with the membrane surface. These peripheral membrane proteins are likewise important and are implicated for example in sensitivity to membrane composition.²⁴ These broad categories of membrane protein break down further into classes with substantial mechanistic differences, from the purely α -helical transmembrane regions most commonly considered by computational methods, transmembrane β -barrels, to peripheral membrane proteins attaching to the membrane with amphipathic helices²⁵ or with hydrophobic loops.²⁶ Transience in certain interac-

tions between these proteins can render them difficult to directly consider *in vitro*,²⁷ but more accessible via computational modeling.²⁸ It is then clear that techniques capable of modeling interactions involving membranes have the potential to elucidate many cellular processes in many biological contexts. Detailed mechanistic understanding of membrane protein complexes currently represents a major knowledge gap in molecular biology and is the subject of much active investigation.²⁹ Recent studies have shown that protein structure prediction methods can enable combinatorial modeling of putative interactions among, for example, the cytochrome c maturation system I proteins of *E. coli*.²⁸

Several computational techniques have been developed which predict the modes of interaction of proteins in a membrane environment. In the case of G-protein coupled receptors (GPCRs), highly specialized approaches have been developed and applied to model their oligomerization.^{30–33} For more general proteins, MPDock, part of the RosettaMP software collection, takes a specification of membrane chemistry and dimensions to model the assembly of transmembrane complexes from bound experimental structures.³⁴ In their benchmark, the bound subunits were pulled apart, repacked according to the algorithm of Rosetta, and then pulled back together. Memdock carries out rigid-body docking of α -helical membrane proteins by constraining their orientation, before a finer-grained refinement procedure, and makes the assumption that the whole-input center of mass roughly coincides with the membrane midplane.^{35,36} JabberDock for membrane proteins represents the subunits with a volume map incorporating dynamics information from an expensive 70 nanosecond molecular dynamics simulation and uses particle swarm optimization techniques to explore the pose space, which required strict constraints on the pose space.³⁷ In summary, MPDock requires bound native structures as input and requires manual processing, Memdock was only tested on α -helical proteins and requires assumptions about the membrane location, and is further unable to include soluble regions of proteins in the modeling, and JabberDock requires substantial molecular dynamics calculations and is limited to exploring a tight region of the search space the precludes tolerance for misorientation. These existing analyses require orientations taken from bound complex structures, rather than orientations predicted separately for each subunit. In the context of blind modeling, however, the precise orientations which individual subunits take on upon binding are not known. A method for blind membrane docking should predict the membrane orientations of the input subunits and have a docking search space which is narrow enough to exclude prediction errors, but broad enough to tolerate reasonable errors in orientation predictions.

Mem-LZerD, which we developed in this work, is based on the LZerD rigid-body docking method and its extensions.^{14–17,38–40} LZerD uses a soft surface representation of the protein subunits based on geometric hashing⁴¹ and 3D Zernike descriptors (3DZDs),^{42–47} which allows for fast generation of docking poses without considering side-chain repacking. The geometric hashing procedure originally used internally by regular LZerD also admits more site-specific orientation and translation restrictions. Mem-LZerD targets any transmembrane protein complexes, as well as peripheral membrane proteins. The geometric hashing data structure used by Mem-LZerD is newly augmented with the positioning of each sample point relative to the membrane generated by the Positioning of Proteins in Membranes (PPM) algorithm,^{48,49} as well as the angular orientation in the same membrane. Pruning the search space in this way, skipping infeasible poses, yielded a running time 74 times faster on average compared to regular LZerD without search space or model constraints. This speedup highlights that for membrane proteins, most of the calculation time in regular LZerD is spent on infeasible poses, e.g. those which are sideways or entirely outside the membrane, which then contaminates the output model ranking.

Mem-LZerD yielded acceptable models in the CAPRI criteria within the top 10 models for 13 of 21 (61.9%) unbound docking targets of the Memdock benchmark set, which is a greater fraction than successfully modeled by existing methods Memdock or JabberDock. Previous studies used knowledge of the ground truth (i.e. correct) subunit orientations, taken as those found in the OPM database, rather than predicting them as in the Mem-LZerD protocol. When assuming the ground truth orientations, Mem-LZerD successfully modeled 14 of 21 targets (66.7%). On our separate transmembrane protein benchmark set, Mem-LZerD successfully modeled 35 of the 44 (79.5%) unbound benchmark docking targets, while on a preoriented peripheral membrane protein benchmark set, Mem-LZerD successfully modeled 54 of 92 (58.7%). Mem-LZerD has been incorporated into the LZerD webserver, available at <https://lzerd.kiharalab.org>. We further show that the protocol pipeline of the Orientations of Proteins in Membranes/Positioning of Proteins in Membranes (OPM/PPM) suite,⁴⁸ Mem-LZerD, and CHARMM-GUI⁵⁰ can produce ready-for-simulation files of sampled binding poses in explicit lipid membranes.

Results and discussion

Overview of Mem-LZerD

The docking procedure of Mem-LZerD follows the observation that presence of a lipid membrane restricts the poses two proteins may take on when they assemble into a complex. Mem-LZerD builds

assembled complex models by independently orienting input proteins using PPM,⁴⁸ then applying a geometric hashing-based procedure adapted from LZerD with augmentations that allow pruning the search space to sample membrane-bound poses, before a final clustering and scoring procedure yield the final outputs. The overall procedure of Mem-LZerD is diagrammed in Figure 1a.

The first step is to orient both the receptor and ligand protein subunits separately in the membrane using PPM, which minimizes its computed free energy of transfer between the aqueous and membrane environments. PPM has been tested on thousands of membrane proteins and validated against hundreds of experimental results.⁴⁸ PPM, however, does not in general generate precisely the orientations the subunits take on in their bound states, but the search and sampling procedure of Mem-LZerD is sufficiently tolerant that such precision is not necessary. Transmembrane proteins are immersed in the membrane and therefore more clearly restrained by it, while peripheral membrane proteins merely interact with the membrane surface and are less clearly restrained.

In the next stage, Mem-LZerD searches the pose space using the LZerD algorithm while accounting for the restricted pose space using an augmented data structure. Unlike the data structures used in original LZerD, the newly augmented geometric hashing in Mem-LZerD facilitates restricting the search space without the expense of running the full pipeline and post-filtering the output. Docking poses were restricted in terms of two parameters, the difference of the heights of two subunit sample points relative to the membrane and the tilt angles between the two subunits as illustrated in Figure 1b. These simple criteria facilitate efficient search of the pose space without wasting time calculating intermediate scores for infeasible poses.

In Figure 2a and b, the distributions of the height difference and the tilt angle of docked models for cytochrome c oxidase subunit 1 and subunit 3 (PDB 1M56 chains A and C) are shown. In Figure 2a, docking was performed with predicted orientations, while correct orientations of the subunits were used in Figure 2b. Both panels clearly illustrate that models of at least acceptable CAPRI quality tend to be found closer to the origin, i.e. docking conformations with a small height difference and tilt angle difference, regardless of whether the basis of the constraints used is a blind or non-blind orientation. In Figure 2c, to select optimal cutoff values for the height and tilt angle differences, an exhaustive parameter sweep was carried out for height difference cutoffs of 4 Å, 8 Å, and 12 Å and tilt difference cutoffs of 0.2 radians, 0.4 radians, and 0.8 radians on the training set of 20 nonredundant transmembrane protein complexes taken from the OPM database.⁴⁹ For this sweep, predicted subunit orientations were used. For each combination of

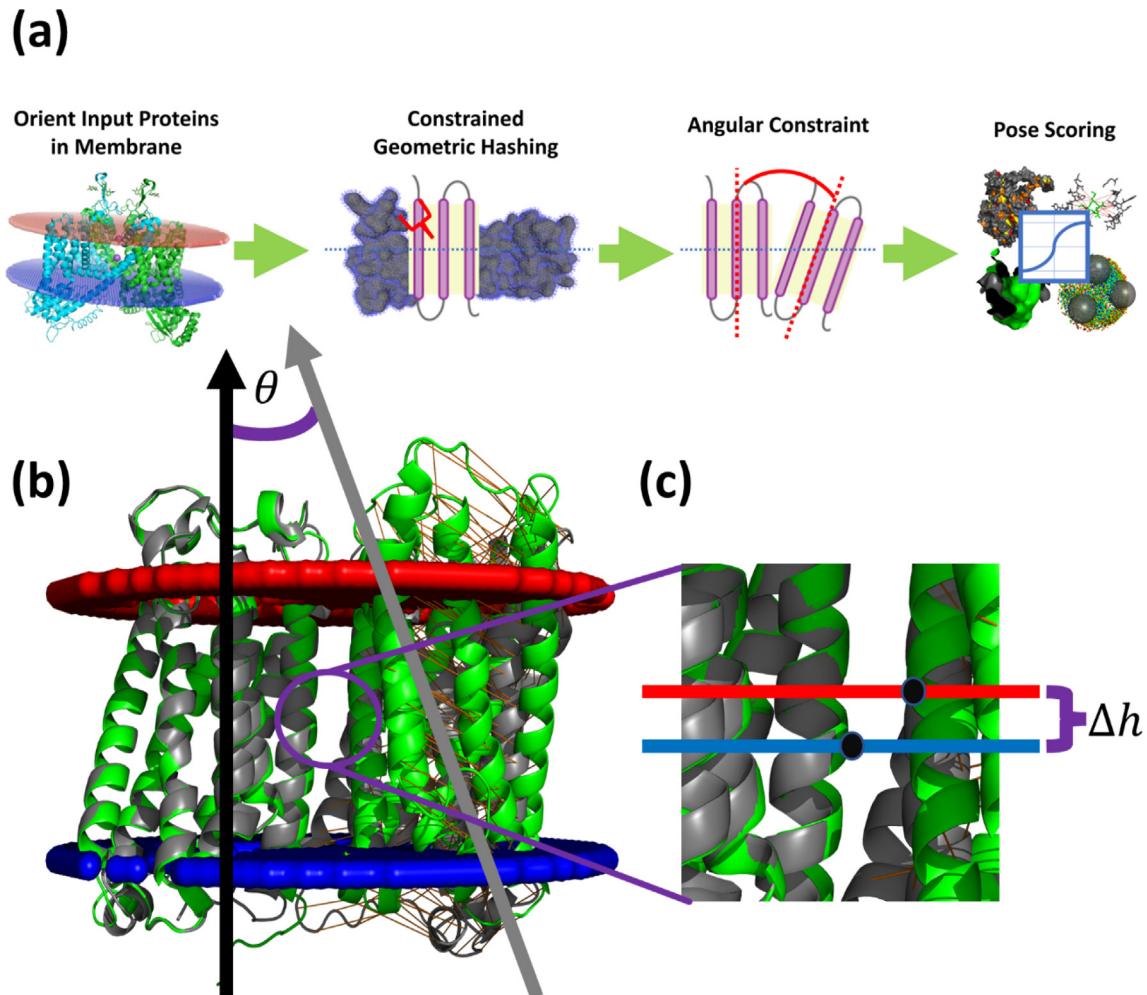


Figure 1. Illustration of Mem-LZerD. **(a)** Overall flow of the Mem-LZerD method. **(b)** Tilt constraint used by Mem-LZerD. The black and grey arrows represent the axes of the receptor and the ligand subunits that are perpendicular to the membrane planes as oriented by PPM, while θ is the tilt angle between those axes. The red and blue boundaries represent the inner and outer membrane surfaces. The black arrow is perpendicular to the membrane, which indicates that the receptor is placed in the membrane in the orientation predicted by PPM. On the other hand, the gray angle for the ligand subunit, is tilted in the figure, which indicates that the docked pose of the ligand is tilted from the predicted orientation by PPM. **(c)** The height constraint used by Mem-LZerD. The black dots represent points on the molecular surfaces of the receptor and ligand subunits sampled by the initial stages of LZerD. A black dot is defined as the middle point between the two sampled surface points used to define the docked pose of the subunits. Each surface point has a height value relative to the membrane midplane, and thus a pose using the two surface points is assigned with a height which is the greater of the heights of the two surface points that define the orientation. The height difference of the two black dots, Δh , is used to prune the docking search.

