

PAPER

PMIpred: A physics-informed web server for quantitative Protein-Membrane Interaction prediction

Niek van Hilten^{1,†}, Nino Verwei^{1,†}, Jeroen Methorst¹, Carsten Nase²,
Andrius Bernatavicius^{3,4} and Herre Jelger Risselada^{1,2,*}

¹Leiden Institute of Chemistry, Leiden University, Einsteinweg 55, 2333 CC, Leiden Netherlands, ²Department of Physics, Technical University Dortmund, Otto-Hahn-Strasse 4, 44227, Dortmund, Germany, ³Leiden Institute of Advanced Computer Science, Leiden University, Niels Bohrweg 1, 2333 CA, Leiden, Netherlands and ⁴Leiden Academic Centre for Drug Research, Leiden University, Einsteinweg 55, 2333 CC, Leiden, Netherlands

*Corresponding author. jelger.risselada@tu-dortmund.de

†These authors contributed equally.

FOR PUBLISHER ONLY Received on Date Month Year; revised on Date Month Year; accepted on Date Month Year

Abstract

Motivation: Many membrane peripheral proteins have evolved to transiently interact with the surface of (curved) lipid bilayers. Currently, methods to quantitatively predict sensing and binding free energies for protein sequences or structures are lacking, and such tools could greatly benefit the discovery of membrane-interacting motifs, as well as their *de novo* design.

Results: Here, we trained a transformer neural network model on molecular dynamics data for >50,000 peptides that is able to accurately predict the (relative) membrane-binding free energy for any given amino acid sequence. Using this information, our physics-informed model is able to classify a peptide's membrane-associative activity as either non-binding, curvature sensing, or membrane binding. Moreover, this method can be applied to detect membrane-interaction regions in a wide variety of proteins, with comparable predictive performance as state-of-the-art data-driven tools like DREAMM, PPM3, and MODA, but with a wider applicability regarding protein diversity, and the added feature to distinguish curvature sensing from general membrane binding.

Availability: We made these tools available as a web server, coined Protein-Membrane Interaction predictor (PMIpred), which can be accessed at <https://pmipred.fkt.physik.tu-dortmund.de>.

Key words: curvature sensing, membrane binding, peripheral membrane proteins, web server

Introduction

Peripheral membrane proteins (PMPs) are a class of proteins that transiently adhere to the surface of the biological membranes that encapsulate and compartmentalise cells. These protein-membrane interactions are crucial for the protein's subcellular localisation and related function, e.g. in membrane remodelling, transport, or lipid metabolism (Monje-Galvan and Klauda [2016], Whited and Johs [2015]).

The membrane partitioning of PMPs is often regulated by distinct protein regions that have evolved to optimally guide these proteins to their functional locations. One striking example is the amphipathic lipid packing sensing (ALPS) motif that facilitates specific binding to curved membranes and not to flat ones. ALPS was first described as part of the ArfGAP1 protein that regulates COPI coat assembly in a curvature-specific manner (Bigay et al. [2003]). With this original ALPS motif as a template, Drin *et al* defined a set of physicochemical criteria (e.g. prevalence of certain residues, net charge, hydrophobicity, hydrophobic moment) that allowed

for a bioinformatic screening of protein databases, which accelerated the discovery of many other curvature-sensing proteins (Drin et al. [2007]).

Despite this breakthrough, there are many membrane-active protein motifs that *do* display similar curvature-related behaviour, but that *do not* fulfill Drin's criteria. For example, the AH repeats of α -synuclein specifically bind to curved anionic membranes (Davidson et al. [1998]), but strongly differ from ALPS in their amino acid (AA) composition (Pranke et al. [2011]). Hence, it is suggested that curvature sensing is governed by a subtle balance between the chemical properties of the hydrophobic and the hydrophilic faces (Drin and Antonny [2010]) that is hard to generalise using simple physicochemical descriptors.

The example of curvature sensing shows that it is not trivial to predict which regions of a given protein structure are responsible for membrane interactions; let alone distinguishing curvature specificity at the same time. Previously developed methods include data-informed classifiers that predict which

1 AAs in a protein structure are membrane-interacting (e.g.
 2 MODA (Kufareva et al. [2014]) and DREAMM (Chatzigeorgoulas
 3 and Cournia [2022b,a])) or describe the general protein
 4 orientation with respect to a (curved) membrane (e.g. PPM3
 5 (Lomize et al. [2022])). Although these methods perform well
 6 for their respective protein targets, they lack the quantitative
 7 nature that is required to detect the more subtle interactions
 8 involved in curvature sensing.

9 In this paper, we present an alternative method that is
 10 able to predict protein-membrane interactions for peptides and
 11 protein structures from a bottom-up physics-based perspective.
 12 Our approach involves a transformer neural network (NN)
 13 model that is trained on a large set of molecular dynamics (MD)
 14 data on relative binding free energies for short peptides (24
 15 residues) interacting with model membranes. Since our data
 16 were initially gathered during the optimisation of curvature
 17 sensing (van Hilten et al. [2023]), our model is additionally able
 18 to distinguish curvature- (or, equivalently, lipid packing defect-
 19) sensing from general membrane binding. Moreover, because
 20 of its physics-informed character, it is – compared to previous
 21 data-driven methods – better able to generalise membrane-
 22 interaction features across a wide variety of protein families,
 23 including lipid kinases, ALPS-containing curvature sensors,
 24 α -synuclein, and N-BAR proteins. To enable researchers to
 25 access our tool without requiring any programming skills,
 26 we incorporated it in our Protein-Membrane Interaction
 27 predictor (“PMIpred”) web server, that is available at
 28 <https://pmipred.fkt.physik.tu-dortmund.de>.

29 Approach

Theoretical background

31 In previous work, we developed a free-energy calculation
 32 method to efficiently quantify curvature-sensing free energy
 33 ($\Delta\Delta F$) for a peptide within (coarse-grained) MD simulations
 34 (van Hilten et al. [2022]). $\Delta\Delta F$ is the difference in affinity
 35 for a curved/stretched membrane versus a flat equilibrium
 36 membrane. Since binding to the lipid packing defects at curved
 37 membrane surfaces negates an energetic penalty, $\Delta\Delta F$ has a
 38 negative value for most peptides. The larger the magnitude
 39 ($\Delta\Delta F \ll 0$), the stronger the sensing propensity.

40 Because $\Delta\Delta F$ is linearly related to the sequence length
 41 L , we can extrapolate it to our maximal reference length of
 42 24 AAs, which enables fair comparison between peptides of
 43 different lengths:

$$44 \Delta\Delta F_{L24} = \frac{a_L \cdot 24 + b_L}{a_L \cdot L + b_L} \Delta\Delta F \quad (1)$$

45 In which $a_L = -1.03$ and $b_L = 3.28$. Next, to account for the
 46 fact that the majority of biomembranes is negatively charged
 47 while the membranes in our simulations were neutral, we can
 48 add another correction:

$$49 \Delta\Delta F_{\text{adj}} = \Delta\Delta F_{L24} + c_z z \quad (2)$$

50 With $c_z = -0.93 \text{ kJ mol}^{-1}$, meaning that adding a positive
 51 charge z reduces $\Delta\Delta F$ by 0.93 kJ mol^{-1} and, conversely,
 52 adding a negative charge increases it by 0.93 kJ mol^{-1} .

53 By combining this approach with an evolutionary algorithm
 54 (EA) to efficiently navigate the vast search space (20^{24} possible
 55 combinations), we found that the optimal curvature sensor is
 56 hydrophobic and bulky, indicating that curvature sensing $\Delta\Delta F$
 57 and general membrane binding $\Delta F_{\text{sm}}(R)$ – for binding to a
 58 vesicle with radius R – are strongly correlated:

$$59 \Delta F_{\text{sm}}(R) = \underbrace{\alpha \cdot \Delta\Delta F_{\text{adj}}}_{\text{Flat membrane}} + \beta + \underbrace{\frac{\Delta\Delta F_{\text{adj}}}{\epsilon_{\text{ref}}}}_{\text{Curvature correction}} \left(\frac{1}{R^2} + \frac{2}{R} \right) \quad (3)$$

60 By fitting our data ($\Delta\Delta F$) to thermodynamic integration
 61 simulations that yield ΔF_{sm} , we found that $\alpha = 3.83$ and $\beta =$
 62 12.27 . The relative strain ϵ_{ref} of the target membrane in our
 63 MD simulations equals 0.165 . The estimation of ΔF_{sm} could
 64 be alternatively fine-tuned in future generations of the model
 65 via incorporation of additional empirically derived algorithmic
 66 rules, such as ref. Hristova and White [2005], in order to correct
 67 for systematic force field errors.

68 Finally, we theoretically derived that the membrane-binding
 69 probability P_m relates to this ΔF_{sm} according to a sigmoid-like
 70 switch function:

$$71 P_m = \frac{1}{1 + \frac{N_s}{N_m} e^{\Delta F_{\text{sm}}}} \quad (4)$$

72 In which the ratio between the number of realisations in
 73 solvent (N_s) and on the membrane (N_m) is estimated from
 74 typical experimental setups to be $\frac{N_s}{N_m} = 5.2 \cdot 10^3$. This
 75 curve, plotted in Fig. 1A, captures three distinct regimes: non-
 76 binders (peptides stay in solution, $P_m \rightarrow 0$), curvature sensors
 77 (peptides only bind to curved membranes, $P_m \approx 0.5$), and
 78 binders (peptides bind to any membrane, $P_m \rightarrow 1$). Details
 79 on the derivations of eq. 1-4 and their constants can be found
 80 in our previous publication (van Hilten et al. [2023]).