parameters, the average enrichment factor (EF) and recall were calculated. The enrichment factor is defined as the ratio of the fraction of acceptable models in the docking output after height and tilt constraints are applied relative to the fraction of all acceptable models in the entire pool of docking decoys. Recall is defined as the fraction of acceptable models remaining after the height difference and the tilt angle constraints are applied relative to the entire set of acceptable models. As shown in Figure 2c, a combination of 8 Å height and 0.4 radian tilt differences exhibits the highest EF value

before a dramatic loss of recall is observed: these settings resulted in an EF of 20.3 and a recall of 0.27, while the next-highest EF of 33.1 dropped the recall to 0.06. The highest EF of 135.7 would have yielded a recall of only 0.18. These cutoffs also agree with those found in existing literature, although the height referred to in such past work has been the subunit centroid height and not a property of the molecular surface.³⁶ Thus, while the formulation of height in Mem-LZerD differs substantially from that in existing methods, this parallel illustrates that Mem-LZerD heights can

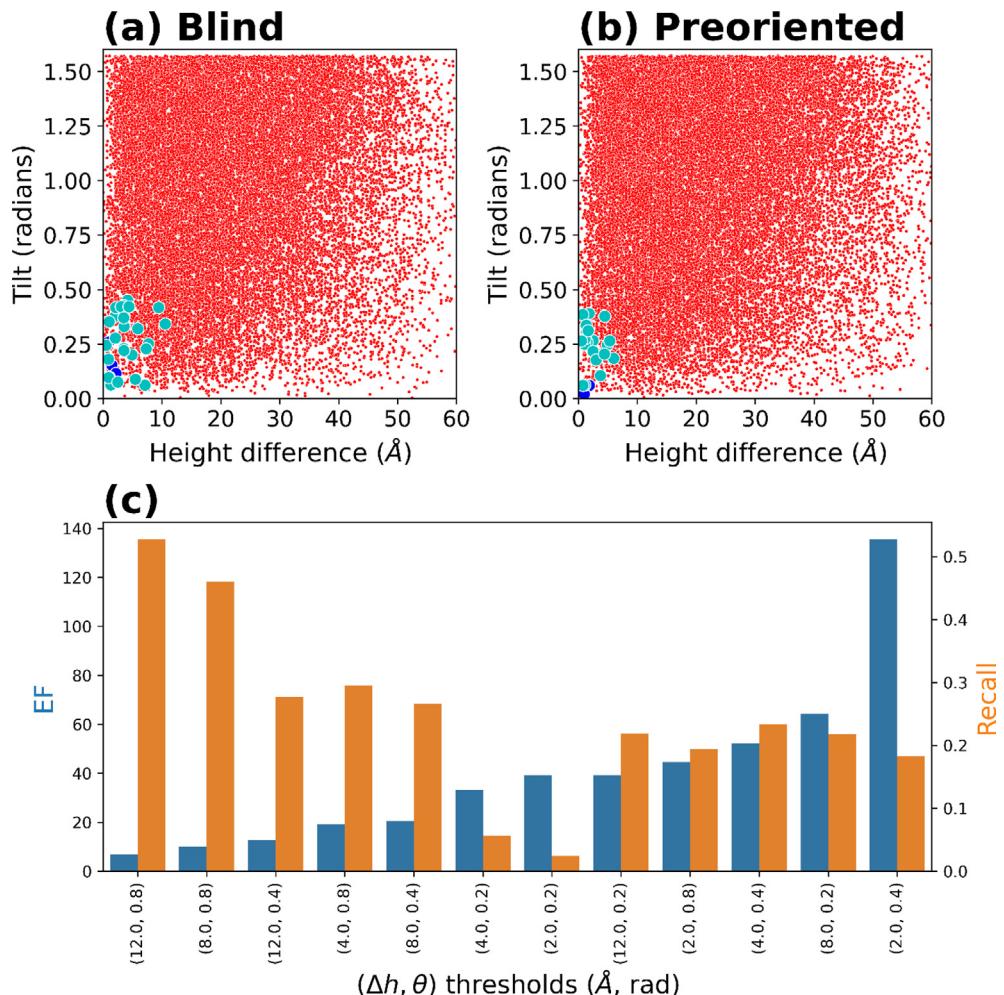


Figure 2. Height difference and tilt angles of docking decoys. Deviation from ideal membrane spatial positioning and orientation using **(a)** blindly oriented and **(b)** preoriented docking inputs for PDB 1M56, as included in the usual top 50,000 models from LZerD. Red: incorrect quality models. Cyan: acceptable quality models. Blue: medium models. The docking of 1M56 did not yield any high-quality models. **(c)** The average enrichment factor (blue) and recall (orange) of 20 targets using different combinations of height difference and tilt angle cutoff values using predicted subunit orientations.

generally be interpreted using the same understanding as of centroid heights. [Supplemental Video S1](#) illustrates the difference between the decoy poses sampled by this procedure and the original LZerD search for an intramembrane aspartate protease complex (PDB 4HYG), using the final chosen constraint cutoffs of 8 Å and 0.4 radians. Using this formulation of constraints, many clearly impossible models, including many sideways or totally dis-immersed poses, which are sampled by LZerD are not sampled by Mem-LZerD. In this way, both the running time of the sampling and the number of raw unclustered decoys were reduced by a factor of 74 on average across our benchmark set. In the case of the decoy distributions shown for target 1M56 in [Figure 2ab](#), the numbers of models were reduced by factors of 66 and 69 for blind and preoriented docking, respectively.

Docking results on the Memdock bound and unbound dataset

Mem-LZerD was trained on a nonredundant set of 20 transmembrane protein complexes and benchmarked on nonredundant sets of 21 and 44 transmembrane protein complexes originally used by Memdock³⁶ and newly constructed for Mem-LZerD, respectively, as described in Methods. Mem-LZerD was additionally benchmarked on a nonredundant set of 92 peripheral membrane protein complexes. The main evaluation results are summarized in [Tables 1](#) and [Table 2](#), covering the existing benchmark set³⁶ and our new benchmark set, respectively. Modeling for a target was considered successful if one of more of the top 10 output models was of at least acceptable CAPRI quality. According to the standard CAPRI criteria, a model

Table 1 Docking performance of the 21 individual targets of the Memdock benchmark set. RFH is the rank of the first CAPRI-acceptable model generated, if any, for a given target. RFH must exist and be at most 10 for modeling to be considered successful. n/a indicates that there was no unbound model modeled for a target in the published study. X indicates that there were no docked models of at least acceptable quality in the analyzed output; for Mem-LZerD, this is limited to 50,000, the pre-clustering model cutoff; for Memdock, this is limited to the top 100 models. We considered only the top 10 models from either method in our evaluation, but scored model ranks worse than 10 are shown here for completeness in parentheses. The I-RMSD columns list the best interface RMSD among the top 10 models. Separate preoriented docking outcomes are not shown for regular LZerD runs since the original LZerD pipeline does not consider input model orientation. ^a These evaluations of Memdock and JabberDock from the paper first describing their use for membrane docking used premade orientations from OPM, which are orientations of the full bound complexes rather than the subunits independently. The Mem-LZerD unbound benchmark otherwise shown is fully unbound docking, with membrane orientations calculated by PPM separately for each subunit without regard to the full complex. However, here we additionally show Mem-LZerD results when using ground truth orientations as in the original Memdock and JabberDock studies.

Target	Bound				Unbound Blind				Unbound Preoriented ^a			
	Mem-LZerD I-RMSD (Å)	Mem-LZerD RFH	Memdock RFH	LZerD RFH	Mem-LZerD I-RMSD (Å)	Mem-LZerD RFH	LZerD RFH	Mem-LZerD I-RMSD (Å)	Mem-LZerD RFH	Memdock RFH	JabberDock RFH	
1e12_A_C	0.6	6	(28)	(268)	10.2	(80)	(97)	3.5	10	(57)	2	
1m0l_A_C	2.2	1	(74)	(1429)	11.4	(28)	(709)	4.9	(17)	(49)	7	
1m56_A_C	1.1	1	2	(97)	1.2	1	(17)	2	2	(11)	X	
1q90_B_S	2.3	2	1	(2661)	2.7	1	(1106)	1.9	1	1	5	
1rc2_A_B	1.2	2	2	(74)	0.8	1	(61)	0.9	1	n/a	n/a	
1rwt_A_C	12.8	(56)	(25)	(19150)	5.8	(1181)	(19985)	19.2	(775)	n/a	n/a	
1u7g_A_C	0.6	2	2	2	1.1	2	(38)	1.3	2	n/a	n/a	
1zoy_C_D	3.5	1	1	(155)	3.8	10	(615)	3	1	1	(159)	
2a65_A_B	5.6	(24)	4	9	4.4	5	(331)	10.8	(11)	n/a	n/a	
2f2b_A_C	3.1	1	1	(145)	3.1	1	(49)	3.2	2	n/a	n/a	
2qjy_A_D	0.7	2	X	(1749)	1.1	2	(2821)	4.3	4	1	3	
3chx_B_J	11.7	(6745)	X	(1005)	14.8	X	X	17.2	X	X	8	
3dh4_A_B	1.6	7	3	(557)	4.6	(13)	(318)	2.1	5	n/a	n/a	
3h5m_A_B	0.7	1	1	(114)	3.8	1	(375)	3.7	1	n/a	n/a	
3hfx_A_B	14	(382)	6	(13054)	17.3	(969)	(11843)	19.3	(236)	n/a	n/a	
3kly_A_B	0.4	1	1	(165)	1.9	1	(366)	2	4	1	5	
3m73_A_K	0.8	1	1	(2635)	0.8	4	(2651)	0.9	1	n/a	n/a	
3oe0_A_B	10.8	(119)	X	(2222)	17.3	X	X	13.5	X	(45)	X	
3rvy_A_B	1.3	1	5	(226)	2	4	(129)	2.1	1	X	X	
3tui_A_B	1	1	1	(233)	1.7	1	(365)	1.6	1	n/a	n/a	
4dkl_A_B	2.2	1	4	(395)	16.4	X	X	18.1	X	X	1	
Total Successes:	16	15	2		13	0		14	4	7		
	(76.2%)	(71.4%)	(9.5%)		(61.9%)	(0.0%)		(66.7%)	(19.0%)	(33.3%)		

Table 2 Docking performance of the 44 individual targets of the Mem-LZerD transmembrane protein benchmark test set. Parentheses indicate that no models of at least CAPRI-acceptable quality were ranked within the top 10. The columns list the best of each measure among the top 10 models.

Target	Bound			Unbound			Regular LZerD Unbound		
	I-RMSD (Å)	L-RMSD (Å)	f_{nat}	I-RMSD (Å)	L-RMSD (Å)	f_{nat}	I-RMSD (Å)	L-RMSD (Å)	f_{nat}
1a0s	1.0	2.3	0.78	(16.9)	(27.8)	0.00	(36.69)	(68.60)	(0.00)
1kqf	1.6	2.7	0.46	2.5	5.1	0.22	(16.29)	(33.42)	(0.00)
1p7b	1.3	2.9	0.62	4.6	8.1	0.12	(17.28)	(32.12)	(0.02)
2yyx	(23.1)	(38.9)	(0.00)	1.5	2.5	0.52	(28.46)	(60.53)	(0.00)
3a7k	1.4	3.2	0.51	2.0	3.6	0.54	(22.47)	(33.77)	(0.00)
3b4r	1.7	2.7	0.53	2.3	5.4	0.49	(16.74)	(33.69)	(0.00)
3b8n	(22.3)	(44.1)	(0.00)	(48.9)	(59.8)	0.00	(47.60)	(51.36)	(0.00)
3pjz	2.0	4.8	0.44	2.2	3.6	0.46	(16.98)	(35.79)	(0.00)
3qnn	1.6	3.7	0.42	1.3	2.6	0.63	1.29	2.57	0.63
3ux4	1.4	3.3	0.63	2.3	5.1	0.51	(14.75)	(28.60)	(0.00)
4ezc	1.8	3.7	0.39	1.4	2.0	0.64	(15.83)	(32.30)	(0.00)
4hkr	1.7	3.1	0.39	(11.2)	(12.2)	(0.10)	(17.99)	(28.63)	(0.00)
4hyg	2.8	5.6	0.46	5.0	6.5	0.51	(24.17)	(58.72)	(0.00)
4o6m	0.6	1.9	0.85	1.1	2.6	0.70	1.06	3.76	0.70
4p6v	0.9	2.0	0.84	2.5	3.8	0.64	(20.67)	(43.74)	(0.00)
4pl0	(16.9)	(32.6)	(0.00)	(21.8)	(36.3)	(0.00)	(30.87)	(57.20)	(0.00)
4r1i	0.7	1.3	0.80	1.2	2.6	0.78	1.17	2.59	0.78
4ri2	2.5	4.8	0.40	4.8	6.4	0.26	(13.04)	(23.82)	(0.03)
4rp8	2.1	4.1	0.58	2.2	4.0	0.55	(23.41)	(56.31)	(0.00)
4wd8	4.2	8.9	0.24	4.6	9.0	0.18	1.96	3.88	0.51
5a1s	0.8	1.4	0.83	2.3	4.5	0.79	2.32	4.54	0.79
5aex	0.5	1.1	0.94	(13.2)	(21.5)	(0.02)	(36.06)	(68.22)	(0.00)
5i6c	1.9	5.8	0.33	2.9	3.9	0.44	(21.19)	(56.41)	(0.00)
5j4i	0.8	1.8	0.75	3.4	4.7	0.48	(21.82)	(61.91)	(0.00)
5o65	0.9	1.9	0.91	1.9	3.4	0.44	(22.93)	(40.19)	(0.00)
5sv9	2.2	4.2	0.44	4.7	6.8	0.23	(23.02)	(42.10)	(0.00)
6c96	(4.6)	(12.2)	(0.16)	(6.9)	(11.6)	(0.02)	(38.81)	(93.12)	(0.00)
6hcp	2.1	6.7	0.48	(40.9)	(43.2)	(0.00)	(39.88)	(70.88)	(0.00)
6k1h	2.9	5.9	0.36	2.0	2.7	0.43	(28.77)	(61.40)	(0.00)
6l85	0.9	2.7	0.90	3.2	4.0	0.43	(7.84)	(13.92)	(0.01)
6n1g	0.6	1.4	0.84	2.7	6.1	0.38	(20.73)	(49.61)	(0.00)
6o58	(20.8)	(67.3)	(0.00)	(23.2)	(46.6)	(0.00)	(24.87)	(53.92)	(0.00)
6rko	0.8	1.8	0.72	4.8	8.3	0.49	(4.77)	(10.08)	(0.49)
6rvx	2.1	6.0	0.45	3.8	7.3	0.22	(33.71)	(60.68)	(0.00)
6vgc	2.8	5.6	0.28	3.6	12.8	0.41	(21.14)	(33.60)	(0.00)
6xjh	1.3	1.9	0.59	3.8	6.0	0.22	(17.85)	(34.48)	(0.00)
6y7f	1.3	5.7	0.63	3.3	14.0	0.38	(13.83)	(49.53)	(0.00)
6ys8	(9.1)	(20.0)	(0.01)	(14.7)	(24.5)	(0.00)	(15.03)	(30.88)	(0.00)
7ck6	2.9	9.6	0.24	1.9	6.0	0.37	(16.94)	(40.07)	(0.00)
7cxr	2.6	5.9	0.45	2.8	5.1	0.52	2.82	5.11	0.31
7czb	2.0	6.4	0.68	3.3	7.4	0.56	(27.62)	(58.81)	(0.00)
7jk8	1.8	3.7	0.69	3.5	5.1	0.28	(16.30)	(32.98)	(0.00)
7lhv	<u>1.0</u>	<u>2.4</u>	<u>0.67</u>	<u>5.3</u>	<u>9.8</u>	<u>0.35</u>	<u>(33.22)</u>	<u>(72.68)</u>	<u>(0.00)</u>
7oqz	<u>1.5</u>	<u>3.2</u>	<u>0.38</u>	<u>2.4</u>	<u>3.9</u>	<u>0.40</u>	<u>(13.53)</u>	<u>(23.66)</u>	<u>(0.05)</u>
Total successes:	38 (86.3%)			35 (79.5%)			6 (13.6%)		

is acceptable if it has a fraction of native contacts (f_{nat}) of at least 0.1 and either an interface RMSD (I-RMSD) of at most 4.0 Å or an L-RMSD of at most 10.0 Å.^{18,21,51,52} A rank-of-first-hit (RFH) value (as shown in Table 1) that is as high as 10 is considered a success for the given method on the given target.

To examine the precision of docking using Mem-LZerD versus past rigid-body methods, we compared the performance of Mem-LZerD against the Memdock benchmark set,³⁶ as shown in

Table 1. The Memdock algorithm is designed specifically for α -helical transmembrane proteins, rather than general membrane proteins which may be β -barrels or not transmembrane at all. Thus, we compare using Memdock's original benchmark set of 21 α -helical transmembrane complexes. Memdock achieved a success rate of 71.4% (15 out of 21 cases) for bound docking, while Mem-LZerD achieved a success rate of 76.2% (16 out of 21 cases; Table 1, Bound column group). On this

same benchmark, the baseline LZerD pipeline, which does not consider membrane constraints, succeeded on no targets for unbound docking and 2 of 21 (9.5%) targets for bound docking. Transmembrane docking targets were thus shown to largely require specialized methods, and Mem-LZerD achieved a higher success rate than such existing specialized methods.

For unbound docking on the Memdock benchmark (Table 1, Unbound Blind and Unbound Preoriented column groups), results from past docking experiments using both Memdock and JabberDock were available, and we evaluated the unbound docking performance of Mem-LZerD against these. Here, Mem-LZerD achieved an overall success rate of 61.9% (Table 1, Unbound Blind column group), while Memdock and JabberDock achieved 19.0% and 33.3% respectively overall (Table 1, Unbound Preoriented column group). The Unbound Preoriented column group, compares unbound docking performance when knowledge of the ground truth subunit orientations is assumed. Here, we calculated preoriented unbound input models by superimposing them individually to their native complex structures as oriented in OPM. The preoriented Mem-LZerD success rate was then 66.7% (14 out of 21) overall. Mem-LZerD thus overall outperformed the 33.3% success rate of JabberDock on preoriented inputs, despite requiring no expensive molecular dynamics simulation, while still performing well on targets with only predicted orientations.

Docking results on the new transmembrane protein dataset

On the additional transmembrane benchmark unbound test set (Table 2), which included both α -helical and non- α -helical transmembrane proteins such as β -barrels, 35 of 44 targets (79.5%) were modeled by Mem-LZerD to at least acceptable CAPRI quality. This is broadly consistent with the performance demonstrated by Mem-LZerD on the smaller Memdock benchmark, which only included pairs of α -helical transmembrane protein chains with an interface immersed in the membrane. Although this new benchmark includes full complexes rather than just pairs of chains and thus might be expected to offer increased opportunities to establish surface complementarity at the protein–protein interface, unconstrained rigid-body docking is not generally able to model these targets. Regular LZerD was only able to model 6 out of 44 (13.6%) targets in this benchmark to at least acceptable CAPRI quality. As indicated by the resulting low or zero f_{nat} values shown in the Regular LZerD Unbound column group of Table 2, the membrane environment severely interferes with recognition of the protein–protein interface. For example, in the case of the zinc metalloprotease mJS2P (PDB

3B4R), discussed in detail in Case Study 1, regular docking missed the interface entirely, while Mem-LZerD yielded a model of acceptable CAPRI quality. Mem-LZerD is of course dependent on the accuracy of its predicted subunit orientations, and this is reflected in the result for the bacterial magnesium transporter MgtE (PDB: 2YVX), where unbound docking succeeded but bound docking did not. In the absence of the full complex structure, PPM did not orient the subunits within the tilt angle cutoff used by Mem-LZerD. Thus, while acceptable poses were sampled in the full LZerD search, they were not emitted by Mem-LZerD as they all violated the tilt angle constraint. On the other hand, PPM oriented the unbound subunit structures appropriately. PPM, while a capable prediction method, is not explicitly designed to yield bound-state orientations from unbound-state structures. Thus, future developments of membrane docking should include an orientation prediction method designed for that purpose.

Docking results on the peripheral membrane protein dataset

On the peripheral benchmark unbound test set (Supplementary Table S1), 15 of 92 targets (16.3%) were modeled to at least acceptable CAPRI quality within the top 10 models. This success rate is substantially lower than what was observed for transmembrane docking (61.9% in Table 1), as might be expected from the intuition that peripheral membrane proteins should be more difficult to orient to the membrane, especially when the subunits are oriented individually. Indeed, an unconstrained docking search sampled poses of at least acceptable CAPRI quality for 75 of 92 peripheral targets (81.5%) out of around 3 million total decoys typically for each target, as shown in the LZerD Sampling column of Supplementary Table S1. Among the sampled decoys generated, the standard LZerD pipeline was able to rank poses of at least acceptable CAPRI quality in the top 10 for 53 of 92 targets (57.6%), which is substantially higher than the 16.3% by Mem-LZerD, demonstrating that methods designed for soluble proteins are better at selecting good peripheral complex poses than the state-of-the-art PPM is at orienting the subunits. Figure 3 illustrates the particular challenges involved in docking peripheral membrane proteins. Figure 3a compares the modeling performance of regular LZerD to that of Mem-LZerD with predicted orientations, in terms of I-RMSD. For 59 out of 92 cases (64.1%), regular LZerD yielded a better (smaller) I-RMSD.

Figure 3b and c show the Mem-LZerD output model quality for peripheral complexes as compared to transmembrane complexes in terms of the I-RMSD and L-RMSD CAPRI measures, respectively. While the distribution of each