Transformer model

81 The randomly initiated evolutionary optimisation of curvature-
 82 sensing peptides we performed previously (van Hilten et al.
 83 [2023]) yielded a large data set of 53,940 unique peptide
 84 sequences (all 24 residues long) and their respective $\Delta\Delta F$
 85 values, as calculated by coarse-grained MD simulations (see S1).
 86 An advantage of such a physics-based data set is that it spans
 87 the entire applicability domain from ‘inactive’ ($\Delta\Delta F = 0.0 \text{ kJ}$
 88 mol^{-1}) to the theoretical optimum ($\Delta\Delta F \approx -30.0 \text{ kJ mol}^{-1}$).
 89 Previously, we showed that a convolutional neural network
 90 (CNN) trained on these data is capable of accurately predicting
 91 a peptide’s $\Delta\Delta F$ value and classifying its membrane-binding
 92 behaviour according to the three aforementioned categories
 93 (non-binder, sensor, binder, see Fig. 1A) (van Hilten et al.
 94 [2023]). Moreover, it did so across different peptide families
 95 and with point-mutation sensitivity.

96 In the current work, we follow-up on this by employing a
 97 transformer model to the same problem. Transformer models
 98 use a so-called ‘attention mechanism’ to learn patterns from
 99 text input, making them especially useful in natural language
 100 processing (Vaswani et al. [2017]). Transformer NNs are able
 101 to process whole sentences (or in our case, peptide sequences)
 102 at once, which gives them two advantages compared to CNNs
 103 and long short-term memory (LSTM) models: (1) they can
 104 process input faster, and (2) they are better able to learn long
 105 range relations between elements, which is important when –
 106 for example – the first and last word of a long input sentence
 107 are conceptually and/or grammatically connected.

108 All details on the transformer architecture, optimisation,
 109 and training are described in the supplementary materials,
 110 section S1.

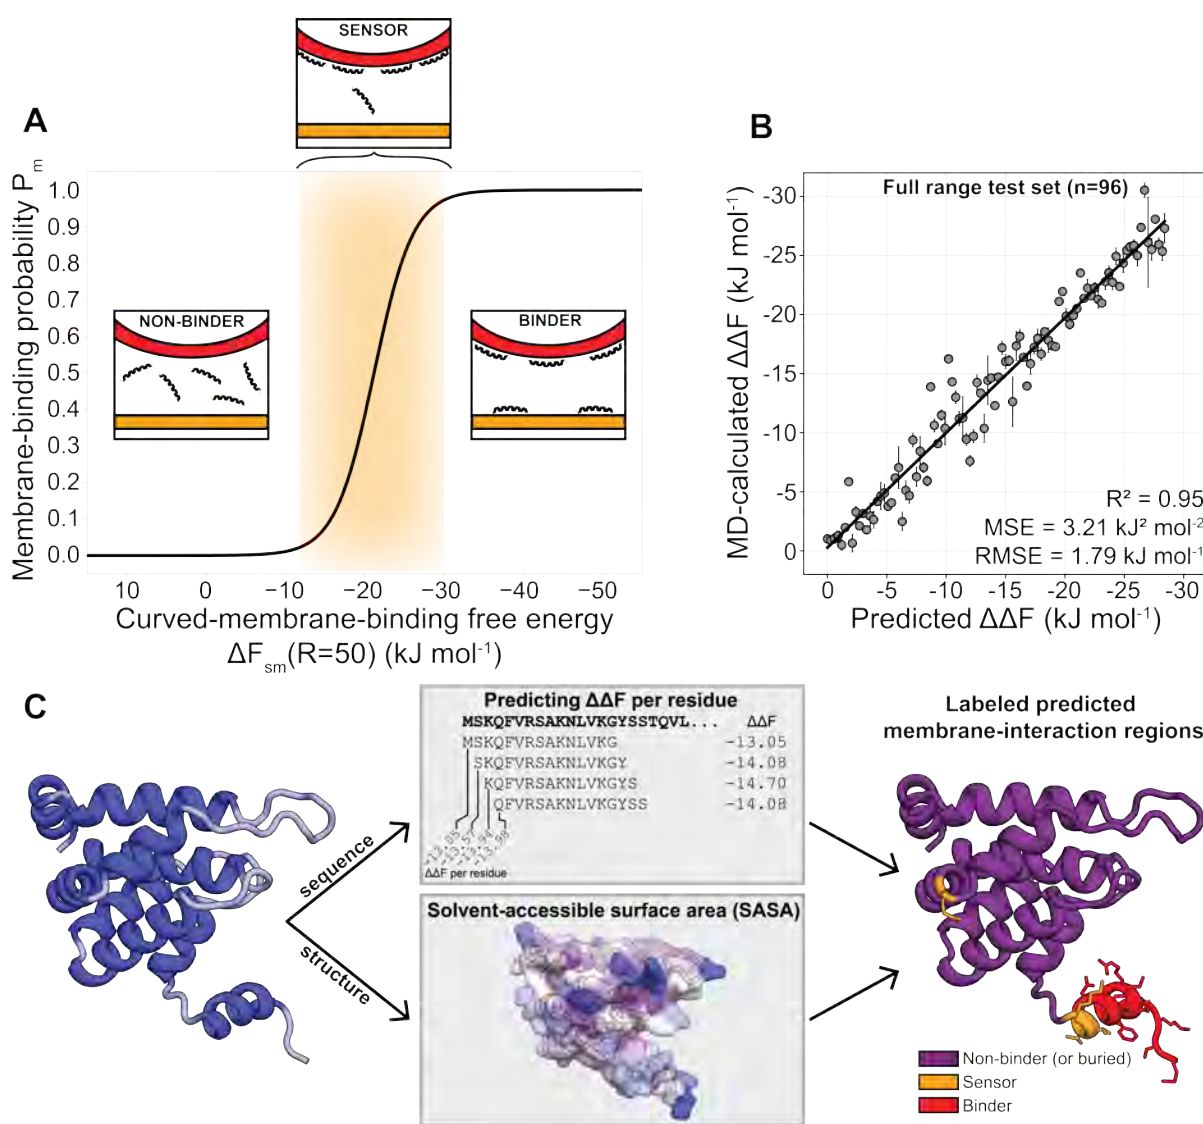


Fig. 1. **A)** The membrane-binding probability P_m as a function of the membrane-binding free energy ΔF_{sm} (at $R = 50$ nm) shows a sharp transition. Please note the distinction between ΔF_{sm} and the *relative curvature-sensing free energy* $\Delta\Delta F$, which are linearly correlated (eq. 3 and Fig. 4C in our previous work (van Hilten et al. [2023])). The orange area marks the 'sensor regime' (van Hilten et al. [2023]). Insets show cartoon explanations of the three classes. **B)** Correlation plot for 96 sequences (all 24 residues long, not part of the training data) spanning the full $\Delta\Delta F$ -range. MD-calculated values are averages and standard deviations for 3 independent replicas (500 ns simulation per run, as described in ref. van Hilten et al. [2023]). **C)** Schematic workflow of predicting membrane-interaction regions in protein structures within PMIPred. The trained transformer model predicts $\Delta\Delta F$ (in kJ mol^{-1}) for overlapping peptide segments from the protein sequence. From the 3D protein structure, per-residue SASA values are calculated. Finally, those two features are combined to yield the final quantitative prediction, as visualised by PMIPred. Non-binding and buried residues are shown in purple. Sensing residues are shown in orange. Binding residues are shown in red. The protein structure shown here is the espin1 ENTH domain (PDB: 5onf).

Our final transformer model is able to predict the $\Delta\Delta F$ calculated from MD simulations with an expected error (root-mean square deviation) of 1.79 kJ mol^{-1} , across the full free-energy range (Fig. 1B), which is similar to the typical standard deviations we get for the MD simulations in the training data.

Identifying membrane-interaction regions in protein structures

In order to utilise our method to detect membrane-interaction regions in PMPs, we apply a sliding window that divides the full

protein sequence into equally sized segments and predicts $\Delta\Delta F$ values for each of them. By default, we consider segments of 15 residues, which is a typical length for a secondary structure element in a protein. A 'per-residue' $\Delta\Delta F$ score is obtained by taking the average of the overlapping segment scores at each position (Fig. 1C). These per-residue values can be interpreted as the contribution of that single amino acid to the overall membrane-binding ability of the protein region it is in. When classifying residues as non-binders, sensors, or binders, we use the $\Delta\Delta F$ -thresholds that follow from the definitions in eq. 1-4:

$\Delta\Delta F > -6.4 \text{ kJ mol}^{-1}$ for non-binders; $-10.0 \leq \Delta\Delta F \leq -6.4 \text{ kJ mol}^{-1}$ for sensors, and $\Delta\Delta F < -10.0 \text{ kJ mol}^{-1}$ for binders.

Next, it is important to consider whether a residue is at the surface of the protein structure (e.g. from the PDB or AlphaFold (AF)), since this is a prerequisite for it to engage in membrane-binding interactions. Our method accounts for this by calculating the solvent-accessible surface area (SASA) for every individual residue from the protein structure (using the Shrake-Rupley algorithm (Shrake and Rupley [1973]) as implemented in BioPython (Cock et al. [2009]), see Fig. 1C), and then taking a local average of the 9-residue vicinity n_{-4} - n_{+4} with the resulting SASA value being assigned to residue n .