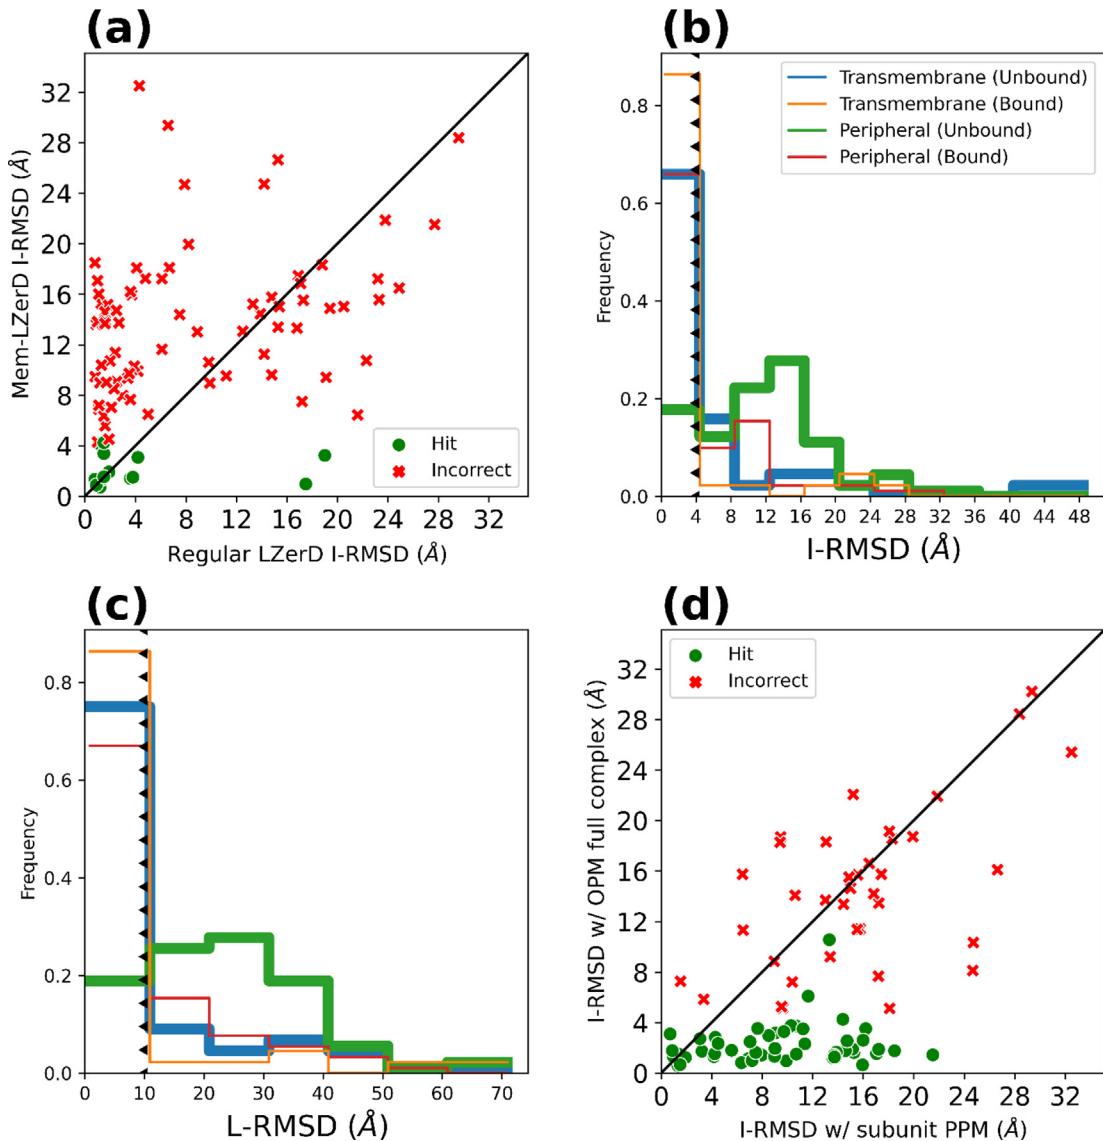


Figure 3. Comparison of peripheral membrane protein docking outcomes. **(a)** I-RMSD of Mem-LZerD using predicted orientations vs regular LZerD. Green: targets where Mem-LZerD was able to generate models of at least acceptable CAPRI quality in the top 10 when using predicted subunit orientations. Red: targets where Mem-LZerD failed to generate models of at least acceptable CAPRI quality in the top 10 when using predicted orientations. **(b)** I-RMSD best in top-10 distribution for transmembrane complexes vs peripheral membrane complexes. The triangle-dashed lines mark the CAPRI acceptable threshold for each metric, with the triangles pointing toward the acceptable side. **(c)** The same as (a), but for L-RMSD. **(d)** Performance of Mem-LZerD on peripheral membrane proteins using predicted versus ground truth membrane orientations. X-axis: the best I-RMSD among the top 10 models yielded by the Mem-LZerD pipeline when using the blind PPM subunit orientations. Y-axis: the best I-RMSD among the top 10 models yielded by the Mem-LZerD pipeline when using ground truth subunit orientations by superimposing to the oriented OPM complex. Green: targets where Mem-LZerD was able to generate models of at least acceptable CAPRI quality in the top 10 when using ground truth orientations. Red: targets where Mem-LZerD failed to generate models of at least acceptable CAPRI quality in the top 10 when using ground truth orientations.

measure peaks within the CAPRI cutoffs of 4 Å and 10 Å, respectively, for transmembrane complexes cases, this is not the case for the peripheral complex targets. For the peripheral targets, the quality distributions peak far outside the corresponding acceptability thresholds.

A comparison in terms of the remaining CAPRI measure, f_{nat} , is shown in [Supplementary Figure S1](#). To investigate what effect improvements to the subunit orientation predictions could have on the accuracy of the docking predictions, we recomputed the peripheral

membrane protein targets using their bound orientations by superimposing the unbound subunits into the native structure taken from the OPM database, in the same way as in the transmembrane comparison (Figure 3d). I-RMSD improved when using bound (ground truth) orientations from OPM for 69 out of 92 cases (75.0%). With the bound orientations, models of at least acceptable CAPRI quality were built for 59.3% of cases within the top 10 (green dots in Figure 3d), while only 16.3% yielded acceptable models within the top 10 when using predicted subunit orientations by PPM. Since peripheral membrane proteins are by definition not immersed in the membrane, they are physically closer to soluble proteins than transmembrane proteins. This disparity confirmed that most of the peripheral targets are not sufficiently well oriented by PPM to dock using this method of constraining the search space. Supplementary Table S1 reflects this comparison in the Unbound Blind Mem-LZerD and Unbound Preoriented Mem-LZerD column groups. Future expansions of PPM or the inclusion of additional information from experiments which provide accurate orientation information, have the potential to change the outcome of peripheral complex docking from likely unsuccessful to likely successful.

Case study 1: Zinc metalloprotease mjS2P (PDB 3B4R)

Highlighted in Figure 4a is the zinc metalloprotease mjS2P complex from *Methanocaldococcus jannaschii*, a protease capable of cleaving proteins in the lipid membrane environment.⁵³ This complex is a key part of signaling by regulated intramembrane proteolysis, a signaling mechanism which has been established as conserved from bacteria to humans.⁵⁴ While this target was not tested in the Memdock or JabberDock studies, a run of AlphaFold as a substitute failed to model this complex, generating a model of incorrect CAPRI quality with an I-RMSD of 19.1 Å, and L-RMSD of 43.2 Å, and an f_{nat} of 0.01, shown in the right of Figure 4a. Although this model looks passably folded, it does not appear to even contain an appropriate transmembrane axis for the full complex. Indeed, the proper transmembrane axes of the two subunits are roughly perpendicular in this model, rather than roughly parallel. Mem-LZerD on the other hand was able to generate a model of acceptable CAPRI quality with an I-RMSD of 2.3 Å, and L-RMSD of 5.4 Å, and an f_{nat} of 0.49, shown in the left of Figure 4a. While the orange deviation lines show that the regions away from the protein–protein interface are not as precisely modeled, the interface residues are properly located. The backbone deviations away from the interface are largely parallel to the plane of the membrane and do not entail substantial disimmersion. The inclusion by Mem-LZerD of

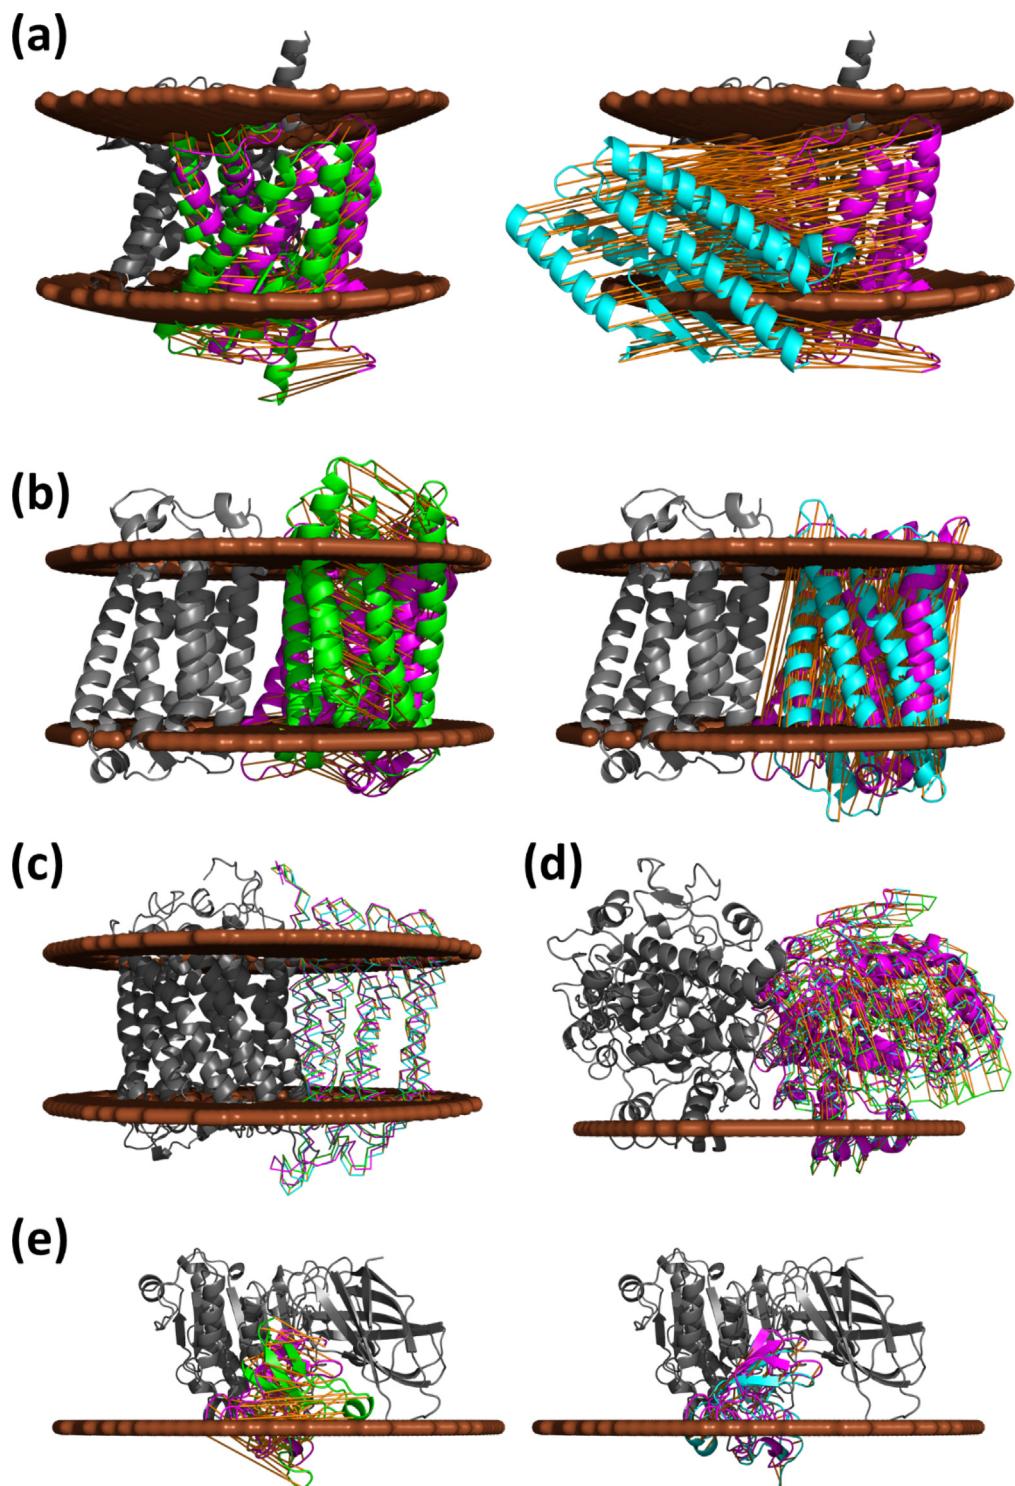
reversed orientations was key to the successful modeling of this target, as otherwise Mem-LZerD would have only yielded incorrect models. Remarkably, all the top 10 models were of acceptable quality. If we consider all the models generated, there were 48 acceptable or better models.

Case study 2: Elongation of very long chain fatty acids protein 7 (PDB 6Y7F)

The next case study, shown in Figure 4b, is the human ELOVL7 elongase complex, which is responsible for catalyzing the fatty acid elongation cycle⁵⁵ and is involved in a variety of serious genetic diseases in humans.^{56–58} In the ELOVL7 homodimer observed by experiment, the substrate binds deep in a narrow tunnel formed by the structure of one subunit. The substrate binding site is not directly formed by the protein–protein interaction interface. The active site and interaction surface of this enzyme were noted to not be well-conserved in 2020 when its PDB entry 6Y7F was deposited, and modern template search techniques such as HHpred,⁵⁹ using database PDB_mmCIF70_17_Apr, found no hits other than PDB 6Y7F itself. While this target was not tested in the Memdock or JabberDock studies, a run of AlphaFold as a substitute failed to model this complex, generating a model of incorrect CAPRI quality with an I-RMSD of 20.2 Å, and L-RMSD of 35.0 Å, and an f_{nat} of 0.00, shown in the right of Figure 4b. While the experimentalists observed the complex in a head-to-tail dimer both in crystal and in solution, AlphaFold instead generated head-to-head dimer models; in other words, AlphaFold placed one of the subunits upside down relative to the native complex structure, as indicated by the orange deviation lines. Mem-LZerD on the other hand was able to generate a model of acceptable CAPRI quality with an I-RMSD of 3.3 Å, and L-RMSD of 14.0 Å, and an f_{nat} of 0.38, shown in the left of Figure 4b. The inclusion by Mem-LZerD of reversed orientations was key to the successful modeling of this target as well, as otherwise Mem-LZerD would have only yielded incorrect models. Among the top 10, three models were of acceptable quality, while 10 acceptable models were generated in the entire model pool.

Case study 3: Cytochrome-c oxidase (PDB 1M56)

The next example (Figure 4c) is a cytochrome-c oxidase complex from *Cereibacter sphaeroides*, a transmembrane enzyme found in the inner mitochondrial and cell membranes of bacteria.⁶⁰ For this target, both Memdock and JabberDock failed to retrieve a model of acceptable CAPRI quality in unbound docking, yielding only models of incorrect CAPRI quality within the top 10 models. Mem-LZerD on the other hand was able to generate a two models of medium CAPRI quality, one with an I-RMSD of 1.5 Å, and L-RMSD of 2.4 Å, and



an f_{nat} of 0.71, and another with an I-RMSD of 1.2 Å, and L-RMSD of 1.9 Å, and an f_{nat} of 0.84. These two models thus both exceeded the criteria for acceptability. Among the top 10, Mem-LZerD generated three models of acceptable or medium quality, while 31 models of acceptable or medium quality were generated in the entire model pool. For this target, AlphaFold was also able to generate a model of high quality, with an I-RMSD of 1.0 Å, an L-RMSD of 1.3 Å, and an f_{nat} of 0.90.

Case study 4: Prostaglandin G/H synthase 2 (PDB 1CX2)

The last two examples are from the peripheral membrane protein dataset. Figure 4d shows a murine prostaglandin G/H synthase 2 (PTGS2) homodimer complex, a peripheral membrane enzyme in the nuclear and endoplasmic reticulum membranes of mice.⁶¹ PTGS2 is involved in inflammatory response pathways, and interactions with homologous human proteins are the molecular basis of many non-steroidal anti-inflammatory drugs (NSAIDs).^{62,63} PTGS2 binds to membranes peripherally via a helical amphipathic domain, which the protein–protein interface of PTGS2 does not include.⁶⁵ For this target, Mem-LZerD using predicted orientations from PPM was able to produce a model of acceptable CAPRI quality with an I-RMSD of 3.2 Å, and L-RMSD of 6.4 Å, and an f_{nat} of 0.26. 31 models of acceptable or medium quality were generated in total. With the correct orientation information taken from the full-complex orientation from OPM, Mem-LZerD yielded a model of medium CAPRI quality with an I-RMSD of 1.7 Å, an L-RMSD of 2.7 Å, and an f_{nat} of 0.51. AlphaFold yielded a model of high quality with an I-RMSD of 0.8 Å, an L-RMSD of 1.2 Å, and an f_{nat} of 0.90.

Case study 5: Pancreatic triacylglycerol lipase with colipase (PDB 1LPA)

The last case study (Figure 4e) is a human pancreatic triacylglycerol lipase (PTL), a peripheral membrane protein involved in the metabolism of fats,^{66,67} complexed with a porcine colipase (CLPS), shown in magenta in Figure 4e, a cofactor which facilitates membrane anchoring.⁶⁸ Due to the presence of amphipathic bile salts in the physiological environment of PTL, the smaller CLPS is required as a cofactor to facilitate a binding mode with the catalytic site of PTL in sufficiently close contact with the membrane⁶⁸ to facilitate its function.^{66,67} For this target, Mem-LZerD using inputs blindly oriented by PPM was unable to produce a model of acceptable CAPRI quality, with the best of the top-10 models only reaching an I-RMSD of 10.7 Å, and L-RMSD of 18.4 Å, and an f_{nat} of 0.02, as shown in the left of Figure 4e. In this incorrect model, CLPS was placed at essentially the correct binding site, but is rotated about an axis roughly parallel to the membrane, completely disrupting the modeling of the native protein–protein interface. The blind orientation predictions of the subunits by PPM did not properly characterize the feasible region of docking poses for this target. When regular LZerD without constraints was used, we found that a correct binding pose with an I-RMSD of 3.7 Å was sampled, but it was not selected by the scoring function within top-10 models, instead yielding an even worse an I-RMSD of 22.3 Å, an L-RMSD of 51.8 Å, and an f_{nat} of 0.00 within the top 10. However, if we consider the entire model pool, there were 19 models of acceptable or medium quality. When instead modeling using the full-complex orientation from OPM, Mem-LZerD yielded a model of medium CAPRI quality with an I-RMSD of 3.7 Å, an L-RMSD of 5.0 Å, and an f_{nat}

Figure 4. Example modeling of membrane protein complexes. Gray: the native receptor structure. Magenta: the native ligand binding pose. Brown: the membrane locations calculated by PPM. Orange: $\text{C}\alpha$ deviations of one subunit when the other is superimposed. (a) Zinc metalloprotease mjs2P (PDB: 3B4R). Left: Mem-LZerD yielded a model (green) with an I-RMSD of 2.3 Å, an L-RMSD of 5.4 Å, and an f_{nat} of 0.49. Right: AlphaFold yielded a model (cyan) with an I-RMSD of 19.1 Å, and L-RMSD of 43.2 Å, and an f_{nat} of 0.01. (b) Elongation of very long chain fatty acids protein 7 (PDB: 6Y7F). Left: Mem-LZerD yielded a model (green) with an I-RMSD of 3.3 Å, an L-RMSD of 14.0 Å, and an f_{nat} of 0.38. Right: AlphaFold yielded a model (cyan) with an I-RMSD of 20.2 Å, and L-RMSD of 35.0 Å, and an f_{nat} of 0.00. (c) Poses of medium CAPRI quality generated by Mem-LZerD for cytochrome-c oxidase (PDB: 1M56). Cyan: first hit, an I-RMSD of 1.5 Å, an L-RMSD of 2.4 Å, and an f_{nat} of 0.71. Green: second hit, with an I-RMSD of 1.2 Å, an L-RMSD of 1.9 Å, and an f_{nat} of 0.84. Orange: displacements between the two poses, with an RMSD of 2.5 Å RMSD and maximum of 3.6 Å. (d) Prostaglandin G/H synthase 2 (PDB: 1CX2). Mem-LZerD yielded a model of acceptable CAPRI quality (green) with an I-RMSD of 3.2 Å, an L-RMSD of 6.4 Å, and an f_{nat} of 0.26, while doping with ground truth subunit orientations yielded a model of medium CAPRI quality (cyan) with an I-RMSD of 1.7 Å, an L-RMSD of 2.7 Å, and an f_{nat} of 0.51. Blind and doped model displacements are shown in orange and red respectively. (e) Pancreatic triacylglycerol lipase/colipase complex (PDB 1LPA). Left: Mem-LZerD yielded an incorrect model with an I-RMSD of 10.7 Å, and L-RMSD of 18.4 Å, and an f_{nat} of 0.02. Right: Mem-LZerD doped with ground truth subunit orientations yielded a model of acceptable CAPRI quality (cyan) with an I-RMSD of 3.7 Å, an L-RMSD of 5.0 Å, and an f_{nat} of 0.54.

of 0.54, as shown in the right of [Figure 4e](#). Thus, in this case, the correct orientation was required for a correct pose to be discovered by the scoring function. AlphaFold yielded a model of medium quality with an I-RMSD of 3.7 Å, an L-RMSD of 3.4 Å, and an f_{nat} of 0.62.

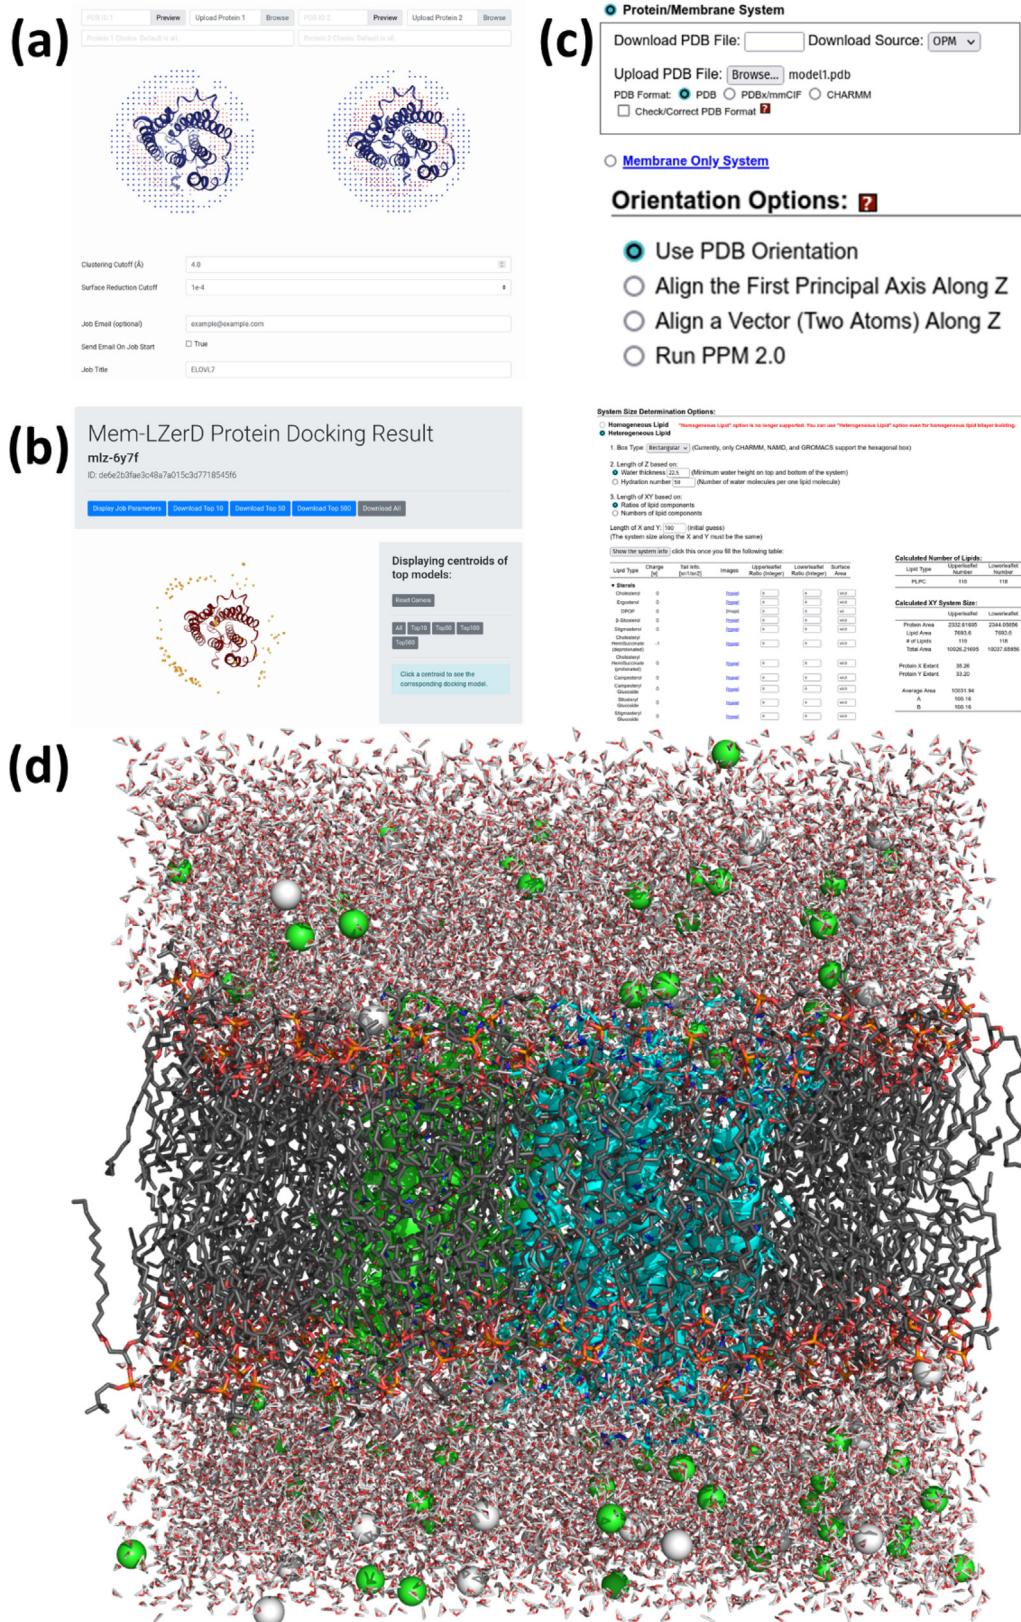
In these five case studies, Mem-LZerD yielded models of acceptable or medium quality within the top 10 in four cases with predicted membrane orientations and in all five cases when membrane orientations are known. AlphaFold produced a model of medium or high quality for three case studies. Although AlphaFold produces accurate models for membrane proteins as it does for other types of proteins in general, Mem-LZerD has notable advantageous features. As Mem-LZerD performs docking with provided membrane orientations, resulting models are intuitive and interpretable in the context of the membrane environment. Therefore, it is suitable for conducting molecular dynamics simulation as we discuss in the subsequent section. Also, Mem-LZerD produces multiple distinct acceptable poses, which may serve as multiple reasonable initial conditions of molecular dynamics simulations or to consider an ensemble of docking conformations that may be useful to study encounter complexes⁶⁹ or cell-scale molecular crowding dynamics.⁷⁰ Also, Mem-LZerD can consider user-input constraints, such as residue-residue distances, as it is equipped with the LZerD program. AlphaFold can be accurate, as mentioned above, but often only produces one pose among its models, and the models are difficult to modify in cases where the generated model is obviously incorrect. Unless a specific variant of the AlphaFold program, such as Distance-AF,⁷¹ is used, AlphaFold does not accept additional biological information for such integrative modeling.