For intuitive visualisation, the predicted sensing/binding activity can be mapped and colour-coded onto the 3D protein structure (right panel in Fig. 1C). To this end, we use the B-factor field in the output PDB-file that is commonly used for applying colour schemes in molecular visualisation software (e.g. PyMOL (Schrodinger, LLC [2015])). Predicted non-binding (or inaccessible), sensing, and binding regions get assigned B-factors of 0.0, 0.5, and 1.0, respectively.

Results and Discussion

Predicting membrane interaction for peptides

To apply the trained transformer model on individual peptide sequences, PMIPred allows users to input a string of 7-24 natural AAs (single-letter abbreviations). Also, they need to specify if their target membrane is neutral (e.g. pure POPC) or negatively charged (as is common in nature), see Fig. 2A. In case of the latter, eq. 2 is applied. Upon submission, the web server will produce:

1. A helical wheel diagram of the input sequence, generated by the modLAMP python module (Müller et al. [2017]) (Fig. 2B). This diagram visualises the spatial orientation of the AA side chains once folded into an α -helix. This can be especially useful to, for example, identify a clustering of hydrophobic residues at one side of the helix, which increases the peptide's amphipathicity and generally favours interactions with lipid bilayer surfaces. Please note that PMIPred – just like the MD simulations that the model is trained on – always assumes an α -helical folding for input sequences. Additional MD data on helical and unstructured peptides suggest that membrane-binding free energies are dominated by the general AA composition, and not secondary structure (see S2). Nevertheless, if an input peptide is known to have a non-helical fold (and, potentially, related function), the PMIPred results should be interpreted with caution.
2. Predicted free energies and several commonly used physicochemical descriptors (Fig. 2C). From top to bottom: (1) The NN-predicted curvature-sensing free energy ($\Delta\Delta F$). The greater the magnitude of $|\Delta\Delta F|$, the higher the preference for the peptide to bind to (strongly curved) membranes. (2) The length-adjusted $\Delta\Delta F_{L24}$ (eq. 1), that enables fair comparison for peptides with different lengths. (3) The charge-adjusted $\Delta\Delta F_{adj}$ (eq. 2) that additionally corrects for charge-sensing effects on negatively charged membranes. (4) $\Delta F_{sm}(R = 50)$, i.e. the predicted membrane-binding free energy to vesicles with a 50 nm radius (eq. 3). Depending on the choice of target membrane (neutral or negatively charged), $\Delta\Delta F_{L24}$ or $\Delta\Delta F_{adj}$ is used for the calculations, respectively. Finally, (5) – and much like tools such as HeliQuest (Gautier et al. [2008]) – PMIPred computes some physicochemical descriptors: sequence length, charge, mean hydrophobicity (using Fauchère & Pliška's hydrophobicity scale (Fauchère and Pliška [1983])), and Eisenberg's hydrophobic moment (Eisenberg et al. [1982]).
3. The membrane-binding probability P_m as a function of the curved-membrane-binding free energy ΔF_{sm} at a 50 nm vesicle radius, as also plotted in Fig. 1A. PMIPred depicts the position of the query sequence on this curve (Fig. 2D).
4. The membrane-binding probability P_m as a function of vesicle radius R (and thus membrane curvature R^{-1}), see Fig. 2E. Peptides that fall in the purple area of this plot are predicted to be non-binders, that prefer to stay in solution regardless of the curvature of available membranes ($P_m \rightarrow 0$). Conversely, the red zone would represent peptides that bind to any membrane, regardless of curvature ($P_m \rightarrow 1$). The orange transition zone is the sensing regime, where the binding probability strongly depends on the membrane curvature. In this plot, a peptide (and its associated $\Delta\Delta F$ value) is represented by a curve, as demonstrated for our example peptide.

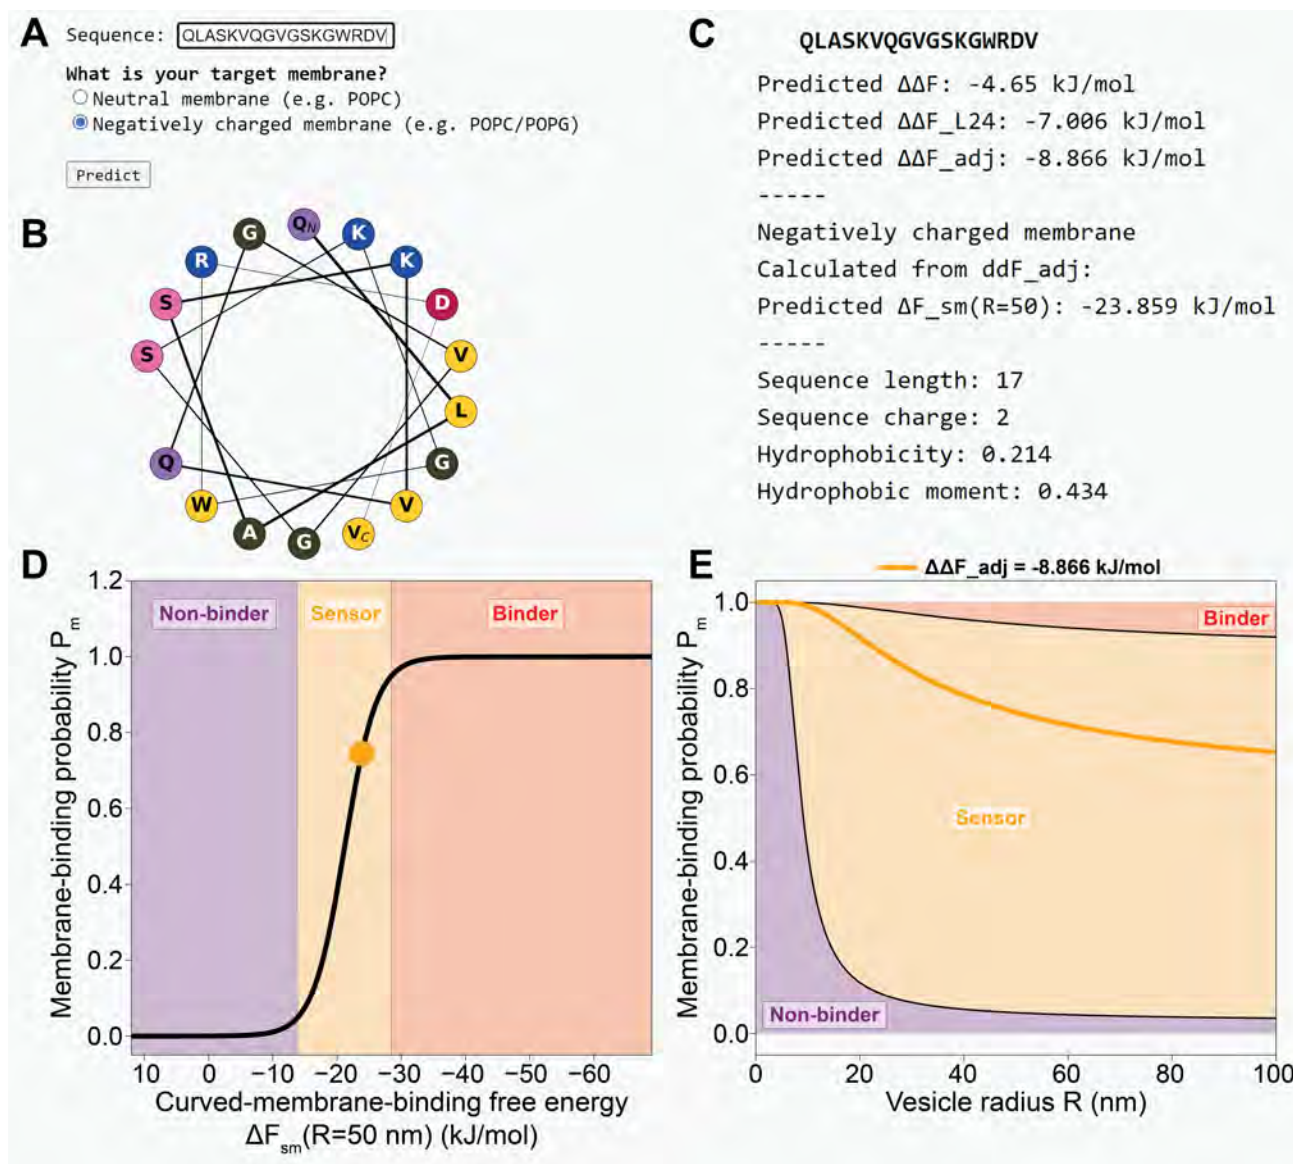


Fig. 2. **A**) Users can input a peptide sequence (part of ArfGAP1's ALPS2 motif (Mesmin et al. [2007]), in this example) and chose their desired target membrane. **B**) Helical wheel diagram, generated by modIAMP (Müller et al. [2017]). **C**) Predicted free-energy values and physicochemical characteristics. **D**) The position of the input sequence on the non-binding \rightarrow binding continuum. Predicted sensors fall in the orange transition zone. **E**) Membrane-binding probability P_m as a function of vesicle radius R . Again, sensors fall in the orange area.

Screening proteins for membrane-interaction motifs

As outlined previously, PMIPred can also be applied to predict which regions in PMPs take part in curvature-sensing or membrane-binding behaviour. To demonstrate this feature, we highlight three example protein structures from different protein families.

ArfGAP1

First, we examined ArfGAP1 (Bigay et al. [2003]), which comprises a structured Zn-finger domain to control GTP hydrolysis in Arf1 on the Golgi membrane (Cukierman et al. [1995]). This catalytic domain is followed by a disordered region that contains two ALPS motifs (Mesmin et al. [2007]), which fold into α -helices upon binding to curved membranes. Hereby, they couple the COPI-induced curvature to the

catalytic activity and control the transport vesicle budding process. In absence of a PDB entry for the full ArfGAP1 protein, we downloaded the predicted structure from the AlphaFold (AF) database (Jumper et al. [2021], Varadi et al. [2022]), which accurately models the distinct structured and unstructured domains (although for the latter, the relative orientation between the flexible loops is likely wrong). We applied our protein screening tool to ArfGAP1 and found that it correctly identifies the two curvature-sensing motifs ALPS1 and ALPS2 (Fig. 3A). Parts of both ALPS motifs were labeled as 'binders', although they only slightly surpassed the sensing \rightarrow binding threshold at $\Delta\Delta F = -10.0$ for most positions (see data repository at <https://github.com/nvanhiltten/PMIPred>). Interestingly, our model also identified a third region (residue 134-151, directly adjacent to the catalytic Zn-finger domain) to potentially contribute to (curvature-specific) membrane

interactions. To the best of our knowledge, this segment is not described to be involved in membrane interaction in the current literature, which could mean that either this is a false positive identification, or it is an actual undiscovered membrane-interacting motif that would be interesting to examine further.

Epsin1 ENTH

The second protein we studied is the N-terminal ENTH domain from epsin1, which is responsible for its binding to anionic phosphatidylinositol 4,5-bisphosphate (PI(4,5)P₂) lipids in the inner leaflet of the plasma membrane (Itoh et al. [2001]). Once bound, epsin1 facilitates endocytosis by inducing positive membrane curvature and driving the assembly of clathrin coat components (Ford et al. [2002]). Previous studies have shown that ENTH favours curved membranes (Capraro et al. [2010]), but also interacts with flat membranes to drive curvature by shallow insertion of its N-terminal helix, named helix-0 (Stahelin et al. [2003], Kweon et al. [2006], Lai et al. [2012]). Consistent with these findings, our screening method successfully highlighted helix-0 as a membrane binder (Fig. 3B). Besides the crucial insertion of hydrophobic residues (e.g. L6 (Ford et al. [2002])), helix-0 has a net charge of +4 to complement the negatively charged PI(4,5)P₂-rich membrane domains that it binds to. The correct identification of this helix in the ENTH domain indicates that, even though the simulations underlying our model were all performed on neutral POPC membranes, the charge correction (eq. 2) we apply is sufficiently accurate to distill these effects.

Bin1 N-BAR

Finally, we screened the N-BAR domain from Bin1, which is a crescent-shaped protein that polymerises to drive the remodelling of membranes into tubes. Hereto, we took the predicted structure of the full Bin1 protein from the AF database, and isolated the N-BAR domain (residue 1-239). Our screen highlighted part of the N-terminal helix (just like for ENTH, named helix-0), classifying it as a curvature sensor (Fig. 3C). When comparing this to the extensive literature on Bin1's helix-0, we see that – contrary to our method's classification – it behaves as a binder in liposome experiments, which show that its membrane binding is not curvature-dependent (Fernandes et al. [2008]) and that both the full Bin1 N-BAR domain and the amphipathic helix-0 alone can induce tubulation (Löw et al. [2008]). Given that the BAR domain by itself (so, without helix-0), also induces liposome tubulation (Peter et al. [2004]), the current understanding on this is that helix-0 is responsible for the initial docking and positive curvature induction, which then triggers binding of the full crescent-shaped BAR domain and the subsequent polymerisation of the full protein scaffold (Adam et al. [2015], Simunovic et al. [2015]). Additionally, our screening tool predicted three other sensing/binding domains that were not sufficiently solvent exposed (SASA < 0.8 nm²) to be highlighted in the results (see data repository at <https://github.com/nvanhilten/PMIpred>). Two of these domains are indeed located at the membrane-binding face of the concave N-BAR domain.

When taking together our screening results for ArfGAP1, epsin1, and Bin1, we conclude that our method was indeed able to correctly pick out the important regions from these proteins (Fig. 3A-C). However, the sensing/binding classification is not perfect. This likely stems from a combination of inaccuracies in the model and inconsistencies between experimental setups. After all, curvature sensing

and membrane binding will be strongly affected by liposome formulations, peptide concentrations, and potential synergistic effects between proteins that may differ between experiments and are not accounted for in our simulations.

Comparison with other tools

Current tools that predict membrane-interaction sites in proteins (DREAMM (Chatzigoulas and Cournia [2022b,a]), PPM3 (Lomize et al. [2022]), and MODA (Kufareva et al. [2014])) produce output on the single-residue level, whereas PMIpred quantifies membrane-binding *regions*; averaging the free energy over the local surrounding of a given amino acid. This complicates a one-to-one comparison with these qualitative prediction methods on the single-residue level (e.g. as done in ref. (Chatzigoulas and Cournia [2022b])). Therefore, we proceeded to benchmark PMIpred against the other web servers in the following manner: we composed a test set of 27 PMPs, including all 19 proteins from Chatzigoulas & Cournia's test set (Chatzigoulas and Cournia [2022b]) (mainly comprising lipid kinases and lipid transfer proteins) and 8 known lipid packing defect sensing proteins (2 PDB structures and 6 AF predictions); see Table S2 for details. We treated the PDB and AF structures in the latter test set separately, because DREAMM, PPM3, and MODA were not designed to be applied to (partly unstructured) AF predictions.

For all these protein structures, we segmented the sequences into consecutive fragments of 15 AAs (matching PMIpred's default window size), which were marked *positive* if they contained > 1 residue that was predicted to be membrane-interacting (either sensing or binding for PMIpred), and *negative* if none of the 15 residues in the fragment was labeled as a 'hit'. We argue that, although it provides fundamentally different insights than e.g. DREAMM, this 'protein-region'-way of looking at membrane-binding interactions is more akin to the way a structural biologist may look at a protein structure, with loops or helices (rather than independent individual residues) dynamically inserting into the membrane leaflet.

We performed this segment-based binary classification analysis and calculated the Matthews correlation coefficient (MCC, see S4) for the four methods on the different test sets (Fig. 3D). We observed that all tools score similarly on the full test set (MCC ≈ 0.45). When only considering the proteins in Chatzigoulas & Cournia's test set, we see that PMIpred is outperformed, mainly by DREAMM and PPM3, due to the fact that these tools are – by design – focused on predicting the single-residue protein-membrane interactions that are typical for the types of PMPs in this test set. In contrast, PMIpred performed much better when predicting the (often less strictly defined) membrane-interacting regions of lipid packing defect sensors, both for the PDB and the AF structures. Details on all results described in this section are provided in Table S3 and the data are accessible at <https://github.com/nvanhilten/PMIpred>.

In general, we note that the majority of erroneous predictions are false positives rather than false negatives, with a false discovery rate (FDR) of 60% and a false omission rate (FOR) of only 4.4% for PMIpred (and similar values for the other web servers, see S4).

Taken together, we argue that PMIpred's predictions are generally on par with state-of-the-art tools like DREAMM, PPM3, and MODA, with the important improvement that PMIpred additionally provides *quantitative* information on membrane-interaction propensity per residue in actual thermodynamic terms (free energies) in the sense that one