Preparation of molecular dynamics with membrane docking in the LZerD webserver

Due to the limited ability of models like AlphaFold to sample many interaction poses out-of-the-box, biologists interested in modeling the dynamics of proteins in a membrane may need to arrange individual protein chains to begin modeling the states desired, or even reasonable states. This difficulty is exemplified by Case Study 2 above, where despite experimental evidence to the contrary, the AlphaFold models of ELOVL7 took on entirely opposite orientations. Mem-LZerD can assemble ELOVL7 into reasonable bound poses which can then be modeled in explicit lipid solvent and simulated. Mem-LZerD is freely accessible via the LZerD webserver for protein docking, available at <https://lzerd.kiharalab.org>,^{16,17} which allows users without high-performance computers or without computing experience to easily model the

assembly of protein complexes. [Figure 5a](#) shows the input preparation for modeling poses of ELOVL7 using the LZerD webserver. Input files for docking can be taken directly from the OPM website (<https://opm.phar.umich.edu/>)⁴⁸ or can be oriented using the PPM website or locally-installed application. The webserver allows users to specify residue-based distance constraints on the docking, and more information on this and all functionality can be found in previous works^{16,17} and the About pages on the webserver. However, in this case the default options are sufficient, and we can proceed to submit the docking. Shown in [Figure 5b](#) is the results page of this docking. Note that the ligand pose centroids, shown as spheres, highlight the focus of the membrane docking search. From here, we can download our desired models and upload them to CHARMM-GUI (<https://www.charmm-gui.org>).⁵⁰ [Figure 5c](#) highlights the steps of processing a docked model using CHARMM-GUI. In this instance, we selected PLPC lipids and the NAMD output format. [Figure 5d](#) shows ELOVL7 solvated in the periodic membrane, and the left of [Supplemental Video S2](#) shows the equilibration and a 1 ns production molecular dynamics trajectory of this system. With these together, we see that the dynamics simulation proceeds smoothly.

To demonstrate the general reliability of Mem-LZerD in generating models suitable for molecular dynamics simulation, we considered models immersed in an explicit PLPC membrane and using the same parameters as above. For each case study complex, simulation runs were carried out starting from both a model of at least acceptable quality and an incorrect model (2 runs per case study) for a total of 10 simulations overall among the five case studies. While protein docking methods can be vulnerable to steric clashes, which can unphysically disturb or crash a molecular dynamics simulation, the initial and final scoring procedures of Mem-LZerD automatically discourage severe clashes between chains through consideration of excluded volume¹⁴ and atomic distance-based potentials,¹⁵ respectively. Following the cutoff used in CAPRI,¹⁸ two residues are considered clashing if any of their nonbonded heavy atoms are within 3 Å of each other. While 3 Å is too close for normal nonbonded atomic contacts, any number of clashes near this cutoff are easily resolved during minimization within molecular dynamics runs.¹⁵ [Supplemental Videos S3, S2, S4, S5, and S6](#) show the simulation trajectories for each case study complex in that order respectively, separately initialized from the noted acceptable (left) and incorrect (right) models. In all cases, the simulations completed without numerical divergences, crashes, or other issues, thus demonstrating that Mem-LZerD models are generally suitable for initializing simulations.



Materials and methods

Dataset construction

The transmembrane dataset used in this work was constructed by starting from 4218 nominally transmembrane protein structures from the OPM database.⁴⁹ Monomers were excluded, and entries were clustered at 25% sequence identity using MMseqs2,⁷² yielding 390 complexes. Structure-based clustering with an MM-align⁷³ TM-score cutoff of 0.50 yielded 85 complexes. From these, complexes which were not oriented inside the membrane and for which the publication did not indicate that the complex was in fact studied in a membrane were excluded, or which were otherwise beyond the scope of Mem-LZerD, yielding the final transmembrane dataset of 64 complexes. 20 complexes from this dataset were separated to serve as the training set for the scoring function, while the remaining 44 complexes served as the test set. For all targets, AlphaFold v2.2.2¹² with v2 weights and without template search was used to generate unbound subunit models.

The peripheral membrane complex dataset was constructed in a similar manner to the transmembrane dataset. Starting from 1607 complexes listed by OPM as “monotopic/peripheral”, 147 complexes with two subunits and each interacting with the membrane were extracted. After clustering as with the transmembrane dataset, 92 complexes remained, constituting the final peripheral membrane complex dataset. Peripheral complexes were not used to train the scoring function, and thus this dataset was not partitioned.

Rigid-body docking of subunits with regular LZerD

Each ligand protein structure subunit was docked with the receptor subunit using LZerD.^{14,74} LZerD is a shape complementarity-based rigid-body docking algorithm which employs a soft representation of the molecular surface to tolerate some small differences between the bound and unbound conformations. Using geometric hashing-based surface searching, LZerD generates many candidate dock-

ing poses, which are then scored according to the shape complementarity of their interfaces. These initial candidate poses, before any truncation or clustering, are referred to as the poses sampled by LZerD. For a given pair of input subunits, the set of docked models generated by LZerD was ordered by the LZerD shape score and truncated to a set of 50,000 models. This set was then clustered with an RMSD cutoff of 4.0 Å. The final stage of the regular LZerD pipeline is a rescoring according to the ranksum scoring function.^{15,75,76}

Membrane-aware pose constraints

The orientations in the membrane of models to be docked were predicted using the Positioning of Proteins in Membranes (PPM) software package^{48,49} using default settings. PPM outputs an oriented version of each input structure and also predicts the interior and exterior membrane surfaces. The separately oriented subunits were assumed to share a membrane midplane. Membrane type and thickness can be considered explicitly at this stage. In this study, the default membrane parameters of PPM were used. These predicted orientations are then used to augment the LZerD docking search. In LZerD, sample points on the surfaces of the receptor and the ligand proteins are matched to each other in the geometric hashing discussed above. To accomplish this, reference frames of, say, the ligand protein, are constructed using pairs of nearby points on the surface in combination with their surface normals. The other surface points of the ligand in a fixed neighborhood of the origin of this reference frame are translated and rotated such that their basis is changed to the reference frame, resulting in a fingerprint of the molecular surface. This fingerprint is then matched with fingerprints calculated from the receptor protein surface. This procedure was implemented using a space-partitioning tree data structure,⁷⁷ which is designed to efficiently search for points in Euclidean space by querying only relevant regions of space. By selecting appropriate planes with which to divide the space of hashed points, individual queries in the accumulation and recognition stage of the geometric hashing are made in sublinear time.

Figure 5. Example workflow using the LZerD webserver to prepare and run a molecular dynamics simulation of a docked model of target elongation of very long chain fatty acids protein 7 (PDB 6Y7F) in the membrane. (a) PPM-oriented receptor and ligand protein models uploaded to the LZerD webserver submission page. Default settings are used in the absence of residue distance information. (b) Mem-LZerD docking results shown in the LZerD webserver results viewer. Each sphere shows the centroid of a sampled pose. (c) Certain options used in generating NAMD input files using CHARMM-GUI. A PPM-oriented PDB file was uploaded, and its precomputed orientation was kept. PLPC lipids and a box size of 100 were used to construct the explicit simulation-ready files. (d) The final prepared coordinates for simulation.

In Mem-LZerD, this data structure is augmented with the distance of each point to the membrane midplane, which defines the height of the point. These height values for each point come directly from the output of PPM and are fixed once the geometric hashing data structure has been built at the beginning of the LZerD search. During a query to the geometric hashing data structure, hashed points tagged with a height different from the query height (Δh in Figure 1b) by more than the cutoff will not be emitted, and thus not contribute to the match vote accumulation of the geometric hashing. Unlike the 3-dimensional coordinates of the sample points, the height value is not transformed during the change of basis, and the internal structure thus differs from otherwise similar lifting transform techniques that bridge related entities in computational geometry.⁷⁸ This augmentation facilitates efficient imposition of constraints in terms of the relative midplane distance, essentially pruning the search space analyzed. For Mem-LZerD we used a relative midplane distance of 8 Å. As surface matches satisfying the relative midplane distance constraint are found, their docked orientations relative to their predicted orientations were further checked to satisfy a relative midplane axis constraint of 0.4 radians (θ in Figure 1c). The base orientation axes are fixed at the start of docking based on PPM, and this tilt angle check does not use the properties of the geometric hashing data structure, as it requires that the full rigid-body transformation have already been calculated. While these are the same numerical thresholds as were previously used by Memdock, we note that they are comparable but different quantities than those constrained in Memdock. We show a validation of these hyperparameters on the training set in Figure 2c based on the recall of models of at least acceptable CAPRI quality and the enrichment factor (EF, the ratio of the fraction of hits in the decoy set after constraining to the fraction of hits in the decoy set before constraining) of the pose space constrained to. When the parameter combinations were ordered by their corresponding EF values, we found that the recall of the constrained decoy set declined dramatically. As illustrated in Figure 2ab, nearly all of the search space outside this feasible region consists of incorrect models. Additionally, the implementation of the height constraint into the geometric hashing results in a running time 74 times faster compared to regular LZerD. We note that in cases where the membrane of the target membrane protein to model is known to restrict the orientations of proteins more rigidly, it can make sense for users to select a tighter tilt angle cutoff.