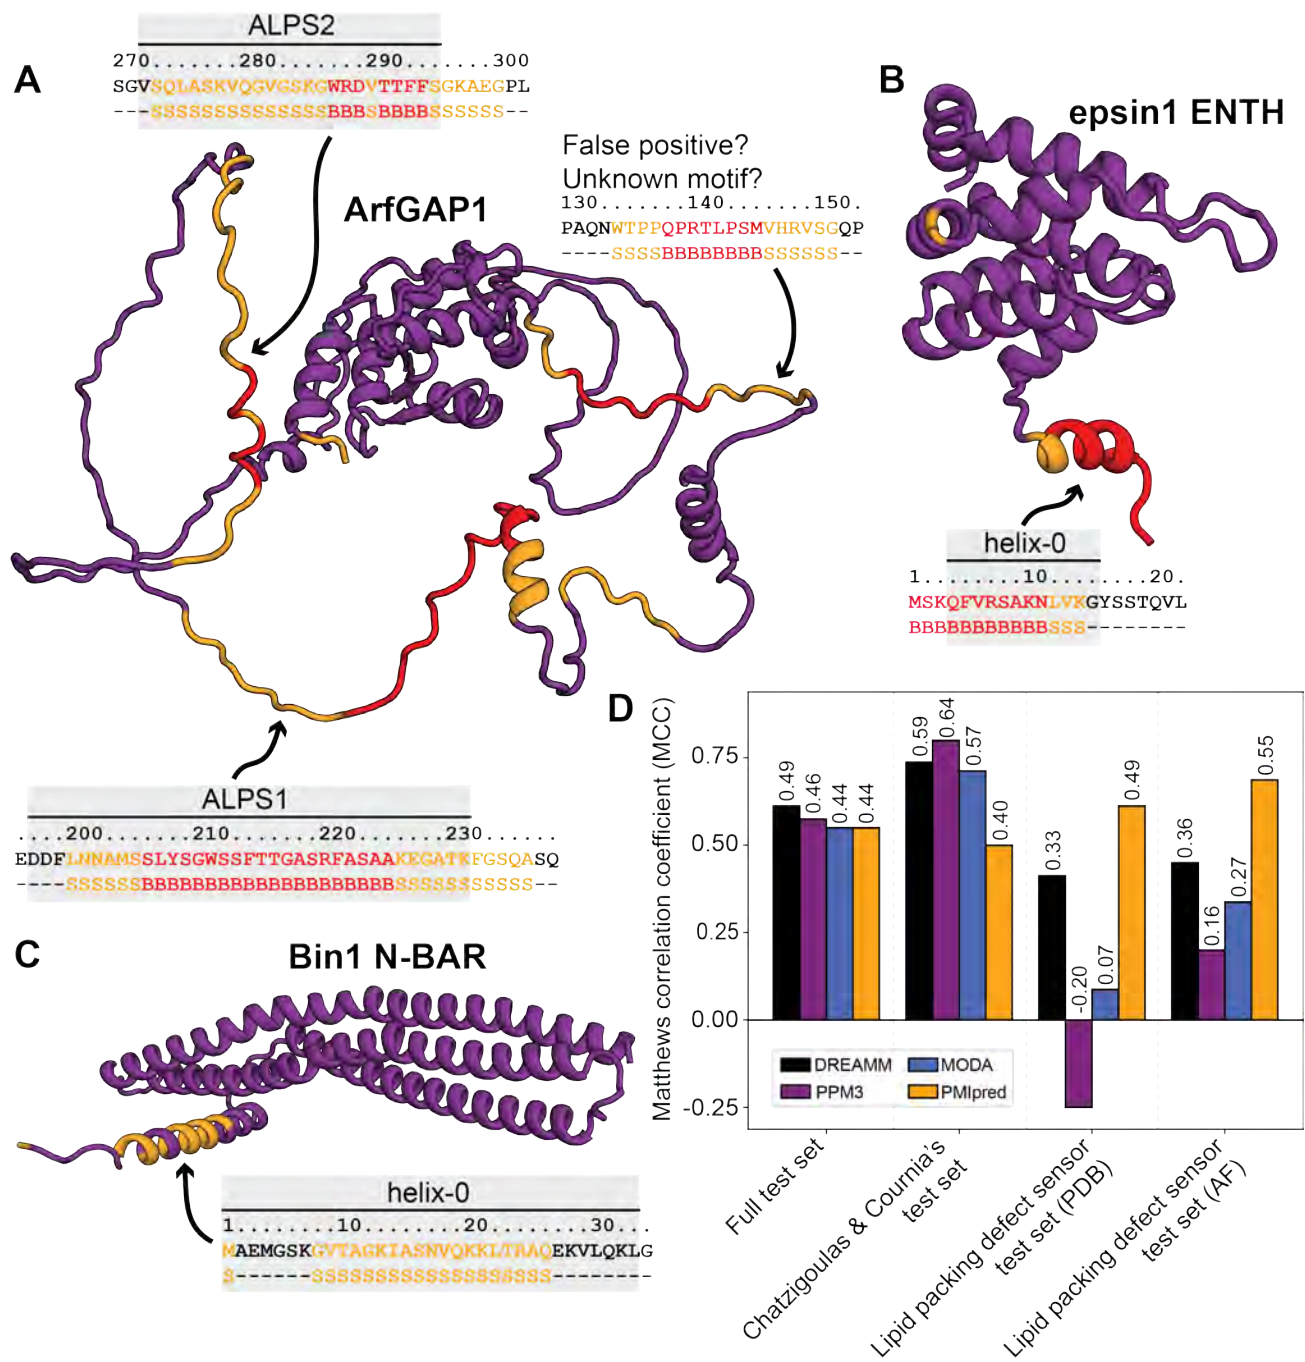


Fig. 3. A-C) Screening results for ArfGAP1 (A, AF: Q8N6T3), the ENTH domain from epsin1 (B, PDB: 5onf), and the N-BAR domain from Bin1 (C, AF: Q9BTH3). Insets show sequences and corresponding predicted classification of putative active regions. Purple: non-binding or SASA below threshold (-). Orange: curvature sensing (S). Red: membrane binding (B). The known membrane-interaction regions are marked in gray. **D)** Comparing PMIpred to DREAMM, PPM3, and MODA.

can directly rank and quantitatively compare its numerical predictions. In contrast, DREAMM and PPM3 act as binary classifiers and MODA only yields membrane-interaction 'likelihoods', that are hard to interpret biophysically. We should note, however, that direct comparison of the predicted $\Delta\Delta F$ values with experiments is complicated for two reasons: (1) its value is strongly dependent on the degree of stretching (curvature) in the target membrane (as discussed previously

(van Hilten et al. [2022])), and (2) the contribution of the entropy is typically underestimated in free energies that are calculated from simulations.

Because PMIpred is trained on bottom-up physics-based data, it is likely more self-consistent at generalising those concepts to undiscovered or *de novo* designed protein structures that are highly distinct from typical known PMPs, compared to data-informed models that are trained on protein databases.

To this end, we suggest that the physics-informed transformer neural network we developed here may be able to additionally boost the performance of prediction methods based on ensemble machine learning approaches such as DREAMM by serving as an additional, independent (strong) classifier, especially since these approaches are notably orthogonal. Finally, we emphasise that users now have multiple different tools to their disposal and encourage them to use these various approaches to compare results and draw conclusions based on the ensemble of results.

Try it yourself

Our protein screening tool is accessible within the PMIPred web server. Here, users can submit a PDB-file or PDB ID and also tune the size of the sliding window and the SASA threshold value (Fig. 4A). Similar to PMIPred's peptide module, users can specify if their target membrane is neutral (e.g. POPC) or negatively charged, as is common in biology. Again, if the latter is chosen, charge correction (eq. 2) is applied to the $\Delta\Delta F$ value to account for additional coulomb interactions and concomitant 'charge sensing'.

After submission, the results are displayed in the browser (Fig. 4B) and available to be downloaded as a zip-file. The processed protein is displayed interactively on the results page, using the aforementioned colouring: purple, orange, and red, for the non-binding (or buried) regions, predicted sensing domains, and membrane-binding domains, respectively. The data underlying this colouring are listed per segment (Fig. 4C) and per residue (Fig. 4D).

We temporarily store the latest 100 protein structure submissions on our local server under randomized file names to ensure anonymity.

Cautionary remarks & recommendations

Before concluding this paper, we wish to point out several cases in which PMIPred's results should be interpreted with some extra caution.

1. We note that shorter peptides have a higher tendency to be classified as binders (Fig. S5A). This effect can not be explained by poorer accuracy in the transformer's prediction for shorter sequences, because the normalised error is constant over the full length range (7-24 residues, Fig. S5B). Alternatively, we suggest that this bias is likely introduced by our length extrapolation step (eq. 1) that intrinsically has a more pronounced effect for shorter peptides. Therefore, we advise users to interpret predictions on short peptides (≤ 12 residues) with some caution.
2. Our charge-correction (eq. 2) is based on a set of benchmark peptides from the experimental literature with charges ranging from $z = -1$ to $z = +4$ (van Hilten et al. [2023]). Predictions for peptides outside this charge range should be considered with some extra caution.
3. The interplay between membrane binding and peptide folding is also linked to the lipid composition of the membrane. This aspect is simplified in the MD simulations on which our model is trained, because we only used pure POPC and POPC/POPE membranes. If a protein's membrane binding behaviour is linked to interactions with a specific lipid species, PMIPred may underestimate the membrane affinity.
4. The N- and C-terminal residues of proteins are more likely to be identified as (false) positives by the PMIPred protein screening module, for two reasons: (1) the termini of

proteins tend to have a high solvent-accessibility. And, (2) the per-residue $\Delta\Delta F$ is averaged over fewer positions, and thus possibly less accurate. Please note that the same applies to gaps in protein chains.

5. As mentioned before, PMIPred always assumes an α -helical fold. Although our MD simulations suggest that folding (helical *vs* random coil) is a less determinant factor than AA composition (Fig. S4), we recommend users to take extra caution when working with peptides that feature a known functional non-helical secondary structure. Of course, this also applies to non-helical domains in full-length proteins.
6. When using predicted 3D structures (e.g. from AF), we advise to only submit structures with high confidence – $\text{pLDDT} \geq 70$ –, unless a section is known to be unstructured.
7. Protein-protein interaction sites may be falsely identified as membrane-binding regions by PMIPred, within the protein context. This misinterpretation is due to the fact that both phenomena are driven mostly by hydrophobic residues being exposed at the surface of the protein. For example, we observed this for the RAS-binding domain of PI(4,5)P₂ 3-kinase, which included two exposed leucines L233 and L234 (see Table S3). We should emphasize here, however, that our method is based on peptide-data, and that the membrane-interaction predictions on the *peptide fragments* that make up a protein-protein interaction domain may still be valid.
8. Applying PMIPred to lipid-anchored proteins (e.g. KRAS in section S6) requires some extra caution, because our web server only considers the protein's amino acids and neglects lipidic membrane anchors (e.g. Cys-linked farnesyl or palmitoyl moieties).
9. Although PMIPred is not designed for it, it can also be applied to transmembrane proteins, since the physicochemical properties driving membrane-surface binding and membrane internalisation are similar (hydrophobicity, mostly). We do note, however, that it may miss transmembrane regions for multi-pass proteins (e.g. GPCRs) due to the crowded protein environment and consequent low 'solvent'-accessibility. Generally, we found that PMIPred behaves similar to DREAMM in this regard (Fig. S7). When working on transmembrane proteins, we recommend using PMIPred in conjunction with other web servers that are specifically developed for this (e.g. DeepTMHMM (Hallgren et al. [2022]), TMAAlphaFold (Dobson et al. [2023]), or others).

Conclusions

We have developed a transformer neural network that, with the sole input requirement of an amino acid sequence, can predict the relative binding free energy ($\Delta\Delta F$) for peptides interacting with curved membranes. Additionally, it provides a basic estimation of ΔF through linear interpolation. We implemented this model into a user-friendly web server, named PMIPred (<https://pmipred.fkt.physik.tu-dortmund.de>), where researchers can readily predict sensing free energies for any peptide, without requiring any simulation or coding experience.

As a second module in PMIPred, users can screen protein structures for putative curvature-sensing or membrane-binding motifs. Hereto, a sliding window is applied to the protein sequence to calculate a free-energy contribution for every

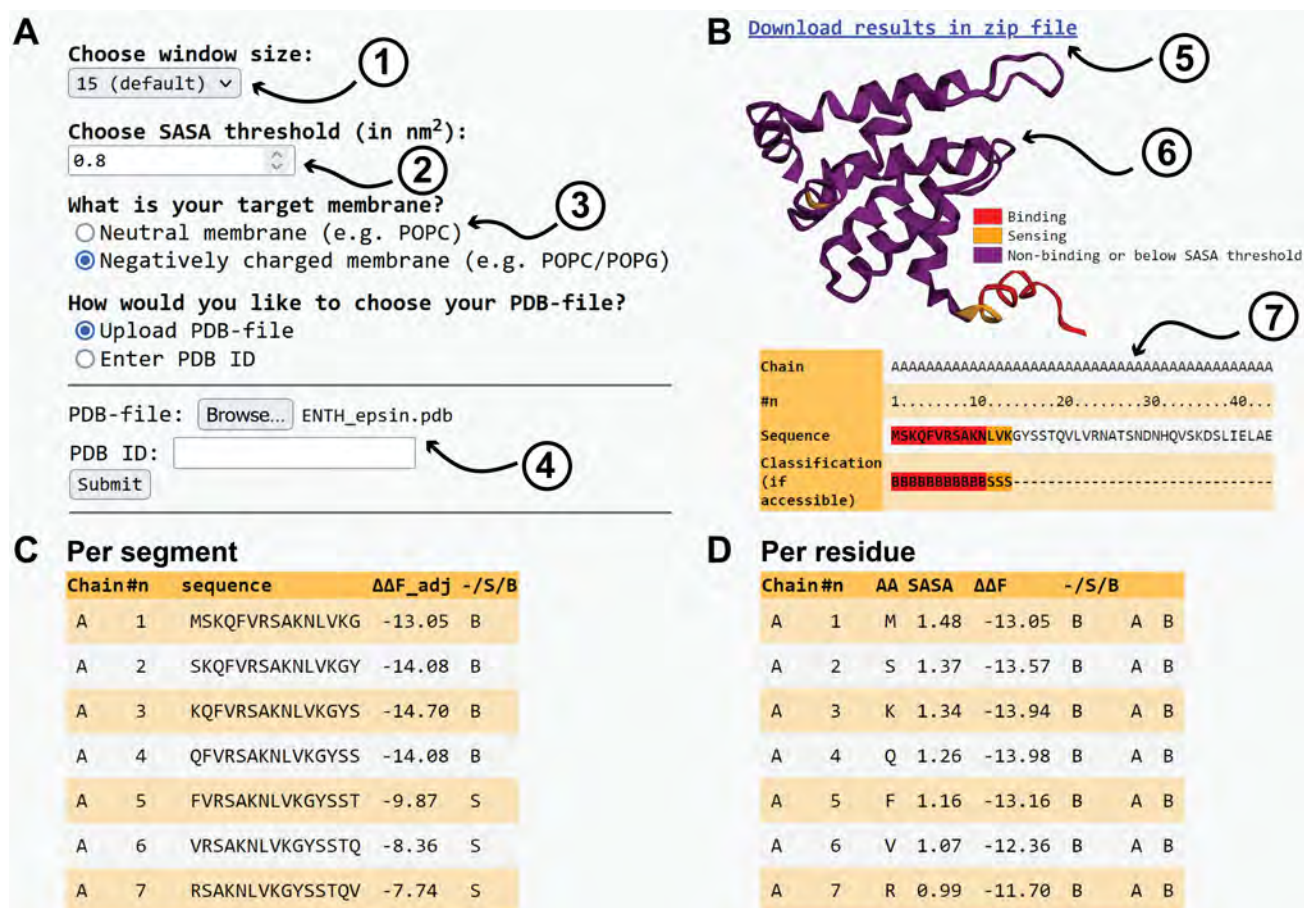


Fig. 4. **A)** The module takes four inputs: (1) The size of the sliding window (from 7-24 residues, default=15). (2) The SASA threshold (default=0.8 nm²), above which a residue is considered sufficiently solvent-accessible to contribute to membrane interactions. (3) The target membrane. (4) An uploaded PDB-file or PDB ID. **B)** Output includes: (5) A downloadable zip-file containing all output files that are also visualised on this page. (6) An interactive visualisation of the protein structure using 3Dmol (Rego and Koes [2015]). Predicted non-binding, sensing, and binding regions are coloured purple, orange, and red, respectively. (7) An output table containing the full protein sequence and the annotated classification: B=binder, S= sensor, -=non-binder or not accessible (SASA below threshold). **C)** A table containing the free energies and classifications of the individual segments that were screened. **D)** A more detailed look into the data underlying the table in B, including the SASA, $\Delta\Delta F$ values, and classification per residue.

residue. By coupling this to the solvent-accessible surface area (SASA) of that residue, the module is able to map and visualise the predicted membrane-interaction activity onto the protein structure. We applied this method to a comprehensive and diverse test set of known PMPs, and show that PMIpred is on par with the state-of-the-art tools DREAMM, PPM3, and MODA. In contrast to these other web servers, PMIpred provides quantitative predictions in terms of relative membrane-interaction free energies that can be directly interpreted biophysically, and additionally enables users to distinguish curvature-sensing from membrane-binding motifs. PMIpred also quantifies the contribution/weight of each individual amino acids to the predicted membrane binding, which guides the design of point mutations to experimentally probe protein-membrane interactions.

Since this method is trained on physics-based data (where experimental findings contribute only indirectly through the parametrisation of the MD force field), we argue that it should be able to characterise membrane-interaction behaviour from a broadly applicable fundamental perspective. In contrast, most current approaches are based on physicochemical descriptors, sequence ‘grammar rules’ that define known curvature-sensing

motifs (e.g. ALPS (Drin et al. [2007])) but are unable to generalise their predictions to unrelated membrane-interacting motifs (e.g. α -synuclein and helix-0 of N-BAR and ENTH). When it comes to predicting the membrane interaction sites in PMPs, current methods are mainly data-driven, and thus rely on experimental findings on natural proteins. Although these tools are already remarkably powerful, we argue that the unique physics-informed predictions by PMIpred can additionally boost the performance of existing ensemble classifiers such as DREAMM, especially because it provides insights from a completely different angle and the data sources are orthogonal. Ultimately, by making PMIpred easily accessible for both computational and experimental biologists, we aim to accelerate the discovery and characterisation of curvature-sensing and membrane-binding peptides and protein motifs.

Competing interests

No competing interest is declared.

Author contributions statement

NvH and HJR designed the research and wrote the manuscript. NvH, NV, JM, and AB wrote the code. NV and AB designed, trained, and optimised the transformer model. NV and NvH designed the web server. CN was responsible for hosting the web server, and provided technical assistance.

Acknowledgments

The authors would like to thank Agata Witkowska for fruitful discussions and feedback. Art Hoti is thanked for thoroughly testing the web server. The Dutch Research Organisation NWO (Snellius@Surfsara) and the HLRN Göttingen/Berlin are acknowledged for the provided computational resources. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2033 - 390677874 - RESOLV. We also thank the NWO Vidi scheme (project number 723.016.005) for funding. We gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).