Docked model scoring

For selecting docked poses, we constructed a combined scoring function derived from methods

used in the scoring of IDP-LZerD and Flex-LZerD.³⁸⁻⁴⁰ This scoring function uses a bagged decision tree estimator⁷⁹ to combine the knowledge-based scoring functions GOAP⁸⁰, DFIRE,⁸¹ and ITScorePro,⁸² which are the three re-scoring functions combined in the usual ranksum score for competitive LZerD docking,^{15,75,76} as well as LZerD initial docking score, the cluster size from the 4 Å RMSD clustering normally used with LZerD, a membrane transfer energy defined as the sum of calculated side chain transfer free energies from water to the membrane environment⁸³ weighted by relative surface area, and order statistic terms ($OS_{(1)}$, $OS_{(2)}$, and $OS_{(3)}$) highlighting the most-outlying values among the other scoring terms for each model. The LZerD shape complementarity scores and the cluster sizes are part of the regular LZerD output.

The order statics terms allow models with any outlying component scores a better chance at inclusion among the top-ranked models, even if some other component scores are low or generally disagree with each other. Each model score for all docked poses of a target is first standardized in the usual way into a Z-score. For example, the mean μ and standard deviation σ of all ITScorePro scores x_i of models of sucrose-specific porin are calculated, and the Z-score of model i will be $z_i = (x_i - \mu)/\sigma$. For scoring terms designed such that more positive values are more favorable, e.g. the LZerD shape complementarity term and cluster sizes, the values were multiplied by -1 so that lower values would uniformly be more favorable across all component scores. Thus, for all the final standardized scoring terms, a more negative Z-score is more favorable. From these, the order statistic terms were calculated for the Z-scores of the scoring terms. To compute e.g. the order statistic $OS(1)$ of a model, we first check the Z-scores of all the scoring terms for the model and select the smallest (the most favorable) Z-score. Following the same logic, $OS(2)$ and $OS(3)$ are the second and the third smallest Z-scores among all the Z-scores of all the scoring terms. A first-order version of this approach was previously used successfully in IDP-LZerD for disordered proteins,⁴⁰ and a third-order version was used successfully in Flex-LZerD for flexible proteins.^{38,39} Here, we use the third-order version. For example, if a model has Z-scores of -1.7 , -2.4 , and -0.5 for three component scores, respectively, and some arbitrary positive-valued Z-scores (i.e. insignificant scores) for the remaining component scores, then $OS(1)$, $OS(2)$, and $OS(3)$ for that model will be -2.4 , -1.7 , and -0.5 , respectively. The three order statistic terms are calculated across all the component scores.

The component scores GOAP, DFIRE, ITScorePro, ranksum, the LZerD shape score, LZerD cluster size, membrane transfer energy, $OS_{(1)}$, $OS_{(2)}$, and $OS_{(3)}$, were then finally

combined in the bagged estimator. Bagging is an ensemble method which constructs multiple decision trees using subsets of the training data, and then averages their individual outputs to produce a more robust classifier.⁷⁹ To optimize the bagged ensemble, docked models were generated for each target in the training dataset using the same procedure as above. CAPRI statistics were then calculated for each model. Each model was labeled positive if it was of acceptable CAPRI quality, and negative if it was not. Each model was then weighted to balance the number of positive and negative cases within each target. Then, the bagged estimator was trained using all the models for all the training targets. The resulting scoring function, with ranksum component scores used to break ties, was used to rescore the docked models after the clustering stage.

Funding

This work used Anvil at Purdue University through allocation CIS230021 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296. This work was partly supported by the National Institutes of Health (R01GM123055, R01GM133840, 3R01 GM133840-02S1) and the National Science Foundation (DMS2151678, DBI2003635, CMMI1825941, MCB2146026, and MCB1925643). CC was supported by NIGMS-funded predoctoral fellowship to CC (T32 GM132024). KH was supported by the Science & Engineering Research Board (SERB), Government of India through the Overseas Visiting Doctoral Fellowship program (OVDF). The contents of this work are solely the responsibility of the authors and do not necessarily represent the official views of the funding agencies.

CRediT authorship contribution statement

Charles Christoffer: Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft.
Kannan Harini: Data curation, Investigation.
Gupta Archit: Data curation, Investigation.
Daisuke Kihara: Conceptualization, Funding acquisition, Methodology, Project administration, Resources, Supervision, Writing – review & editing.

DATA AVAILABILITY

Data will be made available on request.

DECLARATION OF COMPETING 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.

Acknowledgements

The authors thank Daipayan Sarkar for his contribution to an earlier form of the project. The authors are grateful to Information Technology at Purdue, West Lafayette, Indiana for providing computational resources. This publication was made possible with computational resource support from the Purdue Institute of Inflammation, Immunology and Infectious Disease (PI4D).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jmb.2024.168486>.

Received 8 November 2023;
 Accepted 5 February 2024;
 Available online 8 February 2024

Keywords:

protein structure prediction;
 protein–protein docking;
 membrane protein docking;
 membrane protein assembly

References

1. Berman, H.M., Battistuz, T., Bhat, T.N., Bluhm, W.F., Bourne, P.E., Burkhardt, K., et al., (2002). The Protein Data Bank. *Acta Crystallogr. D Biol. Crystallogr.* **58**, 899–907.
2. Berman, H., Henrick, K., Nakamura, H., (2003). Announcing the worldwide Protein Data Bank. *Nature Struct. Biol.* **10**, 980.
3. Peplow, M., (2020). Cryo-electron microscopy reaches resolution milestone. *ACS Cent. Sci.* **6**, 1274–1277.
4. Aderinwale, T., Christoffer, C.W., Sarkar, D., Alhabati, E., Kihara, D., (2020). Computational structure modeling for diverse categories of macromolecular interactions. *Curr. Opin. Struct. Biol.* **64**, 1–8.
5. Mintseris, J., Pierce, B., Wiehe, K., Anderson, R., Chen, R., Weng, Z., (2007). Integrating statistical pair potentials into protein complex prediction. *Proteins* **69**, 511–520.
6. Dominguez, C., Boelens, R., Bonvin, A.M., (2003). HADDOCK: a protein-protein docking approach based on biochemical or biophysical information. *J. Am. Chem. Soc.* **125**, 1731–1737.
7. Kozakov, D., Hall, D.R., Xia, B., Porter, K.A., Padhorny, D., Yueh, C., et al., (2017). The ClusPro web server for protein–protein docking. *Nature Protoc.* **12**, 255–278.

8. Lyskov, S., Gray, J.J., (2008). The RosettaDock server for local protein-protein docking. *Nucleic Acids Res.* **36**, W233–W238.
9. Ritchie, D.W., Venkatraman, V., (2010). Ultra-fast FFT protein docking on graphics processors. *Bioinformatics* **26**, 2398–2405.
10. Torchala, M., Moal, I.H., Chaleil, R.A., Fernandez-Recio, J., Bates, P.A., (2013). SwarmDock: a server for flexible protein-protein docking. *Bioinformatics* **29**, 807–809.
11. de Vries, S., Zacharias, M., (2013). Flexible docking and refinement with a coarse-grained protein model using ATTRACT. *Proteins* **81**, 2167–2174.
12. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., et al., (2021). Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589.
13. Evans, R., O'Neill, M., Pritzel, A., Antropova, N., Senior, A., Green, T., et al., (2021). Protein complex prediction with AlphaFold-Multimer. *bioRxiv*. 2021.10.04.463034.
14. Venkatraman, V., Yang, Y.D., Sael, L., Kihara, D., (2009). Protein-protein docking using region-based 3D Zernike descriptors. *BMC Bioinf.* **10**, 407.
15. Christoffer, C., Terashi, G., Shin, W.H., Aderinwale, T., Maddhuri Venkata Subramaniya, S.R., Peterson, L., et al., (2020). Performance and enhancement of the LZerD protein assembly pipeline in CAPRI 38–46. *Proteins* **88**, 948–961.
16. Christoffer, C., Chen, S., Bharadwaj, V., Aderinwale, T., Kumar, V., Hormati, M., et al., (2021). LZerD webserver for pairwise and multiple protein-protein docking. *Nucleic Acids Res.* **49**, W359–W365.
17. Christoffer, C., Bharadwaj, V., Luu, R., Kihara, D., (2021). LZerD Protein-Protein Docking Webserver Enhanced With de novo Structure Prediction. *Front. Mol. Biosci.* **8**, 724947.
18. Lensink, M.F., Nadzirin, N., Velankar, S., Wodak, S.J., (2020). Modeling protein-protein, protein-peptide, and protein-oligosaccharide complexes: CAPRI. *Proteins* **88**, 916–938.
19. Lensink, M.F., Brysbaert, G., Nadzirin, N., Velankar, S., Chaleil, R.A.G., Gerguri, T., et al., (2019). Blind prediction of homo- and hetero-protein complexes: The CASP13–CAPRI experiment. *Proteins* **87**, 1200–1221.
20. Koukos, P.I., Faro, I., van Noort, C.W., Bonvin, A., (2018). A membrane protein complex docking benchmark. *J. Mol. Biol.* **430**, 5246–5256.
21. Lensink, M.F., Brysbaert, G., Mauri, T., Nadzirin, N., Velankar, S., Chaleil, R.A.G., et al., (2021). Prediction of protein assemblies, the next frontier: The CASP14–CAPRI experiment. *Proteins* **89**, 1800–1823.
22. Kolbe, M., Besir, H., Essen, L.O., Oesterhelt, D., (2000). Structure of the light-driven chloride pump halorhodopsin at 1.8 Å resolution. *Science* **288**, 1390–1396.
23. Kouryama, T., Kanada, S., Takeguchi, Y., Narusawa, A., Murakami, M., Ihara, K., (2010). Crystal structure of the light-driven chloride pump halorhodopsin from *Natronomonas pharaonis*. *J. Mol. Biol.* **396**, 564–579.
24. Turnbull, A.P., Rafferty, J.B., Sedelnikova, S.E., Slabas, A. R., Schierer, T.P., Kroon, J.T., et al., (2001). Analysis of the structure, substrate specificity, and mechanism of squash glycerol-3-phosphate (1)-acyltransferase. *Structure* **9**, 347–353.
25. Gulsevin, A., Meiler, J., (2021). Prediction of amphipathic helix-membrane interactions with Rosetta. *PLoS Comput. Biol.* **17**, e1008818.
26. Fuglebakk, E., Reuter, N., (2018). A model for hydrophobic protrusions on peripheral membrane proteins. *PLoS Comput. Biol.* **14**, e1006325.
27. Sutherland, M.C., Jarodsky, J.M., Ovchinnikov, S., Baker, D., Kranz, R.G., (2018). Structurally mapping endogenous heme in the CcmCDE membrane complex for cytochrome c biogenesis. *J. Mol. Biol.* **430**, 1065–1080.
28. Gao, M., Nakajima, A.D., Parks, J.M., Skolnick, J., (2022). AF2Complex predicts direct physical interactions in multimeric proteins with deep learning. *Nature Commun.* **13**, 1744.
29. Wu, R., Bakelar, J.W., Lundquist, K., Zhang, Z., Kuo, K.M., Ryoo, D., et al., (2021). Plasticity within the barrel domain of BamA mediates a hybrid-barrel mechanism by BAM. *Nature Commun.* **12**, 7131.
30. Kaczor, A.A., Jorg, M., Capuano, B., (2016). The dopamine D2 receptor dimer and its interaction with homobivalent antagonists: homology modeling, docking and molecular dynamics. *J. Mol. Model.* **22**, 203.
31. Kaczor, A.A., Guixa-Gonzalez, R., Carrio, P., Poso, A., Dove, S., Pastor, M., et al., (2015). Multi-component protein – protein docking based protocol with external scoring for modeling dimers of G protein-coupled receptors. *Mol. Inform.* **34**, 246–255.
32. Kaczor, A.A., Guixa-Gonzalez, R., Carrio, P., Obiol-Pardo, C., Pastor, M., Selent, J., (2012). Fractal dimension as a measure of surface roughness of G protein-coupled receptors: implications for structure and function. *J. Mol. Model.* **18**, 4465–4475.
33. Altuwaijry, N.A., Baron, M., Wright, D.W., Coveney, P.V., Townsend-Nicholson, A., (2017). An ensemble-based protocol for the computational prediction of helix-helix interactions in G protein-coupled receptors using coarse-grained molecular dynamics. *J. Chem. Theory Comput.* **13**, 2254–2270.
34. Alford, R.F., Koehler, L.J., Weitzner, B.D., Duran, A.M., Tilley, D.C., Elazar, A., et al., (2015). An integrated framework advancing membrane protein modeling and design. *PLoS Comput. Biol.* **11**, e1004398.
35. Hurwitz, N., Wolfson, H.J., (2021). Memdock: an alpha-helical membrane protein docking algorithm. *Methods Mol. Biol.* **2315**, 111–117.
36. Hurwitz, N., Schneidman-Duhovny, D., Wolfson, H.J., (2016). Memdock: an alpha-helical membrane protein docking algorithm. *Bioinformatics* **32**, 2444–2450.
37. Rudden, L.S.P., Degiacomi, M.T., (2021). Transmembrane protein docking with JabberDock. *J. Chem. Inf. Model.* **61**, 1493–1499.
38. Christoffer, C., Kihara, D., (2022). Domain-based protein docking with extremely large conformational changes. *J. Mol. Biol.* **434**, 167820.
39. Christoffer, C., Kihara, D., (2022). Modeling protein-nucleic acid complexes with extremely large conformational changes using Flex-LZerD. *Proteomics* e2200322.
40. Peterson, L.X., Roy, A., Christoffer, C., Terashi, G., Kihara, D., (2017). Modeling disordered protein interactions from biophysical principles. *PLoS Comput. Biol.* **13**, e1005485.
41. Wolfson, H.J., Rigoutsos, I., (1997). Geometric hashing: an overview. *IEEE Comput. Sci. Eng.* **4**, 10–21.
42. Canterakis, N., (1999). 3D Zernike moments and Zernike affine invariants for 3D image analysis and recognition. *In 11th Scandinavian Conf on Image Analysis: Citeseer*.