References

- J. Adam, N. Basnet, and N. Mizuno. Structural insights into the cooperative remodeling of membranes by amphiphysin/bin1. *Sci. Rep.*, 5(1):1–15, 2015.
- H. Ago, M. Oda, M. Takahashi, H. Tsuge, N. Ochi, S. and Katunuma, M. Miyano, and J. Sakurai. Structural basis of the sphingomyelin phosphodiesterase activity in neutral sphingomyelinase from *Bacillus cereus*. *J. Biol. Chem.*, 281(23):16157–16167, 2006.
- T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. *Optuna: A Next-generation Hyperparameter Optimization Framework*, pages 2623–2631. 2019.
- J. Baylon, I. Lenov, S. Sligar, and E. Tajkhorshid. Characterizing the membrane-bound state of cytochrome P450 3a4: structure, depth of insertion, and orientation. *J. Am. Chem. Soc.*, 135(23):8542–8551, 2013.
- J. Bigay, P. Gounon, S. Robineau, and B. Antonny. Lipid packing sensed by ArfGAP1 couples COPI coat disassembly to membrane bilayer curvature. *Nature*, 426(6966):563–566, 2003.
- J. Bigay, J. Casella, G. Drin, B. Mesmin, and B. Antonny. ArfGAP1 responds to membrane curvature through the folding of a lipid packing sensor motif. *EMBO J.*, 24(13):2244–2253, 2005.
- B. Capraro, Y. Yoon, W. Cho, and T. Baumgart. Curvature sensing by the epsin n-terminal homology domain measured on cylindrical lipid membrane tethers. *J. Am. Chem. Soc.*, 132(4):1200–1201, 2010.
- A. Chatzigoulas and Z. Cournia. DREAMM: a web-based server for drugging protein-membrane interfaces as a novel workflow for targeted drug design. *Bioinformatics*, 38(24):5449–5451, 2022a.
- A. Chatzigoulas and Z. Cournia. Predicting protein-membrane interfaces of peripheral membrane proteins using ensemble machine learning. *Brief. Bioinform.*, 23(2):bbab518, 2022b.
- X. Chen, D. Wolfgang, and N. Sampson. Use of the parallax-quench method to determine the position of the active-site loop of cholesterol oxidase in lipid bilayers. *Biochemistry*, 39(44):13383–13389, 2000.
- P. Cock, T. Antao, J. Chang, B. Chapman, C. Cox, A. Dalke, I. Friedberg, T. Hamelryck, F. Kauff, B. Wilczynski, and M. de Hoon. Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics*, 25(11):1422–1423, 2009.
- B. Córscico, H. Liou, and J. Storch. The alpha-helical domain of liver fatty acid binding protein is responsible for the diffusion-mediated transfer of fatty acids to phospholipid membranes. *Biochemistry*, 43(12):3600–3607, 2004.
- J. Costeira-Paulo, J. Gault, G. Popova, M. Ladds, I. van Leeuwen, M. Sarr, A. Olsson, D. Lane, S. Lafn, E. Marklund, and M. Landreh. Lipids shape the electron acceptor-binding site of the peripheral membrane protein dihydroorotate dehydrogenase. *Cell Chem. Biol.*, 25(3):309–317.e4, 2018.
- H. Cui, E. Lyman, and G. Voth. Mechanism of membrane curvature sensing by amphipathic helix containing proteins. *Biophys. J.*, 100(5):1271–1279, 2011.
- E. Cukierman, I. Huber, M. Rotman, and D. Cassel. The ARF1 GTPase-activating protein: zinc finger motif and golgi complex localization. *Science*, 270(5244):1999–2002, 1995.
- W. Davidson, A. Jonas, D. Clayton, and J. George. Stabilization of alpha-synuclein secondary structure upon binding to synthetic membranes. *J. Biol. Chem.*, 273(16):9443–9449, 1998.
- L. Dobson, L. Szekeres, C. Gerdán, T. Langó, A. Zeke, and G. Tusnády. TmAlphaFold database: membrane localization and evaluation of AlphaFold2 predicted alpha-helical transmembrane protein structures. *Nucleic Acids Res.*, 51(D1):D517–D522, 2023.
- G. Drin and B. Antonny. Amphipathic helices and membrane curvature. *FEBS Lett.*, 584(9):1840–1847, 2010.
- G. Drin, J. Casella, R. Gautier, T. Boehmer, T. Schwartz, and B. Antonny. A general amphipathic alpha-helical motif for sensing membrane curvature. *Nat. Struct. Mol. Biol.*, 14(2):138–146, 2007.
- P. Dubovskii, D. Dementieva, E. Bocharov, Y. Utkin, and A. Arseniev. Membrane binding motif of the p-type cardiotoxin. *J. Mol. Biol.*, 305(1):137–149, 2001.
- D. Eisenberg, R. Weiss, and T. Terwilliger. The helical hydrophobic moment: a measure of the amphiphilicity of a helix. *Nature*, 299(5881):371–374, 1982.
- L. Falomir-Lockhart, L. Laborde, P. Kahn, J. Storch, and B. Córscico. Protein-membrane interaction and fatty acid transfer from intestinal fatty acid-binding protein to membranes. support for a multistep process. *J. Biol. Chem.*, 281(20):13979–13989, 2006.
- W. Fan, A. Nassiri, and Q. Zhong. Autophagosome targeting and membrane curvature sensing by barkor/atg14(1). *Proc. Natl. Acad. Sci.*, 108(19):7769–7774, 2011.
- J. Fauchère and V. Pliška. Hydrophobic parameters-pi of amino-acid side-chains from the partitioning of n-acetyl-amino-acid amides. *Eur. J. Med. Chem.*, 18:369–375, 1983.
- L. Feng, W. Chan, S. Roderick, and D. Cohen. High-level expression and mutagenesis of recombinant human phosphatidylcholine transfer protein using a synthetic gene: evidence for a c-terminal membrane binding domain. *Biochemistry*, 39(50):15399–15409, 2000.
- F. Fernandes, L. Loura, F. Chichón, J. Carrascosa, A. Fedorov, and M. Prieto. Role of helix 0 of the N-BAR domain in membrane curvature generation. *Biophys. J.*, 94(8):3065–3073, 2008.

- 1
2 M. Ford, I. Mills, B. Peter, Y. Vallis, G. Praefcke, P. Evans,
3 and H. McMahon. Curvature of clathrin-coated pits driven
4 by epsin. *Nature*, 419(6905):361–366, 2002.
- 5 S. Gabelli, C.-H. Huang, D. Mandelker, O. Schmidt-Kittler,
6 B. Vogelstein, and L. Amzel. Structural effects of oncogenic
7 PI3K α mutations. *Curr. Top. Microbiol. Immunol.*, 347:
8 43–53, 2010.
- 9 R. Gautier, B. Douguet, D. and Antonny, and G. Drin.
10 HELIQUEST: a web server to screen sequences with specific
11 α -helical properties. *Bioinformatics*, 24(18):2101–2102,
12 2008.
- 13 J. Hallgren, K. Tsirigos, M. Pedersen, J. J. Almagro A.,
14 P. Marcatili, H. Nielsen, A. Krogh, and O. Winther.
15 DeepTMHMM predicts alpha and beta transmembrane
16 proteins using deep neural networks. *bioRxiv*, 2022.
- 17 K. Hristova and S. H. White. An experiment-based algorithm
18 for predicting the partitioning of unfolded peptides into
19 phosphatidylcholine bilayer interfaces. *Biochemistry*, 44(37):
20 12614–12619, 2005.
- 21 D. Iaea, I. Dikiy, I. Kiburu, D. Eliezer, and F. Maxfield.
22 STAR4D membrane interactions and sterol binding.
23 *Biochemistry*, 54(30):4623–4636, 2015.
- 24 T. Itoh, S. Koshiba, T. Kigawa, A. Kikuchi, S. Yokoyama,
25 and T. Takenawa. Role of the ENTH domain in
26 phosphatidylinositol-4,5-bisphosphate binding and
27 endocytosis. *Science*, 291(5506):1047–1051, 2001.
- 28 M. Jensen, V. Bhatia, C. Jao, J. Rasmussen, S. Pedersen,
29 K. Jensen, R. Langen, and D. Stamou. Membrane curvature
30 sensing by amphipathic helices: a single liposome study using
31 α -synuclein and annexin B12. *J. Biol. Chem.*, 286(49):
32 42603–42614, 2011.
- 33 Y. Ju, Z. Liu, Z. Zhang, L. Duan, Q. Liu, Q. Gu,
34 C. Zhang, J. Xu, and H. Zhou. Membrane binding of the
35 insertion sequence of proteus vulgaris l-amino acid deaminase
36 stabilizes protein structure and increases catalytic activity.
37 *Sci. Rep.*, 7(1):13719, 2017.
- 38 J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov,
39 O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek,
40 A. Potapenko, A. Bridgland, C. Meyer, S. Kohl, A. Ballard,
41 A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler,
42 T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski,
43 M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein,
44 D. Silver, O. Vinyals, A. Senior, K. Kavukcuoglu, P. Kohli,
45 and D. Hassabis. Highly accurate protein structure
46 prediction with AlphaFold. *Nature*, 596(7873):583–589,
47 2021.
- 48 H. Jung, J. Lee, S. Kim, Y.-J. Eu, S. Shin, M. Milescu,
49 K. Swartz, and J. Kim. Solution structure and lipid
50 membrane partitioning of VSTx1, an inhibitor of the KvAP
51 potassium channel. *Biochemistry*, 44(16):6015–6023, 2005.
- 52 Y. Jung, B. Kong, S. Moon, S.-H. Yu, J. Chung, C. Ban, W.-J.
53 Chung, S.-G. Kim, and D.-H. Kweon. Envelope-deforming
54 antiviral peptide derived from influenza virus M2 protein.
55 *Biochem. Biophys. Res. Commun.*, 517(3):507–512, 2019.
- 56 I. Karageorgos, V. Silin, N. Zvonok, J. Marino, D. Janero,
57 and A. Makriyannis. The role of human monoacylglycerol
58 lipase (hMAGL) binding pocket in breakup of unsaturated
59 phospholipid membranes. *Anal. Biochem.*, 536:90–95, 2017.
- 60 E. Karanasios, G.-S. Han, Z. Xu, G. Carman, and
S. Siniosoglou. A phosphorylation-regulated amphipathic
helix controls the membrane translocation and function of the
yeast phosphatidate phosphatase. *Proc. Natl. Acad. Sci.*,
107(41):17539–17544, 2010.
- P. Kiser, E. Farquhar, W. Shi, X. Sui, M. Chance, and
K. Palczewski. Structure of RPE65 isomerase in a lipidic
matrix reveals roles for phospholipids and iron in catalysis.
Proc. Natl. Acad. Sci. U. S. A., 109(41):E2747–56, 2012.
- I. Kufareva, M. Lenoir, F. Dancea, P. Sridhar, E. Raush,
C. Bissig, J. Gruenberg, R. Abagyan, and M. Overduin.
Discovery of novel membrane binding structures and
functions. *Biochem. Cell Biol.*, 92(6):555–563, 2014.
- D.-H. Kweon, Y.-K. Shin, J. Shin, J.-H. Lee, J.-B. Lee, J.-
H. Seo, and Y. Kim. Membrane topology of helix 0 of the
epsin n-terminal homology domain. *Mol. Cells*, 21(3):428–
435, 2006.
- C.-L. Lai, C. Jao, E. Lyman, J. Gallop, B. Peter, H. McMahon,
R. Langen, and G. Voth. Membrane binding and self-
association of the epsin n-terminal homology domain. *J. Mol.
Biol.*, 423(5):800–817, 2012.
- C. Lau, G. King, and M. Mobli. Molecular basis of the
interaction between gating modifier spider toxins and the
voltage sensor of voltage-gated ion channels. *Sci. Rep.*, 6
(1):34333, 2016.
- J. Lee, H. Yang, M. Georgescu, A. Di Cristofano, T. Maehama,
Y. Shi, J. Dixon, P. Pandolfi, and N. Pavletich. Crystal
structure of the PTEN tumor suppressor: implications for
its phosphoinositide phosphatase activity and membrane
association. *Cell*, 99(3):323–334, 1999.
- A. Lomize, S. Todd, and I. Pogozheva. Spatial arrangement
of proteins in planar and curved membranes by PPM 3.0.
Protein Sci., 31(1):209–220, 2022.
- A. Lookene, N. Groot, J. Kastelein, G. Olivecrona, and
T. Bruin. Mutation of tryptophan residues in lipoprotein
lipase. effects on stability, immunoreactivity, and catalytic
properties. *J. Biol. Chem.*, 272(2):766–772, 1997.
- C. Löw, U. Weininger, H. Lee, K. Schweimer, I. Neundorf,
A. Beck-Sickingler, R. Pastor, and J. Balbach. Structure and
dynamics of helix-0 of the n-bar domain in lipid micelles and
bilayers. *Biophys. J.*, 95(9):4315–4323, 2008.
- B. Matthews. Comparison of the predicted and observed
secondary structure of T4 phage lysozyme. *Biochim.
Biophys. Acta - Protein Struct.*, 405(2):442–451, 1975.
- B. Mesmin, G. Drin, S. Levi, M. Rawet, D. Cassel, J. Bigay,
and B. Antonny. Two lipid-packing sensor motifs contribute
to the sensitivity of ArfGAP1 to membrane curvature.
Biochemistry, 46(7):1779–1790, 2007.
- V. Monje-Galvan and J. Klauda. Peripheral membrane
proteins: Tying the knot between experiment and
computation. *Biochim. Biophys. Acta*, 1858(7):1584–1593,
2016.
- L. Monticelli, S. Kandasamy, X. Periolo, R. Larson,
D. Tieleman, and S. Marrink. The martini coarse-grained
force field: Extension to proteins. *J. Chem. Theory
Comput.*, 4(5):819–834, 2008.
- A. Müller, G. Gabernet, J. Hiss, and G. Schneider. modAMP:
Python for antimicrobial peptides. *Bioinformatics*, 33(17):
2753–2755, 2017.
- S. Neumann, M. Opacić, R. Wechselberger, H. Sprong, and
M. Egmond. Glycolipid transfer protein: clear structure and
activity, but enigmatic function. *Adv. Enzyme Regul.*, 48(1):
137–151, 2008.
- H. Nguyen, Y. Afkari, H. Senoo, H. Sesaki, P. Devreotes, and
M. Iijima. Mechanism of human PTEN localization revealed
by heterologous expression in dictyostelium. *Oncogene*, 33
(50):5688–5696, 2014.
- H. Ohvo-Rekilä and P. Mattjus. Monitoring glycolipid transfer
protein activity and membrane interaction with the surface

- plasmon resonance technique. *Biochim. Biophys. Acta*, 1808 (1):47–54, 2011.
- S.-I. Ozawa, T. Kimura, H. Nozaki, T. and Harada, I. Shimada, and M. Osawa. Structural basis for the inhibition of voltage-dependent K^+ channel by gating modifier toxin. *Sci. Rep.*, 5 (1):14226, 2015.
- P. Papadopoulou, R. van der Pol, N. van Hilten, W. van Os, R. Pattipeiluhu, G. Arias-Alpizar, R. Knol, W. Noteborn, M. Moradi, M. Ferraz, J. Aerts, N. Sommerdijk, F. Campbell, H. Risselada, G. Sevink, and A. Kros. Phase-separated lipid-based nanoparticles: Selective behavior at the nano-bio interface. *Adv. Mat.*, 2023.
- B. Peter, H. Kent, I. Mills, Y. Vallis, P. Butler, P. Evans, and H. McMahon. BAR domains as sensors of membrane curvature: the amphiphysin BAR structure. *Science*, 303 (5657):495–499, 2004.
- S. Phillips, K. Ile, M. Boukhelifa, R. Huijbregts, and V. Bankaitis. Specific and nonspecific membrane-binding determinants cooperate in targeting phosphatidylinositol transfer protein beta-isoform to the mammalian trans-golgi network. *Mol. Biol. Cell*, 17(6):2498–2512, 2006.
- I. Pranke, V. Morello, J. Bigay, K. Gibson, J.-M. Verbavatz, B. Antonny, and C. Jackson. α -synuclein and ALPS motifs are membrane curvature sensors whose contrasting chemistry mediates selective vesicle binding. *J. Cell Biol.*, 194(1):89–103, 2011.
- J. Rawls, W. Knecht, K. Diekert, R. Lill, and M. Löffler. Requirements for the mitochondrial import and localization of dihydroorotate dehydrogenase. *Eur. J. Biochem.*, 267(7): 2079–2087, 2000.
- N. Rego and D. Koes. 3Dmol.js: molecular visualization with WebGL. *Bioinformatics*, 31(8):1322–1324, 2015.
- Schrödinger, LLC. The PyMOL molecular graphics system, version 1.8. 2015.
- A. Shrake and J. Rupley. Environment and exposure to solvent of protein atoms. lysozyme and insulin. *J. Mol. Biol.*, 79(2): 351–371, 1973.
- M. Simunovic, G. Voth, A. Callan-Jones, and P. Bassereau. When physics takes over: BAR proteins and membrane curvature. *Trends Cell Biol.*, 25(12):780–792, 2015.
- P. Souza, R. Alessandri, J. Barnoud, S. Thallmair, I. Faustino, F. Grünewald, I. Patmanidis, H. Abdizadeh, B. Bruininks, T. Wassenaar, P. Kroon, J. Melcr, V. Nieto, V. Corradi, H. Khan, J. Domański, M. Javanainen, H. Martinez-Seara, N. Reuter, R. Best, I. Vattulainen, L. Monticelli, X. Periole, D. Tieleman, A. de Vries, and S. Marrink. Martini 3: a general purpose force field for coarse-grained molecular dynamics. *Nat. Methods*, 18(4):382–388, 2021.
- R. Stahelin, F. Long, B. Peter, D. Murray, P. De Camilli, H. McMahon, and W. Cho. Contrasting membrane interaction mechanisms of AP180 n-terminal homology (ANTH) and epsin n-terminal homology (ENTH) domains. *J. Biol. Chem.*, 278(31):28993–28999, 2003.
- T. Sugiki, K. Takeuchi, T. Yamaji, Y. Takano, T. and Tokunaga, K. Kumagai, K. Hanada, H. Takahashi, and I. Shimada. Structural basis for the golgi association by the pleckstrin homology domain of the ceramide trafficking protein (CERT). *J. Biol. Chem.*, 287(40):33706–33718, 2012.
- X. Sui, P. Kiser, J. v. Lintig, and K. Palczewski. Structural basis of carotenoid cleavage: from bacteria to mammals. *Arch. Biochem. Biophys.*, 539(2):203–213, 2013.
- B. C. Tan, K. Cline, and D. R. McCarty. Localization and targeting of the VP14 epoxy-carotenoid dioxygenase to chloroplast membranes. *Plant J.*, 27(5):373–382, Sept. 2001.
- T. Tubiana, I. Sillitoe, C. Orengo, and N. Reuter. Dissecting peripheral protein-membrane interfaces. *PLOS Comput. Biol.*, 18(12):e1010346, 2022.
- N. van Hilten, K. Stroh, and H. Risselada. Efficient quantification of lipid packing defect sensing by amphipathic peptides: Comparing martini 2 and 3 with CHARMM36. *J. Chem. Theory Comput.*, 18(7):4503–4514, 2022.
- N. van Hilten, J. Methorst, N. Verwei, and H. Risselada. Physics-based generative model of curvature sensing peptides; distinguishing sensors from binders. *Sci. Adv.*, 9 (11):eade8839, 2023.
- S. Varadi, M. and Anyango, M. Deshpande, S. Nair, C. Natassia, G. Yordanova, D. Yuan, O. Stroe, G. Wood, A. Laydon, A. Židek, T. Green, K. Tunyasuvunakool, S. Petersen, J. Jumper, E. Clancy, R. Green, A. Vora, M. Lutfi, M. Figurnov, A. Cowie, N. Hobbs, P. Kohli, G. Kleywegt, E. Birney, D. Hassabis, and S. Velankar. AlphaFold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res.*, 50(D1):D439–D444, 2022.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *Adv. Neural Inf. Process. Syst.*, 30:6000–6010, 2017.
- A. Whited and A. Johs. The interactions of peripheral membrane proteins with biological membranes. *Chem. Phys. Lipids*, 192:51–59, 2015.
- W. Wimley. Describing the mechanism of antimicrobial peptide action with the interfacial activity model. *ACS Chem. Biol.*, 5(10):905–917, 2010.
- Y. Xu and B. Mitra. A highly active, soluble mutant of the membrane-associated (s)-mandelate dehydrogenase from *Pseudomonas putida*. *Biochemistry*, 38(38):12367–12376, 1999.