43. Aderinwale, T., Bharadwaj, V., Christoffer, C., Terashi, G., Zhang, Z., Jahandideh, R., et al., (2022). Real-time structure search and structure classification for AlphaFold protein models. *Commun Biol.* **5**, 316.

44. Shin, W.H., Christoffer, C.W., Wang, J., Kihara, D., (2016). PL-PatchSurfer2: improved local surface matching-based virtual screening method that is tolerant to target and ligand structure variation. *J. Chem. Inf. Model.* **56**, 1676–1691.

45. Esquivel-Rodriguez, J., Xiong, Y., Han, X., Guang, S., Christoffer, C., Kihara, D., (2015). Navigating 3D electron microscopy maps with EM-SURFER. *BMC Bioinf.* **16**, 181.

46. Xiong, Y., Esquivel-Rodriguez, J., Sael, L., Kihara, D., (2014). 3D-SURFER 2.0: web platform for real-time search and characterization of protein surfaces. *Methods Mol. Biol.* **1137**, 105–117.

47. La, D., Esquivel-Rodriguez, J., Venkatraman, V., Li, B., Sael, L., Ueng, S., et al., (2009). 3D-SURFER: software for high-throughput protein surface comparison and analysis. *Bioinformatics* **25**, 2843–2844.

48. Lomize, A.L., Todd, S.C., Pogozheva, I.D., (2022). Spatial arrangement of proteins in planar and curved membranes by PPM 3.0. *Protein Sci.* **31**, 209–220.

49. Lomize, M.A., Pogozheva, I.D., Joo, H., Mosberg, H.I., Lomize, A.L., (2012). OPM database and PPM web server: resources for positioning of proteins in membranes. *Nucleic Acids Res.* **40**, D370–D376.

50. Lee, J., Cheng, X., Swails, J.M., Yeom, M.S., Eastman, P. K., Lemkul, J.A., et al., (2016). CHARMM-GUI input generator for NAMD, GROMACS, AMBER, OpenMM, and CHARMM/OpenMM simulations using the CHARMM36 additive force field. *J. Chem. Theory Comput.* **12**, 405–413.

51. Lensink, M.F., Velankar, S., Baek, M., Heo, L., Seok, C., Wodak, S.J., (2018). The challenge of modeling protein assemblies: the CASP12-CAPRI experiment. *Proteins* **86**, 257–273.

52. Lensink, M.F., Velankar, S., Wodak, S.J., (2017). Modeling protein-protein and protein-peptide complexes: CAPRI. *Proteins* **85**, 359–377.

53. Feng, L., Yan, H., Wu, Z., Yan, N., Wang, Z., Jeffrey, P.D., et al., (2007). Structure of a site-2 protease family intramembrane metalloprotease. *Science* **318**, 1608–1612.

54. Brown, M.S., Ye, J., Rawson, R.B., Goldstein, J.L., (2000). Regulated intramembrane proteolysis: a control mechanism conserved from bacteria to humans. *Cell* **100**, 391–398.

55. Nie, L., Pascoa, T.C., Pike, A.C.W., Bushell, S.R., Quigley, A., Ruda, G.F., et al., (2021). The structural basis of fatty acid elongation by the ELOVL elongases. *Nature Struct. Mol. Biol.* **28**, 512–520.

56. Zhang, K., Kniazeva, M., Han, M., Li, W., Yu, Z., Yang, Z., et al., (2001). A 5-bp deletion in ELOVL4 is associated with two related forms of autosomal dominant macular dystrophy. *Nature Genet.* **27**, 89–93.

57. Cadieux-Dion, M., Turcotte-Gauthier, M., Noreau, A., Martin, C., Meloche, C., Gravel, M., et al., (2014). Expanding the clinical phenotype associated with ELOVL4 mutation: study of a large French-Canadian family with autosomal dominant spinocerebellar ataxia and erythrokeratoderma. *JAMA Neurol.* **71**, 470–475.

58. Di Gregorio, E., Borroni, B., Giorgio, E., Lacerenza, D., Ferrero, M., Lo, B.N., et al., (2014). ELOVL5 mutations cause spinocerebellar ataxia 38. *Am. J. Hum. Genet.* **95**, 209–217.

59. Zimmermann, L., Stephens, A., Nam, S.Z., Rau, D., Kubler, J., Lozajic, M., et al., (2018). A completely reimplemented MPI bioinformatics toolkit with a new HHpred server at its core. *J. Mol. Biol.* **430**, 2237–2243.

60. Svensson-Ek, M., Abramson, J., Larsson, G., Tornroth, S., Brzezinski, P., Iwata, S., (2002). The X-ray crystal structures of wild-type and EQ(I-286) mutant cytochrome c oxidases from Rhodobacter sphaeroides. *J. Mol. Biol.* **321**, 329–339.

61. Spencer, A.G., Woods, J.W., Arakawa, T., Singer, I.I., Smith, W.L., (1998). Subcellular localization of prostaglandin endoperoxide H synthases-1 and -2 by immunoelectron microscopy. *J. Biol. Chem.* **273**, 9886–9893.

62. Duggan, K.C., Walters, M.J., Musee, J., Harp, J.M., Kiefer, J.R., Oates, J.A., et al., (2010). Molecular basis for cyclooxygenase inhibition by the non-steroidal anti-inflammatory drug naproxen. *J. Biol. Chem.* **285**, 34950–34959.

63. Hogestatt, E.D., Jonsson, B.A., Ermund, A., Andersson, D. A., Bjork, H., Alexander, J.P., et al., (2005). Conversion of acetaminophen to the bioactive N-acylphenolamine AM404 via fatty acid amide hydrolase-dependent arachidonic acid conjugation in the nervous system. *J. Biol. Chem.* **280**, 31405–31412.

64. Rowlinson, S.W., Kiefer, J.R., Prusakiewicz, J.J., Pawlitz, J.L., Kozak, K.R., Kalgutkar, A.S., et al., (2003). A novel mechanism of cyclooxygenase-2 inhibition involving interactions with Ser-530 and Tyr-385. *J. Biol. Chem.* **278**, 45763–45769.

65. Kurumbail, R.G., Stevens, A.M., Gierse, J.K., McDonald, J. J., Stegeman, R.A., Pak, J.Y., et al., (1996). Structural basis for selective inhibition of cyclooxygenase-2 by anti-inflammatory agents. *Nature* **384**, 644–648.

66. Eydoux, C., De Caro, J., Ferrato, F., Boullanger, P., Lafont, D., Laugier, R., et al., (2007). Further biochemical characterization of human pancreatic lipase-related protein 2 expressed in yeast cells. *J. Lipid Res.* **48**, 1539–1549.

67. van Bennekom, A.M., Fisher, E.A., Blaner, W.S., Harrison, E.H., (2000). Hydrolysis of retinyl esters by pancreatic triglyceride lipase. *Biochemistry* **39**, 4900–4906.

68. van Tilburgh, H., Egloff, M.P., Martinez, C., Rugani, N., Verger, R., Cambillau, C., (1993). Interfacial activation of the lipase-procolipase complex by mixed micelles revealed by X-ray crystallography. *Nature* **362**, 814–820.

69. Kozakov, D., Li, K., Hall, D.R., Beglov, D., Zheng, J., Vakili, P., et al., (2014). Encounter complexes and dimensionality reduction in protein-protein association. *Elife* **3**, e01370.

70. Vakser, I.A., Grudinin, S., Jenkins, N.W., Kundrotas, P.J., Deeds, E.J., (2022). Docking-based long timescale simulation of cell-size protein systems at atomic resolution. *Proc. Natl. Acad. Sci. USA* **119**, e2210249119.

71. Zhang, Y., Zhang, Z., Kagaya, Y., Terashi, G., Zhao, B., Xiong, Y., et al., (2023). Distance-AF: modifying predicted protein structure models by alphafold2 with user-specified distance constraints. *bioRxiv*.

72. Steinegger, M., Soding, J., (2018). Clustering huge protein sequence sets in linear time. *Nature Commun.* **9**, 2542.

73. Mukherjee, S., Zhang, Y., (2009). MM-align: a quick algorithm for aligning multiple-chain protein complex

structures using iterative dynamic programming. *Nucleic Acids Res.* **37**, e83.

74. Esquivel-Rodriguez, J., Filos-Gonzalez, V., Li, B., Kihara, D., (2014). Pairwise and multimeric protein-protein docking using the LZerD program suite. *Methods Mol. Biol.* **1137**, 209–234.

75. Peterson, L.X., Kim, H., Esquivel-Rodriguez, J., Roy, A., Han, X., Shin, W.H., et al., (2017). Human and server docking prediction for CAPRI round 30–35 using LZerD with combined scoring functions. *Proteins* **85**, 513–527.

76. Peterson, L.X., Shin, W.H., Kim, H., Kihara, D., (2018). Improved performance in CAPRI round 37 using LZerD docking and template-based modeling with combined scoring functions. *Proteins* **86**, 311–320.

77. Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.Y., (1998). An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM* **45**, 891–923.

78. Edelsbrunner, H., Seidel, R., (1986). Voronoi diagrams and arrangements. *Discrete & Computational Geometry* **1**, 25–44.

79. Breiman, L., (1996). Bagging predictors. *Mach. Learn.* **24**, 123–140.

80. Zhou, H., Skolnick, J., (2011). GOAP: a generalized orientation-dependent, all-atom statistical potential for protein structure prediction. *Biophys. J.* **101**, 2043–2052.

81. Zhou, H., Zhou, Y., (2002). Distance-scaled, finite ideal-gas reference state improves structure-derived potentials of mean force for structure selection and stability prediction. *Protein Sci.* **11**, 2714–2726.

82. Huang, S.Y., Zou, X., (2011). Statistical mechanics-based method to extract atomic distance-dependent potentials from protein structures. *Proteins* **79**, 2648–2661.

83. MacCallum, J.L., Bennett, W.F., Tielemans, D.P., (2007). Partitioning of amino acid side chains into lipid bilayers: results from computer simulations and comparison to experiment. *J. Gen. Physiol.* **129**, 371–377.